ARX® CLI Reference

810-0025-00

Publication Date

This manual was published on May 13, 2013.

Legal Notices

Copyright

Copyright 2004-5/13/13, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

Access Policy Manager, Advanced Client Authentication, Advanced Routing, APM, Application Security Manager, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Management Pack, F5 Networks, F5 World, Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, IBR, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iApps, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, Local Traffic Manager, LTM, Message Security Manager, MSM, OneConnect, OpenBloX, OpenBloX [DESIGN], Packet Velocity, Policy Enforcement Manager, PEM, Protocol Security Manager, PSM, Real Traffic Policy Builder, Rosetta Diameter Gateway, ScaleN, Signaling Delivery Controller, SDC, SSL Acceleration, StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TMOS, Traffic Management Operating System, Traffix Diameter Load Balancer, Traffix Systems, Traffix Systems (DESIGN), Transparent Data Reduction, UNITY, VAULT, VIPRION, vCMP, virtual Clustered Multiprocessing, WA, WAN Optimization Manager, WebAccelerator, WOM, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by U.S. Patents 7,877,511; 7,958,347. This list is believed to be current as of May 13, 2013.

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Acknowledgments

This product includes software from several third-party vendors. Each vendor is listed below with the applicable copyright.

Copyright (c) 1990, 1993, 1994, 1995 The Regents of the University of California. All rights reserved.

Copyright 2000 by the Massachusetts Institute of Technology. All Rights Reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

Copyright 1993 by OpenVision Technologies, Inc.

Copyright (C) 1998 by the FundsXpress, INC.

All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

Copyright (c) 1995-2001 International Business Machines Corporation and others

All rights reserved.

Copyright (c) 1990-2003 Sleepycat Software. All rights reserved.

Copyright (c) 1995, 1996 The President and Fellows of Harvard University. All rights reserved.

Copyright (c) 1998-2004 The OpenSSL Project. All rights reserved.

Unless otherwise noted, the companies, organizations, products, domain names, email addresses, logos, people, places, and events depicted in examples herein are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

Revision History

October 2005 - Rev A

March 2006 - Rev B, updates for Software Release 2.1.4

April 2006 - Rev C, Software Release 2.3

August 2006 - Rev D, updates for Software Release 2.4

September 2006 - Rev E, updates for Software Release 2.4.1

October 2006 - Rev F, updates for Software Release 2.4.2

January 2007 - Rev G, updates for Software Release 2.4.3

March 2007 - Rev H, updates for Software Release 2.5.0

May 2007 - Rev J, updates for Software Release 2.5.1

August 2007 - Rev K, add forest-to-forest trusts for Software Release 2.6.0

November 2007 - Rev L, minor output changes for Software Release 2.7.0

February 2008 - Rev M, minor show-policy changes for Software Release 2.7.1

March 2008 - Rev N, convert to F5 format; updates for Software Release 3.1.0

June 2008 - Rev P, updates for Software Release 3.2.0

June 2008 - Rev Q, updates for Software Release 4.0.0

July 2008 - Rev R, formatting change and minor updates for Software Release 4.0.1

September 2008 - Rev S, minor updates for Software Release 3.2.1

October 2008 - Rev T, CIFS features for Software Release 4.1.0

June 2009 - Rev U, updates for Software Release 5.00.000

November 2009 - Rev V, updates for Software Release 5.01.000

January 2010 - Rev W, changes for Release 5.01.000 update
March 2010 - Rev X, updates for Software Release 5.01.005
August 2010 - Rev Y, updates for Software Release 5.01.007
October 2010 - Rev Z, updates for Software Release 5.02.000
January 2011 - Rev AA, updates for Software Release 5.03.000
April 2011 - Rev AB, updates for Software Release 6.00.000
September 2011 - Rev AC, updates for Software Release 6.01.000
December 2011 - Rev AD, updates for Software Release 6.01.001
January 2012 - Rev AE, minor updates for Software Release 6.01.001, hot fix
April 2012 - Rev AF, updates for Software Release 6.02.000
October 2012 - Rev AG, updates for Software Release 6.03.000
June 2013 - Rev AH, updates for Software Release 6.04.000

Table of Contents

1
Introduction

The ARX 1-3
Back-end Storage and Servers 1-3
Front-end Services 1-3
Policy 1-3
Resilient Overlay Network (RON) 1-4
Audience for this Manual 1-5
CLI Overview 1-6
Exec Mode 1-6
Priv-exec Mode 1-6
Exiting a Mode 1-7
Prompts 1-7
The no Convention 1-7
The enable/no enable Convention 1-7
Getting Started 1-9
Entering Cfg or Gbl Mode 1-9
Document Conventions 1-10
Command Definitions 1-11
Contacting Customer Service 1-12

2
Command Keys and Shortcuts

Keys 2-3

3
Terminal Control

save profile 3-3
show history 3-4
show terminal 3-5
terminal beta 3-7
terminal character-set 3-8
terminal clear 3-9
terminal confirmation 3-10
terminal expand-prompt 3-11
terminal expert 3-12
terminal history 3-13
terminal length 3-14
terminal logging 3-15
terminal stop-on-error 3-16
terminal timeout 3-17
terminal width 3-18

4
Administrative Users

group 4-3
group (gbl-user) 4-5
password 4-6
password (gbl-user) 4-7
role 4-8
ssh-key 4-9
show group all 4-11
show group roles 4-12

show group users 4-13
show ssh-user 4-14
show users 4-16
user 4-17
user (gbl-group) 4-18
windows-domain (gbl-group) 4-19

5
Chassis Management

Slot Locations 5-3
clear metalog usage 5-4
clear nvr 5-5
clock set 5-7
clock timezone 5-8
dual-reboot 5-10
hostname 5-11
login-banner 5-12
probe metalog latency 5-13
reload 5-14
resource-profile 5-16
show baudrate 5-18
show chassis 5-19
show clock 5-35
show hostname 5-36
show memory usage 5-37
show metalog usage 5-39
show processors 5-43
show processors usage 5-47
show uptime 5-52
show version 5-53
shutdown 5-57

6
RAID Management

raid offline 6-3
raid rebuild 6-4
raid rebuild-rate 6-6
raid silence 6-7
raid verification-mode 6-8
raid verification-rate 6-10
raid verify 6-12

7
File Management

at 7-3
clear at 7-5
copy ftp 7-6
copy {nfs|cifs} 7-8
copy ron 7-11
copy scp 7-13
copy smtp 7-16
copy tftp 7-18
delete 7-20

grep 7-22
ip ftp-user 7-26
ip ron-user 7-27
ip scp-user 7-28
move 7-29
move ... ftp 7-31
move ... {nfs|cifs} 7-33
move ... scp 7-35
move ... tftp 7-37
pause 7-39
remark 7-40
rename 7-41
run 7-42
show at 7-43
show capture 7-44
show configs 7-48
show cores 7-50
show diag-info 7-52
show directories 7-53
show license 7-57
show logs 7-59
show releases 7-60
show replicated-configs 7-62
show reports 7-64
show reports status 7-70
show scripts 7-71
show software 7-73
show stats-logs 7-75
tail 7-79
truncate-report 7-80
wait-for report 7-81

8
Master Key

show master-key 8-3

9
Layer 2

blocked-vlan (cfg-vlan) 9-3
channel 9-4
clear counters channel 9-6
clear counters gigabit 9-7
clear counters lacp 9-8
clear counters ten-gigabit 9-9
clear counters redundancy network 9-10
description (cfg-channel) 9-11
description (cfg-if-gig) 9-12
description (cfg-if-ten-gig) 9-13
description (cfg-vlan) 9-14
forward-delay 9-15
flowcontrol 9-16
hello-time 9-17
interface gigabit 9-18
interface ten-gigabit 9-19

ip private vlan 9-20
jumbo mtu 9-22
lacp active 9-23
lacp passive 9-25
lacp rate 9-27
load-balance 9-29
mac-address aging-time 9-31
max-age 9-32
members (cfg-channel) 9-33
members (cfg-vlan) 9-35
priority (cfg-channel) 9-37
priority (cfg-stp) 9-38
protocol (cfg-stp) 9-39
redundancy protocol 9-40
redundancy protocol (cfg-channel) 9-42
show channel 9-44
show channel ... stats 9-51
show interface gigabit 9-54
show interface ten-gigabit 9-57
show load-balancing 9-59
show mac-address-table 9-60
show mac-address-table summary 9-62
show redundancy network 9-63
show spanning-tree detailed 9-65
show spanning-tree interface 9-66
show spanning-tree summary 9-67
show vlan 9-68
show vlan summary 9-70
shutdown (cfg-channel) 9-71
shutdown (cfg-if-gig) 9-72
shutdown (cfg-if-ten-gig) 9-73
shutdown (cfg-stp) 9-74
spanning-tree 9-75
spanning-tree cost 9-77
spanning-tree edgeport 9-78
spanning-tree priority 9-79
spanning-tree shutdown 9-80
speed (cfg-if-gig) 9-81
switch-forwarding enable 9-82
tag 9-83
trap shutdown 9-85
vlan 9-86
vlan (cfg-channel) 9-88
vlan-tag 9-89

10
Layer 3 (Network Layer)

arp 10-3
arp gratuitous 10-4
clear arp 10-5
clear counters mgmt 10-6
description (cfg-if-vlan) 10-7
description (cfg-mgmt) 10-8
interface mgmt 10-9
interface vlan 10-11

ip address (cfg-if-vlan) 10-12
ip address (cfg-mgmt) 10-14
ip domain-list 10-15
ip name-server 10-16
ip proxy-address 10-17
ip route 10-19
ip route ... per-vlan 10-21
ip route ... source-ip 10-23
ntp server 10-25
redundancy (cfg-if-vlan) 10-26
show arp 10-28
show interface 10-34
show interface mgmt 10-42
show interface vlan 10-44
show ip address 10-45
show ip domain 10-46
show ip proxy-addresses 10-47
show ip route 10-49
show ntp servers 10-55
show ntp status 10-56
shutdown (cfg-if-vlan) 10-58
shutdown (cfg-mgmt) 10-60
speed (cfg-mgmt) 10-61
wait-for ip-routes 10-62

11
Licensing

clear active-license 11-3
license activate 11-4
license activate file 11-6
license create license-dossier 11-8
ping license-server 11-10
show active-license 11-11
show license-dossier 11-14

12
SNMP

show snmp-server 12-3
snmp-server community 12-5
snmp-server contact 12-6
snmp-server host 12-7
snmp-server location 12-8
snmp-server name 12-9
snmp-server traps 12-10
snmp-server trusthost 12-11

13
Email Notifications (SMTP)

clear smtp queue 13-3
description (cfg-email-event) 13-4
email-event 13-5
email-severity 13-7
enable (cfg-email-event) 13-8

from (cfg-smtp) 13-9
group (cfg-email-event) 13-10
group ... event 13-12
mail-server 13-14
mail-to (cfg-email-event) 13-15
maximum age 13-16
retry interval 13-17
show email-event 13-18
show email-severity 13-21
show smtp queue 13-26
show smtp status 13-28
show smtp welcome 13-31
smtp 13-32
smtp retry 13-34
smtp test email-event 13-35
smtp test message 13-37
smtp test server 13-38
smtp welcome 13-39
to 13-41

14
RON

heartbeat failure 14-3
heartbeat interval 14-4
interface ron 14-5
ip address (cfg-ron) 14-6
ip private subnet reassign 14-7
peer address 14-8
rconsole 14-9
ron evict 14-10
ron tunnel 14-11
show ron 14-12
show ron conflicts 14-14
show ron database 14-15
show ron route 14-18
show ron tunnel 14-19
shutdown (cfg-if-vlan-ron-tnl) 14-23

15
Redundant Pairs (HA)

clear counters redundancy 15-3
critical route 15-5
enable (cfg-redundancy) 15-7
nsm binary-core-files 15-9
nsm recovery 15-10
nsm warm-restart 15-11
peer 15-13
quorum-disk 15-14
redundancy 15-16
redundancy force-active 15-19
resilver-timeout 15-21
show nsm 15-22
show nsm warm-restart history 15-24
show redundancy 15-27

show redundancy all 15-28
show redundancy ballots 15-31
show redundancy critical-services 15-33
show redundancy history 15-35
show redundancy license 15-36
show redundancy metalog 15-37
show redundancy peer 15-40
show redundancy quorum-disk 15-41
show redundancy reboot-history 15-44
show redundancy resilver-timeout 15-45
suspend-failover 15-46

16
Active Directory Discovery

active-directory update seed-domain 16-3
description (gbl-proxy-user) 16-8
proxy-user 16-9
show active-directory 16-11
show active-directory status 16-15
show proxy-user 16-20
user (gbl-proxy-user) 16-22
windows-domain (gbl-proxy-user) 16-24

17
Management Access

authentication 17-3
auth-port (gbl-radius) 17-6
clear session 17-7
clear statistics authentication 17-8
key (gbl-radius) 17-9
management access 17-10
permit 17-12
radius-server 17-13
retries (gbl-radius) 17-14
show management access 17-15
show radius-server 17-16
show sessions 17-17
show ssh-host-key 17-19
show statistics authentication 17-21
ssh-host-key 17-22
ssh-host-key generate 17-24
ssh-v1 enable 17-25
timeout (gbl-radius) 17-26

18
CIFS Authentication

active-directory-forest 18-3
active-directory forest-trust 18-5
active-directory update forest 18-7
child-domain 18-11
forest-root 18-13
ip address (gbl-ntlm-auth-srv) 18-15
kerberos auto-realm-traversal 18-17

kerberos health-check threshold 18-18
name-server 18-19
ntlm-auth-db 18-20
ntlm-auth-server 18-22
password 18-24
permit (gbl-mgmt-auth) 18-25
port 18-27
show ntlm-auth-db 18-28
show ntlm-auth-server 18-29
show windows-mgmt-auth 18-34
tree-domain 18-36
user (gbl-mgmt-auth) 18-38
user (gbl-ntlm-auth-db) 18-40
windows-domain (gbl-ntlm-auth-srv) 18-41
windows-mgmt-auth 18-43

19
NFS Access Lists

anonymous-gid 19-3
anonymous-uid 19-4
deny 19-5
description (gbl-nfs-acl) 19-6
ip address (gbl-nis-dom) 19-7
nfs-access-list 19-8
nis domain 19-10
nis domain (gbl-nfs-acl) 19-11
nis update 19-12
permit (gbl-nfs-acl) 19-14
show nfs-access-list 19-17
show nis domain 19-20
show nis netgroup 19-22

20
External Filer

cifs connection-limit 20-3
cifs-port 20-4
description (gbl-filer) 20-5
external-filer 20-6
ext-filer-ip-addrs activate 20-8
filer-type 20-9
ignore-name 20-12
ip address 20-14
ip address ... change-to 20-16
manage snapshots 20-18
nfs tcp connections 20-19
probe authentication 20-20
probe exports 20-22
proxy-user (gbl-filer) 20-25
show exports 20-27
show external-filer 20-34
spn 20-47

21

Namespace

character-encoding nfs 21-3
cifs anonymous-access 21-5
cifs authentication 21-6
cifs filer-signatures 21-8
description (gbl-ns, gbl-ns-...) 21-10
enable (gbl-ns, gbl-ns-vol) 21-11
metadata cache-size 21-14
namespace 21-16
ntlm-auth-db (gbl-ns) 21-18
ntlm-auth-server (gbl-ns) 21-19
protocol 21-21
proxy-user (gbl-ns) 21-22
remove namespace 21-23
sam-reference 21-25
show namespace 21-27
show namespace mapping 21-68
show namespace status 21-71
smb2-allowed 21-77
windows-mgmt-auth (gbl-ns) 21-78

22
Volume

attach 22-3
auto reserve files 22-4
auto sync files 22-5
cancel import 22-7
cifs access-based-enum 22-8
cifs access-based-enum (priv-exec) 22-11
cifs access-based-enum exclude 22-14
cifs case-sensitive 22-16
cifs deny-symlinks 22-18
cifs file-system-name 22-19
cifs notify-change-mode 22-21
cifs oplocks-disable 22-22
cifs path-cache 22-23
compressed-files 22-24
critical 22-25
direct 22-26
enable (gbl-ns-vol-shr) 22-27
filer 22-31
filer-subshares 22-34
freespace adjust 22-41
freespace apparent-size 22-43
freespace calculation dir-master-only 22-45
freespace calculation manual 22-47
freespace cifs-quota 22-49
freespace ignore 22-52
ignore-sid-errors 22-53
import priority 22-54
import rename-directories 22-56
import rename-files 22-59
import skip-managed-check 22-61
import sync-attributes 22-62
managed-volume 22-64

max-volume-groups 22-65
metadata critical 22-66
metadata share 22-67
modify 22-69
named-streams 22-71
nfs-param 22-73
persistent-acls 22-74
reimport-modify 22-76
replica-snap 22-77
reserve files 22-79
restart namespace ... volume 22-80
share 22-81
show share status 22-84
show sid-translation 22-86
show volume-group 22-87
sid-translation 22-92
sparse-files 22-93
unicode-on-disk 22-94
volume 22-96
volume-group 22-98
wait-for shares-online 22-100
wait-for volume-disable 22-101
wait-for volume-enable 22-102

23
Global Server

active-directory alias 23-3
active-directory proxy-user 23-5
description (gbl-gs) 23-6
enable (gbl-gs, gbl-gs-vs) 23-7
global server 23-8
show global server 23-10
show server-mapping 23-13
virtual server 23-21
wait-for vip-disable 23-23
wait-for vip-enable 23-24
windows-domain (gbl-gs) 23-25
wins 23-26
wins-alias 23-27
wins-name 23-29

24
Front-End Services

browsing 24-3
cifs 24-5
cifs rekey 24-8
clear dynamic-dns 24-10
clear nlm locks 24-11
description (gbl-cifs) 24-12
description (gbl-nfs) 24-13
domain-join 24-14
dynamic-dns 24-19
dynamic-dns update 24-21
enable (gbl-cifs, gbl-nfs) 24-22

export (gbl-cifs) 24-23
export (gbl-nfs) 24-27
export offline-access 24-29
nfs 24-31
nfs tcp timeout 24-33
nlm enable 24-34
offline-behavior 24-36
probe delegate-to 24-40
remove namespace ... volume ... exports-only 24-44
show cifs-service 24-45
show cifs-service filer-capabilities 24-54
show cifs-service subshares 24-58
show dynamic-dns 24-62
show global service 24-64
show nfs-service 24-65
show nfs-service mounts 24-70
show nfs tcp 24-72
show nlm client 24-73
show nlm file 24-75
show nlm statistics 24-77
show virtual service 24-79
signatures 24-81
smb2-allowed (gbl-cifs) 24-82
sync cifs delegation 24-83
wins-name-encoding 24-84

25
Schedules

description (gbl-schedule) 25-3
duration 25-4
every 25-5
schedule 25-7
show schedule 25-8
start 25-10
stop 25-11

26
Filesets

age-fileset 26-3
every (...-fs-age) 26-4
filename-fileset 26-6
filesize-fileset 26-7
from fileset (gbl-...-fs-...) 26-8
intersection-fileset 26-9
last 26-10
name 26-11
offline 26-13
path 26-14
policy-age-fileset 26-17
policy-cifs-attributes-fileset 26-18
policy-filename-fileset 26-20
policy-filesize-fileset 26-21
policy-intersection-fileset 26-22
policy-union-fileset 26-23

recurse 26-24
select-files (...-fs-filesize) 26-25
select-files (...-fs-age) 26-26
show policy filesets 26-28
start 26-32
union-fileset 26-33

27
Place Rules

enable (gbl-ns-vol-plc) 27-3
from (gbl-ns-vol-plc) 27-4
inline notify 27-6
inline report 27-7
limit-migrate 27-10
migrate close-file 27-11
migrate hard-links 27-12
migrate retain-files 27-13
place-rule 27-14
policy freespace 27-18
policy freespace (gbl-ns, gbl-ns-vol) 27-20
policy migrate-attempts 27-22
policy migrate-delay 27-23
policy migrate-method 27-24
policy migrate-retry-delay 27-26
policy order-rule 27-27
policy pause 27-29
policy pause (gbl-ns-vol) 27-30
policy treewalk-threads 27-31
remove namespace ... policy-only 27-32
report (gbl-ns-vol-plc) 27-34
schedule (gbl-ns-vol-plc) 27-36
show policy 27-37
show policy files-closed 27-48
source 27-50
target 27-51
tentative 27-52
volume-scan 27-53
wait-for migration 27-54

28
Share Farms

auto-migrate 28-3
balance 28-4
constrain-directories 28-6
constrain-files 28-8
enable (gbl-ns-vol-sfarm) 28-10
policy freespace (gbl-ns-vol-sfarm) 28-11
share (gbl-ns-vol-sfarm) 28-13
share-farm 28-14

29
Shadow Volume

bandwidth-limit 29-3

cifs-8dot3-resolution 29-4
database-location 29-5
delta-threshold 29-6
enable (gbl-ns-vol-shdwcp) 29-7
from fileset (gbl-ns-vol-shdwcp) 29-8
inline-notify 29-9
prune-target 29-10
publish 29-11
report (gbl-ns-vol-shdwcp) 29-12
retry attempts 29-16
retry delay 29-17
schedule (gbl-ns-vol-shdwcp) 29-18
shadow 29-19
shadow-copy-rule 29-21
show shadow 29-24
sid-translation (gbl-ns-vol-shdwcp) 29-29
target (gbl-ns-vol-shdwcp) 29-31

30
Snapshots

clear statistics snapshot 30-3
enable (gbl-ns-vol-...snap) 30-4
exclude 30-5
report (gbl-ns-vol-...snap) 30-7
retain 30-11
schedule (gbl-ns-vol-...snap) 30-12
show snapshots 30-14
snapshot clear 30-22
snapshot consistency 30-24
snapshot create 30-25
snapshot directory cifs-name 30-27
snapshot directory display 30-29
snapshot directory nfs-name 30-31
snapshot manage 30-32
snapshot privileged-access 30-34
snapshot remove 30-36
snapshot replica-snap-rule 30-39
snapshot rule 30-46
snapshot verify 30-53
snapshot vss-mode 30-56
wait-for snapshot create 30-58
wait-for snapshot remove 30-59
wait-for snapshot verify 30-60

31
File Tracking

archive 31-3
cancel snapshot archive 31-4
clear file-history archive 31-5
contents 31-7
description (gbl-archive) 31-9
file-history archive 31-10
find 31-12
location 31-16

show file-history archive 31-18
show file-history archive ... contents 31-21
show file-history virtual-service 31-25
show virtual path-history 31-31

32
Restore Data

cancel restore data 32-3
clear restore data 32-4
restore data 32-6
show restore data 32-11
wait-for restore data 32-14

33
Running and Global-Configs

clear global-config 33-3
copy global-config 33-4
copy running-config 33-8
copy startup-config 33-12
delete startup-config 33-16
restore startup-config 33-17
save boot-config 33-18
show global-config 33-19
show running-config 33-43

34
ARX API and Notification Rules

clear statistics api 34-3
clear statistics notification 34-4
enable (gbl-ns-vol-ntfy) 34-5
notification rule 34-6
report (gbl-ns-vol-ntfy) 34-11
retain (gbl-ns-vol-ntfy) 34-13
schedule (gbl-ns-vol-ntfy) 34-14
show notification 34-15
show statistics api 34-21

35
Disaster Recovery Between ARX Clusters

activate configs 35-3
cluster-name 35-13
config-replication 35-15
description (gbl-cfg-repl) 35-18
enable (gbl-cfg-repl) 35-19
load configs 35-20
remove cluster-config 35-31
report (gbl-cfg-repl) 35-32
schedule (gbl-cfg-repl) 35-34
show cluster 35-35
show config-replication 35-36
target-cluster 35-39
target-file 35-40
user (gbl-cfg-repl) 35-41

36
Software Upgrades

boot system 36-3
firmware upgrade 36-6
show boot 36-8
show chassis software 36-9
show firmware upgrade 36-10

37
Namespace Check (nsck) and Sync

cancel migrate-metadata 37-3
cancel migrate-volume 37-4
cancel nsck report 37-5
cancel sync 37-6
clear nsck 37-7
clear sync 37-8
nsck ... destage 37-10
nsck ... migrate-metadata 37-12
nsck ... migrate-volume 37-14
nsck ... rebuild 37-16
nsck ... report dir-structure 37-18
nsck ... report inconsistencies 37-22
nsck ... report metadata-only 37-27
nsck ... report symlinks 37-30
nsck ... sync directories 37-33
nsck ... sync files 37-34
show nsck 37-36
show sync 37-40
sync directories 37-43
sync files 37-46
sync subshares from-namespace 37-51
sync subshares from-service 37-55
wait-for nsck 37-59
wait-for sync 37-60

38
Basic Troubleshooting Tools

Log Components 38-3
additional-command 38-13
auto-diagnostics 38-15
auto-diagnostics test 38-17
clear health 38-18
collect 38-19
collect logs 38-23
expect monitor 38-25
logging destination 38-26
logging level 38-27
mail-to (gbl-auto-diag) 38-29
management source 38-30
moving-average 38-31
notify 38-33
sampling interval 38-35
schedule (gbl-auto-diag) 38-36
show auto-diagnostics 38-37

show documentation 38-39
show health 38-40
show health time-skew 38-42
show id-mappings 38-44
show logging destination 38-46
show logging levels 38-47
show stats-monitor 38-49
show system tasks 38-51
stats-monitor 38-53
trap 38-54

39
Network Troubleshooting Tools

Network-Log Components 39-3
capture merge 39-5
capture session 39-6
clear statistics filer connections 39-8
clear statistics global server 39-9
drop filer-connections 39-10
expect nslookup 39-12
expect show firewall 39-13
expect show ifconfig 39-16
expect show netstat 39-19
expect traceroute 39-20
expect ttcp 39-21
logging fastpath component 39-23
logging fastpath component ... filter 39-25
logging fastpath processor 39-27
monitor 39-29
ping 39-31
show capture sessions 39-34
show fastpath logging 39-36
show fastpath resources 39-38
show filer connections 39-43
show monitor 39-45
show statistics filer connections 39-46
show statistics global server 39-51

40
Managed-Volume Troubleshooting Tools

cancel migration 40-3
cancel remove 40-4
clear statistics filer 40-5
clear statistics metadata 40-6
clear statistics metalog 40-7
clear statistics migration 40-8
export-mapping 40-9
find 40-12
remove service 40-15
remove-share migrate 40-17
remove-share nomigrate 40-19
remove-share offline 40-21
show policy history 40-23
show policy queue 40-25

show statistics filer 40-27
show statistics metadata 40-37
show statistics metalog 40-41
show statistics migration 40-45
show statistics namespace ... fastpath 40-49
show statistics namespace ... request-detail 40-52
show statistics namespace ... response-detail 40-61
show statistics namespace ... summary 40-67
strict-attribute-consistency 40-72
wait-for remove 40-74

41
CIFS-Service Troubleshooting Tools

cifs promote-subshares 41-3
clear statistics cifs authentication 41-6
clear statistics cifs path-cache 41-7
clear statistics cifs symlinks 41-8
clear statistics cifs work-queues 41-9
clear statistics domain-controller 41-10
clear statistics domain-controller load-balancing 41-11
clear subshare-cache 41-12
close cifs file 41-13
drop cifs-service user-session 41-14
show cifs-service client-activity 41-15
show cifs-service exports 41-21
show cifs-service kerberos-tickets 41-24
show cifs-service open-files 41-27
show cifs-service path-cache 41-30
show cifs-service transactions 41-33
show cifs-service user-sessions 41-34
show fastpath cifs-signatures 41-38
show statistics cifs authentication 41-42
show statistics cifs fastpath 41-47
show statistics cifs path-cache 41-49
show statistics cifs symlinks 41-53
show statistics cifs work-queues 41-57
show statistics domain-controller 41-65
show statistics domain-controller load-balancing 41-69
show subshare-cache 41-72

42
ARX Manager (GUI) Maintenance

cipher 42-3
gui 42-4
ssl 42-5
ssl-key-store 42-6

Alphabetical Command List

activate configs 35-3
active-directory alias 23-3
active-directory forest-trust 18-5
active-directory proxy-user 23-5
active-directory update forest 18-7
active-directory update seed-domain 16-3
active-directory-forest 18-3
additional-command 38-13
age-fileset 26-3
anonymous-gid 19-3
anonymous-uid 19-4
archive 31-3
arp 10-3
arp gratuitous 10-4
at 7-3
attach 22-3
Audience for this Manual 1-5
authentication 17-3
auth-port (gbl-radius) 17-6
auto reserve files 22-4
auto sync files 22-5
auto-diagnostics 38-15
auto-diagnostics test 38-17
auto-migrate 28-3
balance 28-4
bandwidth-limit 29-3
blocked-vlan (cfg-vlan) 9-3
boot system 36-3
browsing 24-3
cancel import 22-7
cancel migrate-metadata 37-3
cancel migrate-volume 37-4
cancel migration 40-3
cancel nsck report 37-5
cancel remove 40-4
cancel restore data 32-3
cancel snapshot archive 31-4
cancel sync 37-6
capture merge 39-5
capture session 39-6
channel 9-4
character-encoding nfs 21-3
child-domain 18-11
cifs 24-5
cifs access-based-enum 22-8
cifs access-based-enum (priv-exec) 22-11
cifs access-based-enum exclude 22-14
cifs anonymous-access 21-5
cifs authentication 21-6
cifs case-sensitive 22-16
cifs connection-limit 20-3
cifs deny-symlinks 22-18
cifs filer-signatures 21-8
cifs file-system-name 22-19

cifs notify-change-mode 22-21
cifs oplocks-disable 22-22
cifs path-cache 22-23
cifs promote-subshares 41-3
cifs rekey 24-8
cifs-8dot3-resolution 29-4
cifs-port 20-4
cipher 42-3
clear active-license 11-3
clear arp 10-5
clear at 7-5
clear counters channel 9-6
clear counters gigabit 9-7
clear counters lacp 9-8
clear counters mgmt 10-6
clear counters redundancy 15-3
clear counters redundancy network 9-10
clear counters ten-gigabit 9-9
clear dynamic-dns 24-10
clear file-history archive 31-5
clear global-config 33-3
clear health 38-18
clear metalog usage 5-4
clear nlm locks 24-11
clear nsck 37-7
clear nvr 5-5
clear restore data 32-4
clear session 17-7
clear smtp queue 13-3
clear statistics api 34-3
clear statistics authentication 17-8
clear statistics cifs authentication 41-6
clear statistics cifs path-cache 41-7
clear statistics cifs symlinks 41-8
clear statistics cifs work-queues 41-9
clear statistics domain-controller 41-10
clear statistics domain-controller load-balancing 41-11
clear statistics filer 40-5
clear statistics filer connections 39-8
clear statistics global server 39-9
clear statistics metadata 40-6
clear statistics metalog 40-7
clear statistics migration 40-8
clear statistics notification 34-4
clear statistics snapshot 30-3
clear subshare-cache 41-12
clear sync 37-8
CLI Overview 1-6
clock set 5-7
clock timezone 5-8
close cifs file 41-13
cluster-name 35-13
collect 38-19
collect logs 38-23
Command Definitions 1-11
compressed-files 22-24

config-replication 35-15
constrain-directories 28-6
constrain-files 28-8
Contacting Customer Service 1-12
contents 31-7
copy {nfs|cifs} 7-8
copy ftp 7-6
copy global-config 33-4
copy ron 7-11
copy running-config 33-8
copy scp 7-13
copy smtp 7-16
copy startup-config 33-12
copy tftp 7-18
critical 22-25
critical route 15-5
database-location 29-5
delete 7-20
delete startup-config 33-16
delta-threshold 29-6
deny 19-5
description (cfg-channel) 9-11
description (cfg-email-event) 13-4
description (cfg-if-gig) 9-12
description (cfg-if-ten-gig) 9-13
description (cfg-if-vlan) 10-7
description (cfg-mgmt) 10-8
description (cfg-vlan) 9-14
description (gbl-archive) 31-9
description (gbl-cfg-repl) 35-18
description (gbl-cifs) 24-12
description (gbl-filer) 20-5
description (gbl-gs) 23-6
description (gbl-nfs) 24-13
description (gbl-nfs-acl) 19-6
description (gbl-ns, gbl-ns-...) 21-10
description (gbl-proxy-user) 16-8
description (gbl-schedule) 25-3
direct 22-26
Document Conventions 1-10
domain-join 24-14
drop cifs-service user-session 41-14
drop filer-connections 39-10
dual-reboot 5-10
duration 25-4
dynamic-dns 24-19
dynamic-dns update 24-21
email-event 13-5
email-severity 13-7
enable (cfg-email-event) 13-8
enable (cfg-redundancy) 15-7
enable (gbl-cfg-repl) 35-19
enable (gbl-cifs, gbl-nfs) 24-22
enable (gbl-gs, gbl-gs-vs) 23-7
enable (gbl-ns, gbl-ns-vol) 21-11
enable (gbl-ns-vol-...snap) 30-4

enable (gbl-ns-vol-ntfy) 34-5
enable (gbl-ns-vol-plc) 27-3
enable (gbl-ns-vol-sfarm) 28-10
enable (gbl-ns-vol-shdwcp) 29-7
enable (gbl-ns-vol-shr) 22-27
every 25-5
every (...-fs-age) 26-4
exclude 30-5
expect monitor 38-25
expect nslookup 39-12
expect show firewall 39-13
expect show ifconfig 39-16
expect show netstat 39-19
expect traceroute 39-20
expect ttcp 39-21
export (gbl-cifs) 24-23
export (gbl-nfs) 24-27
export offline-access 24-29
export-mapping 40-9
external-filer 20-6
ext-filer-ip-addrs activate 20-8
file-history archive 31-10
filename-fileset 26-6
filer 22-31
filer-subshares 22-34
filer-type 20-9
filesize-fileset 26-7
find 31-12
find 40-12
firmware upgrade 36-6
flowcontrol 9-16
forest-root 18-13
forward-delay 9-15
freespace adjust 22-41
freespace apparent-size 22-43
freespace calculation dir-master-only 22-45
freespace calculation manual 22-47
freespace cifs-quota 22-49
freespace ignore 22-52
from (cfg-smtp) 13-9
from (gbl-ns-vol-plc) 27-4
from fileset (gbl-...-fs-...) 26-8
from fileset (gbl-ns-vol-shdwcp) 29-8
Getting Started 1-9
global server 23-8
grep 7-22
group 4-3
group ... event 13-12
group (cfg-email-event) 13-10
group (gbl-user) 4-5
gui 42-4
heartbeat failure 14-3
heartbeat interval 14-4
hello-time 9-17
hostname 5-11
ignore-name 20-12

ignore-sid-errors 22-53
import priority 22-54
import rename-directories 22-56
import rename-files 22-59
import skip-managed-check 22-61
import sync-attributes 22-62
inline notify 27-6
inline report 27-7
inline-notify 29-9
interface gigabit 9-18
interface mgmt 10-9
interface ron 14-5
interface ten-gigabit 9-19
interface vlan 10-11
intersection-fileset 26-9
ip address 20-14
ip address ... change-to 20-16
ip address (cfg-if-vlan) 10-12
ip address (cfg-mgmt) 10-14
ip address (cfg-ron) 14-6
ip address (gbl-nis-dom) 19-7
ip address (gbl-ntlm-auth-srv) 18-15
ip domain-list 10-15
ip ftp-user 7-26
ip name-server 10-16
ip private subnet reassign 14-7
ip private vlan 9-20
ip proxy-address 10-17
ip ron-user 7-27
ip route 10-19
ip route ... per-vlan 10-21
ip route ... source-ip 10-23
ip scp-user 7-28
jumbo mtu 9-22
kerberos auto-realm-traversal 18-17
kerberos health-check threshold 18-18
key (gbl-radius) 17-9
Keys 2-3
lacp active 9-23
lacp passive 9-25
lacp rate 9-27
last 26-10
license activate 11-4
license activate file 11-6
license create license-dossier 11-8
limit-migrate 27-10
load configs 35-20
load-balance 9-29
location 31-16
Log Components 38-3
logging destination 38-26
logging fastpath component 39-23
logging fastpath component ... filter 39-25
logging fastpath processor 39-27
logging level 38-27
login-banner 5-12

mac-address aging-time 9-31
mail-server 13-14
mail-to (cfg-email-event) 13-15
mail-to (gbl-auto-diag) 38-29
manage snapshots 20-18
managed-volume 22-64
management access 17-10
management source 38-30
max-age 9-32
maximum age 13-16
max-volume-groups 22-65
members (cfg-channel) 9-33
members (cfg-vlan) 9-35
metadata cache-size 21-14
metadata critical 22-66
metadata share 22-67
migrate close-file 27-11
migrate hard-links 27-12
migrate retain-files 27-13
modify 22-69
monitor 39-29
move 7-29
move ... {nfs|cifs} 7-33
move ... ftp 7-31
move ... scp 7-35
move ... tftp 7-37
moving-average 38-31
name 26-11
named-streams 22-71
name-server 18-19
namespace 21-16
Network-Log Components 39-3
nfs 24-31
nfs tcp connections 20-19
nfs tcp timeout 24-33
nfs-access-list 19-8
nfs-param 22-73
nis domain 19-10
nis domain (gbl-nfs-acl) 19-11
nis update 19-12
nlm enable 24-34
notification rule 34-6
notify 38-33
nsck ... destage 37-10
nsck ... migrate-metadata 37-12
nsck ... migrate-volume 37-14
nsck ... rebuild 37-16
nsck ... report dir-structure 37-18
nsck ... report inconsistencies 37-22
nsck ... report metadata-only 37-27
nsck ... report symlinks 37-30
nsck ... sync directories 37-33
nsck ... sync files 37-34
nsm binary-core-files 15-9
nsm recovery 15-10
nsm warm-restart 15-11

ntlm-auth-db 18-20
ntlm-auth-db (gbl-ns) 21-18
ntlm-auth-server 18-22
ntlm-auth-server (gbl-ns) 21-19
ntp server 10-25
offline 26-13
offline-behavior 24-36
password 18-24
password 4-6
password (gbl-user) 4-7
path 26-14
pause 7-39
peer 15-13
peer address 14-8
permit 17-12
permit (gbl-mgmt-auth) 18-25
permit (gbl-nfs-acl) 19-14
persistent-acls 22-74
ping 39-31
ping license-server 11-10
place-rule 27-14
policy freespace 27-18
policy freespace (gbl-ns, gbl-ns-vol) 27-20
policy freespace (gbl-ns-vol-sfarm) 28-11
policy migrate-attempts 27-22
policy migrate-delay 27-23
policy migrate-method 27-24
policy migrate-retry-delay 27-26
policy order-rule 27-27
policy pause 27-29
policy pause (gbl-ns-vol) 27-30
policy treewalk-threads 27-31
policy-age-fileset 26-17
policy-cifs-attributes-fileset 26-18
policy-filename-fileset 26-20
policy-filesize-fileset 26-21
policy-intersection-fileset 26-22
policy-union-fileset 26-23
port 18-27
priority (cfg-channel) 9-37
priority (cfg-stp) 9-38
probe authentication 20-20
probe delegate-to 24-40
probe exports 20-22
probe metalog latency 5-13
protocol 21-21
protocol (cfg-stp) 9-39
proxy-user 16-9
proxy-user (gbl-filer) 20-25
proxy-user (gbl-ns) 21-22
prune-target 29-10
publish 29-11
quorum-disk 15-14
radius-server 17-13
raid offline 6-3
raid rebuild 6-4

raid rebuild-rate 6-6
raid silence 6-7
raid verification-mode 6-8
raid verification-rate 6-10
raid verify 6-12
rconsole 14-9
recurse 26-24
redundancy 15-16
redundancy (cfg-if-vlan) 10-26
redundancy force-active 15-19
redundancy protocol 9-40
redundancy protocol (cfg-channel) 9-42
reimport-modify 22-76
reload 5-14
remark 7-40
remove cluster-config 35-31
remove namespace 21-23
remove namespace ... policy-only 27-32
remove namespace ... volume ... exports-only 24-44
remove service 40-15
remove-share migrate 40-17
remove-share nomigrate 40-19
remove-share offline 40-21
rename 7-41
replica-snap 22-77
report (gbl-cfg-repl) 35-32
report (gbl-ns-vol-...snap) 30-7
report (gbl-ns-vol-ntfy) 34-11
report (gbl-ns-vol-plc) 27-34
report (gbl-ns-vol-shdwcp) 29-12
reserve files 22-79
resilver-timeout 15-21
resource-profile 5-16
restart namespace ... volume 22-80
restore data 32-6
restore startup-config 33-17
retain 30-11
retain (gbl-ns-vol-ntfy) 34-13
retries (gbl-radius) 17-14
retry attempts 29-16
retry delay 29-17
retry interval 13-17
role 4-8
ron evict 14-10
ron tunnel 14-11
run 7-42
sampling interval 38-35
sam-reference 21-25
save boot-config 33-18
save profile 3-3
schedule 25-7
schedule (gbl-auto-diag) 38-36
schedule (gbl-cfg-repl) 35-34
schedule (gbl-ns-vol-...snap) 30-12
schedule (gbl-ns-vol-ntfy) 34-14
schedule (gbl-ns-vol-plc) 27-36

schedule (gbl-ns-vol-shdwcp) 29-18
select-files (...-fs-age) 26-26
select-files (...-fs-filesize) 26-25
shadow 29-19
shadow-copy-rule 29-21
share 22-81
share (gbl-ns-vol-sfarm) 28-13
share-farm 28-14
show active-directory 16-11
show active-directory status 16-15
show active-license 11-11
show arp 10-28
show at 7-43
show auto-diagnostics 38-37
show baudrate 5-18
show boot 36-8
show capture 7-44
show capture sessions 39-34
show channel 9-44
show channel ... stats 9-51
show chassis 5-19
show chassis software 36-9
show cifs-service 24-45
show cifs-service client-activity 41-15
show cifs-service exports 41-21
show cifs-service filer-capabilities 24-54
show cifs-service kerberos-tickets 41-24
show cifs-service open-files 41-27
show cifs-service path-cache 41-30
show cifs-service subshares 24-58
show cifs-service transactions 41-33
show cifs-service user-sessions 41-34
show clock 5-35
show cluster 35-35
show config-replication 35-36
show configs 7-48
show cores 7-50
show diag-info 7-52
show directories 7-53
show documentation 38-39
show dynamic-dns 24-62
show email-event 13-18
show email-severity 13-21
show exports 20-27
show external-filer 20-34
show fastpath cifs-signatures 41-38
show fastpath logging 39-36
show fastpath resources 39-38
show file-history archive 31-18
show file-history archive ... contents 31-21
show file-history virtual-service 31-25
show filer connections 39-43
show firmware upgrade 36-10
show global server 23-10
show global service 24-64
show global-config 33-19

show group all 4-11
show group roles 4-12
show group users 4-13
show health 38-40
show health time-skew 38-42
show history 3-4
show hostname 5-36
show id-mappings 38-44
show interface 10-34
show interface gigabit 9-54
show interface mgmt 10-42
show interface ten-gigabit 9-57
show interface vlan 10-44
show ip address 10-45
show ip domain 10-46
show ip proxy-addresses 10-47
show ip route 10-49
show license 7-57
show license-dossier 11-14
show load-balancing 9-59
show logging destination 38-46
show logging levels 38-47
show logs 7-59
show mac-address-table 9-60
show mac-address-table summary 9-62
show management access 17-15
show master-key 8-3
show memory usage 5-37
show metalog usage 5-39
show monitor 39-45
show namespace 21-27
show namespace mapping 21-68
show namespace status 21-71
show nfs tcp 24-72
show nfs-access-list 19-17
show nfs-service 24-65
show nfs-service mounts 24-70
show nis domain 19-20
show nis netgroup 19-22
show nlm client 24-73
show nlm file 24-75
show nlm statistics 24-77
show notification 34-15
show nsck 37-36
show nsm 15-22
show nsm warm-restart history 15-24
show ntlm-auth-db 18-28
show ntlm-auth-server 18-29
show ntp servers 10-55
show ntp status 10-56
show policy 27-37
show policy files-closed 27-48
show policy filesets 26-28
show policy history 40-23
show policy queue 40-25
show processors 5-43

show processors usage 5-47
show proxy-user 16-20
show radius-server 17-16
show redundancy 15-27
show redundancy all 15-28
show redundancy ballots 15-31
show redundancy critical-services 15-33
show redundancy history 15-35
show redundancy license 15-36
show redundancy metalog 15-37
show redundancy network 9-63
show redundancy peer 15-40
show redundancy quorum-disk 15-41
show redundancy reboot-history 15-44
show redundancy resilver-timeout 15-45
show releases 7-60
show replicated-configs 7-62
show reports 7-64
show reports status 7-70
show restore data 32-11
show ron 14-12
show ron conflicts 14-14
show ron database 14-15
show ron route 14-18
show ron tunnel 14-19
show running-config 33-43
show schedule 25-8
show scripts 7-71
show server-mapping 23-13
show sessions 17-17
show shadow 29-24
show share status 22-84
show sid-translation 22-86
show smtp queue 13-26
show smtp status 13-28
show smtp welcome 13-31
show snapshots 30-14
show snmp-server 12-3
show software 7-73
show spanning-tree detailed 9-65
show spanning-tree interface 9-66
show spanning-tree summary 9-67
show ssh-host-key 17-19
show ssh-user 4-14
show statistics api 34-21
show statistics authentication 17-21
show statistics cifs authentication 41-42
show statistics cifs fastpath 41-47
show statistics cifs path-cache 41-49
show statistics cifs symlinks 41-53
show statistics cifs work-queues 41-57
show statistics domain-controller 41-65
show statistics domain-controller load-balancing 41-69
show statistics filer 40-27
show statistics filer connections 39-46
show statistics global server 39-51

show statistics metadata 40-37
show statistics metalog 40-41
show statistics migration 40-45
show statistics namespace ... fastpath 40-49
show statistics namespace ... request-detail 40-52
show statistics namespace ... response-detail 40-61
show statistics namespace ... summary 40-67
show stats-logs 7-75
show stats-monitor 38-49
show subshare-cache 41-72
show sync 37-40
show system tasks 38-51
show terminal 3-5
show uptime 5-52
show users 4-16
show version 5-53
show virtual path-history 31-31
show virtual service 24-79
show vlan 9-68
show vlan summary 9-70
show volume-group 22-87
show windows-mgmt-auth 18-34
shutdown 5-57
shutdown (cfg-channel) 9-71
shutdown (cfg-if-gig) 9-72
shutdown (cfg-if-ten-gig) 9-73
shutdown (cfg-if-vlan) 10-58
shutdown (cfg-if-vlan-ron-tnl) 14-23
shutdown (cfg-mgmt) 10-60
shutdown (cfg-stp) 9-74
sid-translation 22-92
sid-translation (gbl-ns-vol-shdwcp) 29-29
signatures 24-81
Slot Locations 5-3
smb2-allowed 21-77
smb2-allowed (gbl-cifs) 24-82
smtp 13-32
smtp retry 13-34
smtp test email-event 13-35
smtp test message 13-37
smtp test server 13-38
smtp welcome 13-39
snapshot clear 30-22
snapshot consistency 30-24
snapshot create 30-25
snapshot directory cifs-name 30-27
snapshot directory display 30-29
snapshot directory nfs-name 30-31
snapshot manage 30-32
snapshot privileged-access 30-34
snapshot remove 30-36
snapshot replica-snap-rule 30-39
snapshot rule 30-46
snapshot verify 30-53
snapshot vss-mode 30-56
snmp-server community 12-5

snmp-server contact 12-6
snmp-server host 12-7
snmp-server location 12-8
snmp-server name 12-9
snmp-server traps 12-10
snmp-server trusthost 12-11
source 27-50
spanning-tree 9-75
spanning-tree cost 9-77
spanning-tree edgeport 9-78
spanning-tree priority 9-79
spanning-tree shutdown 9-80
sparse-files 22-93
speed (cfg-if-gig) 9-81
speed (cfg-mgmt) 10-61
spn 20-47
ssh-host-key 17-22
ssh-host-key generate 17-24
ssh-key 4-9
ssh-v1 enable 17-25
ssl 42-5
ssl-key-store 42-6
start 25-10
start 26-32
stats-monitor 38-53
stop 25-11
strict-attribute-consistency 40-72
suspend-failover 15-46
switch-forwarding enable 9-82
sync cifs delegation 24-83
sync directories 37-43
sync files 37-46
sync subshares from-namespace 37-51
sync subshares from-service 37-55
tag 9-83
tail 7-79
target 27-51
target (gbl-ns-vol-shdwcp) 29-31
target-cluster 35-39
target-file 35-40
tentative 27-52
terminal beta 3-7
terminal character-set 3-8
terminal clear 3-9
terminal confirmation 3-10
terminal expand-prompt 3-11
terminal expert 3-12
terminal history 3-13
terminal length 3-14
terminal logging 3-15
terminal stop-on-error 3-16
terminal timeout 3-17
terminal width 3-18
The ARX 1-3
timeout (gbl-radius) 17-26
to 13-41

trap 38-54
trap shutdown 9-85
tree-domain 18-36
truncate-report 7-80
unicode-on-disk 22-94
union-fileset 26-33
user 4-17
user (gbl-cfg-repl) 35-41
user (gbl-group) 4-18
user (gbl-mgmt-auth) 18-38
user (gbl-ntlm-auth-db) 18-40
user (gbl-proxy-user) 16-22
virtual server 23-21
vlan 9-86
vlan (cfg-channel) 9-88
vlan-tag 9-89
volume 22-96
volume-group 22-98
volume-scan 27-53
wait-for ip-routes 10-62
wait-for migration 27-54
wait-for nsck 37-59
wait-for remove 40-74
wait-for report 7-81
wait-for restore data 32-14
wait-for shares-online 22-100
wait-for snapshot create 30-58
wait-for snapshot remove 30-59
wait-for snapshot verify 30-60
wait-for sync 37-60
wait-for vip-disable 23-23
wait-for vip-enable 23-24
wait-for volume-disable 22-101
wait-for volume-enable 22-102
windows-domain (gbl-group) 4-19
windows-domain (gbl-gs) 23-25
windows-domain (gbl-ntlm-auth-srv) 18-41
windows-domain (gbl-proxy-user) 16-24
windows-mgmt-auth 18-43
windows-mgmt-auth (gbl-ns) 21-78
wins 23-26
wins-alias 23-27
wins-name 23-29
wins-name-encoding 24-84

1

Introduction

This manual is a reference for F5's Command Line Interface (CLI). The book is organized as a companion to the CLI user's guides (*ARX® CLI Network-Management Guide*, *ARX® CLI Storage-Management Guide*, and *ARX CLI Maintenance Guide*, in that order); for each chapter in the user guides there is a companion chapter in this book. This manual also has some additional chapters for CLI commands outside the scope of the user's guides.

The ARX

The Adaptive Resource Switch (ARX®) is a highly available and scalable solution that brings resource awareness to a file storage infrastructure, and adapts these resources to meet the demands of users and applications in real time. The ARX provides a file-virtualization layer that aggregates the total capacity and performance of your file storage. A *namespace* provides location-independent, transparent mapping of user requests onto the appropriate storage resource. You can configure policies that the switch enforces for the placement, replication and migration of files. Through policy configuration, the ARX adapts to the real-time demands of users and applications. The ARX thereby serves as a *resource proxy* for the files and services behind it.

Back-end Storage and Servers

The Adaptive Resource Switch aggregates heterogeneous file systems and storage into a unified pool of file storage resources. Through this unification, you can manage these resources to adapt to user demands and client applications. File storage assets can be differentiated based on user-defined attributes, enabling a class-of-storage model. You can reclaim stranded capacity through policy implementation for more effective storage utilization, and you can add capacity without disruption. Back-end resources are monitored for availability and performance, as well as user-access patterns that drive policy decisions.

Front-end Services

The Adaptive Resource Switch acts as an in-band file proxy for the Network File System (NFS) and Microsoft's Common Internet File System (CIFS) protocols.
Front-end services provide the file virtualization layer that masks the physical file storage from the user and application. The switch becomes the file access point, as opposed to the actual physical resource, providing file access through a namespace. Users and applications maintain a single consistent file path that is transparently mapped to the proper physical resource where the information resides.

Policy

The Adaptive Resource Switch provides policy-based resource switching. Through *policy* configuration, you can optimize the placement of files onto the appropriate storage resources and automatically adapt these resources based on user and application demand. The ARX performs file replication and migration based on performance, usage or other life-cycle characteristics, enabling you to implement a flexible file services strategy.

Examples of policies include: migrating files to reclaim stranded capacity; migrating files across different tiers of storage based on access patterns and/or value; and replicating frequently accessed files for performance. The result is more efficient utilization and greater flexibility in file storage management.

Resilient Overlay Network (RON)

You can connect multiple ARXes with a Resilient Overlay Network (RON), which can reside on top of any IP network. This provides a network for distributing and accessing file storage. ARXes can replicate storage to other switches in the same RON, updating the replicas periodically as the writable master files change. This is called a *shadow copy*, where a source volume on one switch periodically copies its files to one or more *shadow volumes* on other switches. Clients can access the shadow volumes at multiple geographic locations, independent to where the source volume resides.

Audience for this Manual

This manual is intended for

- network technicians responsible for layer 1 and 2 networks,
- network engineers responsible for the Internet Protocol (IP) layer (layer 3),
- storage engineers who design and manage storage systems (SANs, NASes, and DASes), and
- crypto officers who manage all of the Critical Security Parameters (CSPs) of a network.

The text presumes that all readers are comfortable with a command-line interface (CLI), especially one based on the Cisco IOS.

CLI Overview

The Command-Line Interface (CLI) has its commands grouped into modes.

Modes are structured as a tree with a single root, *exec* mode. This section summarizes the mode structure and explains some CLI conventions.

Exec Mode

When you log into the CLI, you begin in exec mode. If the hostname is "bstnA," the prompt appears as shown below:

```
bstnA>
```

You can access all global commands (such as show commands) from exec mode, and you can use the enable command to enter priv-exec mode.

```
bstnA> enable
```

Global Commands

You can access global commands from any mode, not just exec. Global commands include all show commands and terminal-control commands.

Priv-exec Mode

Priv-exec mode has the following prompt:

```
bstnA#
```

Priv-exec mode contains chassis-management commands, clock commands, and other commands that require privileges but do not change the network or storage configuration.

Priv-exec has two sub modes, cfg and gbl.

Cfg Mode

To enter cfg mode, use the config command:

```
bstnA# config
bstnA(cfg)#
```

Config mode contains all modes and commands for changing the configuration of the local switch, such as network configuration.

Gbl Mode

To enter gbl mode, use the global command:

```
bstnA# global
bstnA(gbl)#
```

Gbl mode controls all parameters that are shared in a redundant pair, such as namespaces and global servers.

Exiting a Mode

From any mode, use the exit command to return to its parent mode. From priv-exec mode, this command exits the CLI; to go from priv-exec mode back to exec mode, use the no enable command.

From any submode of cfg or gbl mode, you can return immediately to priv-exec mode by using the end command or pressing <Ctrl-z>.

Prompts

Prompts contain information about your position in the mode hierarchy as well as the name of the object you are configuring. For example, suppose you use the following command in gbl mode:

```
bstnA(gbl)# namespace wwmed
bstnA(gbl-ns[wwmed])#
```

This command places you into a new mode, as indicated by the new CLI prompt. The prompt shows the name of the mode, “gbl-ns,” and the name of the configuration object, a namespace called “wwmed.” Abbreviations are used for mode names (for example, “ns” instead of “namespace”) to conserve space on the command line.

When you descend to lower modes in the config tree, the prompt offers more information. To extend the previous example, suppose you enter the following command to configure the “/local” volume in the wwmed namespace:

```
bstnA(gbl-ns[wwmed])# volume /local
bstnA(gbl-ns-vol[wwmed~/local])#
```

The tilde character (~) separates a parent object from its child: “wwmed~/local” shows that you are in the “/local” volume under the “wwmed” namespace.

The no Convention

Most config commands have the option to use the “no” keyword to negate the command. For commands that create an object, the no form removes the object. For commands that change a default setting, the no form reverts back to the default. As an example,

```
bstnA(gbl-ns[wwmed])# no volume /local
```

removes the “/local” volume from the “wwmed” namespace.

The enable/no enable Convention

Many objects and configurations require you to enable them using the enable command before they can take effect. Likewise, many objects and configurations require you to first disable them using the no enable command before you can complete a related command or function. The no

enable command does not remove an object; it only disables it until you re-enable it. The enable/no enable commands exist in many modes and submodes in the CLI.

For example, the following command sequence enables the namespace named “wwmed:”

```
bstnA(gbl)# namespace wwmed
bstnA(gbl-ns[wwmed])# enable
bstnA(gbl-ns[wwmed])# ...
```

Getting Started

For the initial login, refer to the instructions for booting and configuring the switch in the appropriate *Hardware Installation Guide*.

For subsequent logins, use the following steps to log into the F5 CLI:

1. If you are on-site, you can connect a serial line to the serial console port. This port is labeled 'Console' or '10101' (depending on your ARX platform). By default, the port is set for 9600 baud, 8, N, 1.

   You can also telnet to the switch's management interface. For example:

   ```
   telnet 10.10.10.10
   ```

   In either case, a login prompt appears:

```
Username:
```

2. Enter your username and password. For example:

```
Username: admin
Password: acopia
```

   The CLI prompt appears:

```
SWITCH>
```

   The name, "SWITCH," is the default hostname. The hostname is reset as part of the initial-boot process, so it is likely that yours will differ.

Entering Cfg or Gbl Mode

The CLI uses two high-level modes for switch configuration: cfg mode for switch configuration and gbl mode for global configuration. Switch configuration applies to the local switch only; chassis, layer-2, and layer-3 commands are all under cfg mode. Global configuration applies to settings that can be shared by both switches in a redundant pair; for example, namespace-configuration commands all appear under gbl mode.

To enter cfg mode, use the **config** command from priv-exec mode:

```
SWITCH> enable
SWITCH# config
SWITCH(cfg)#
```

To enter gbl mode instead, use the **global** command:

```
SWITCH> enable
SWITCH# global
SWITCH(gbl)#
```

The command sequences in this manual all begin in either cfg or gbl mode.

Document Conventions

This manual uses the following conventions:

`this font` represents screen input and output;

- **`bold text`** represents input, and
- *`italic text`* appears for variable input or output.

this font is used for command-syntax definitions, which use the same rules for bold and italic.

Command-syntax definitions also use the following symbols:

- [*optional-argument*] - square brackets ([ ]) surround optional arguments;
- *choice1* | *choice2* - the vertical bar (|) separates argument choices;
- {*choice1* | *choice2* | *choice3*} - curly braces ({ }) surround a required choice;
- [*choice1* | *choice2*]* - an asterisk (*) means that you can choose none of them, or as many as desired (for example, "choice1 choice2" chooses both);
- {*choice1* | *choice2*}+ - a plus sign (+) means that you must choose one or more.

For commands nested in multiple layers of modes, the following convention illustrates the mode path to reach the command:

```
mode1 -> mode2 -> mode3 -> command
```

Command Definitions

Purpose describes the reason for using the command.

Mode is the mode in the CLI (for example, exec, priv-exec, or cfg).

Security Roles is the set of administrative *roles* that are allowed to access the command. Any number of administrative roles can be assigned to a command; and administrator who has one of these roles can access the command. The roles are

1. operator,
2. backup operator,
3. network-technician,
4. network-engineer,
5. storage-engineer, and
6. crypto-officer.

A command that is accessible to the *operator* is also accessible by an administrator with any other role.

Each administrative group has one or more of these roles, assigned by using the role command.

Syntax describes the syntax of the command. This shows all keywords, arguments, and flags, and provides ranges.

Default(s) lists all default states affected by the command and its arguments.

Valid Platforms lists the platforms (ARX-VE, ARX-500, ARX-1500, ARX-2000, ARX-2500, and/or ARX-4000) where this command is supported. This row is omitted for universally-supported commands.

Guidelines provides context for the command, including references to related commands.

For show commands, this section describes the output of the command.

Samples show sample input and output. For situations where the output is too wide for this table, it appears in a separate figure below.

Related Commands are links to commands that you can use in conjunction with the current command.

Contacting Customer Service

You can use the following methods to contact F5 Networks Customer Service:

F5 Networks Online Knowledge Base Online repository of answers to frequently-asked questions.	http://support.f5.com
F5 Networks Services Support Online Online customer support request system	https://websupport.f5.com
Telephone	Follow this link for a list of Support numbers: http://www.f5.com/support/support-services/contact/

2

Command Keys and Shortcuts

Keys

The CLI supports the following key combinations:

<Arrows>	Move the cursor: <Up-Arrow> and <Down Arrow> go to next/previous line in command history.
<Tab>	Expand the command or keyword abbreviation.
\	Ignore any special meaning for the next character.
<Ctrl-a>	Move cursor to beginning of line.
<Ctrl-b>	Move cursor left one character.
<Ctrl-c>	Abort the current command, or exit cfg/gbl mode.
<Ctrl-d>	Delete character at cursor.
<Ctrl-e>	Move cursor to end of line.
<Ctrl-f>	Move cursor right one character.
<Ctrl-h>	Delete character to left of cursor.
<Ctrl-i>	Expand command or keyword abbreviation (same as <Tab>).
<Ctrl-k>	Delete from cursor to end of line.
<Ctrl-l>	Re-display the current line.
<Ctrl-m>	Same as <Enter> or <Return>.
<Ctrl-n>	Display next line in command history.
<Ctrl-p>	Display previous line in command history.
<Ctrl-r>	Re-display the current line.
<Ctrl-t>	Transpose current character with one to left.
<Ctrl-u>	Delete from cursor to beginning of line.
<Ctrl-w>	Delete word to left of cursor.
<Ctrl-x>	Delete everything on the current line without pressing <Enter>.
<Ctrl-z>	Return to priv mode from cfg or gbl mode.
<Esc-b>	Move cursor left one word.
<Esc-d>	Delete from cursor to end of word.
<Esc-f>	Move cursor right one word.

<Esc-c>	Capitalize character at cursor.
<Esc-m>	Toggle *more* support.
<Esc-l>	Change word at cursor to lower case.
<Esc-s>	Toggle line logging suppression.
<Esc-u>	Change word at cursor to upper case.
<Esc-bs>	Delete word to left of cursor.

3

Terminal Control

save profile

Purpose Use the save profile command to save the current terminal settings for future login sessions.

Modes exec and priv-exec

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `save profile`

Default(s) None

Guidelines This command saves your current terminal settings. The settings reside in a file in the configs directory, admin-profile.scr. The next time you log in to the switch, the CLI automatically starts with the same terminal settings. If you later change the profile and use this command again to save the change, the CLI prompts for confirmation before overwriting the admin-profile.scr file; enter **yes** to proceed.

Use the show terminal command to view the current terminal settings. To remove the profile, run delete configs admin-profile.scr.

Since an administrator with a role of "operator" does not have privileges to access priv-exec mode, exec mode also supports this command. The command operates identically in both modes.

Sample

```
bstnA# save profile
```

saves your current terminal settings.

Related Commands show terminal

show history

Purpose Use the show history command to see the list of CLI commands from the current CLI session.

Mode (any)

Security Role(s) operator

Syntax **`show history`**

Default(s) None

Guidelines You can set the number of commands kept in the history with the terminal history command.

Sample

```
bstnA> show history

Command History
---------------
show terminal
no cli-more
show running-config
show namespace

bstnA>
```

Related Commands terminal history

show terminal

Purpose Use the show terminal command to see the terminal settings for the current CLI session.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show terminal`

Default(s) None

Guidelines To save your current terminal settings, use the save profile command.

If Terminal Idle Timeout is enabled, the system logs you out of the CLI after the idle time shown in the Terminal Idle Timeout Value field. Use terminal timeout to set the idle-timeout time.

Terminal Expand Prompt shows the current object name in the CLI prompt. For example, if you are configuring a namespace volume, the volume name appears in the CLI prompt. Use [no] terminal expand-prompt to disable/enable the expanded prompt.

Terminal Confirmation, if enabled, causes the CLI to prompt for confirmation at sensitive configuration commands. If disabled, the CLI performs all commands without prompting for confirmation. Use [no] terminal confirmation to disable/enable Terminal Confirmation.

Expert Mode, if enabled, stops the CLI from prompting for confirmation when a user creates a new global-configuration object. If disabled, the CLI prompts for confirmation for every new policy, namespace, or global-server object. Use [no] terminal expert to disable/enable Expert Mode.

Terminal Length is the number of lines for each page of the CLI's *more* output. (By default, the CLI pipes all of its output through *more*.) Use terminal length to reset the number of lines; use terminal length 0 to disable the *more* feature.

Terminal Width is the maximum number of characters in each line of CLI output. Use terminal width to reset this maximum.

Logfile Directory only appears if terminal logging is enabled. This shows the directory where the current CLI-log file is being written; currently, this is always "reports."

Terminal Logfile is enabled if the CLI is recording its output to a report file. Use terminal logging to enable CLI logging.

Terminal History is the number of CLI commands kept in a history buffer. Use terminal history to set this number. To access the commands, use show history or <up-arrow>.

Script History Logging indicates whether the CLI keeps terminal history for commands in CLI scripts, too. This is enabled or disabled with the script flag in the terminal history command.

Terminal Character Set is either ISO 8859-1 or Unicode/UTF-8. If the latter, the terminal supports multi-byte characters (for Japanese, Chinese, and other languages with multi-byte characters). Use terminal character-set to change this.

Guidelines (Cont.)

Stop Scripts on error is disabled if CLI scripts continue on error. Use terminal stop-on-error to stop CLI scripts that encounter an error.

If terminal beta is enabled, you can preview and test the CLI beta test commands available in later versions of F5 software. Contact your F5 representative for more information on these commands. Use [no] terminal beta to disable/enable previewing the beta test commands.

Sample

```
bstnA> show terminal
Terminal Options
----------------

Terminal Idle Timeout:        disabled
Terminal Idle Timeout value:
Terminal Expand Prompt:       enabled
Terminal Confirmation:        enabled
Expert Mode:                  disabled
Terminal Length:              24 lines,  more disabled
Terminal Width:               80
Terminal Logfile:             disabled
Terminal History:             10 lines
Script History Logging:       disabled
Terminal Character Set:       ISO 8859-1
Stop Scripts on error:        disabled
Beta Commands:                disabled
```

Related Commands

terminal timeout
terminal expand-prompt
terminal confirmation
terminal length
terminal width
terminal logging
terminal history
show history
terminal character-set
terminal beta

terminal beta

Purpose Use the terminal beta command to unlock CLI beta test commands.

Use the no form of the command to lock these commands.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
terminal beta
no terminal beta
```

Default(s) no terminal beta

Guidelines The terminal beta command provides access to the CLI beta test commands to be released in later versions of F5 software. Use terminal beta to unlock these commands for previewing and testing purposes only. Contact your F5 representative for more information on the availability of these commands.

Entering terminal beta causes the beta test commands to appear in Help.

Use show terminal to find the current setting (locked or unlocked) for the beta test commands.

Samples

```
bstnA> terminal beta

This command unlocks UNSUPPORTED features. Are you sure? [yes/no] yes
```

unlocks the beta test commands and features.

```
bstnA> no terminal beta
```

hides all beta commands and features.

Related Commands show terminal

terminal character-set

Purpose Use this command to enable or disable CLI support for Unicode characters (such as the characters used in Chinese or Japanese languages). Match this to the character set used in the local network.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
terminal character-set {iso-8859-1 | unicode-utf-8}
no terminal character-set
```

Default(s) `iso-8859-1`

Guidelines The terminal character set is the one that was used in your network before the ARX was installed. The correct character setting is required so that you can enter back-end filer names into the CLI, and so that the CLI can properly display filer names and service names in its show commands.

The **iso-8859-1** setting is for single-byte-character languages like English. The other setting, **unicode-utf-8**, supports all languages including those with characters larger than one byte.

You can use Unicode characters for naming objects such as filers and back-end shares. They are not used in CLI prompts, help, or error messages.

To view the current character set, use show terminal.

Samples

```
bstnA> terminal character-set unicode-utf-8
```

allows multi-byte characters.

```
bstnA> no terminal character-set
```

restricts the CLI to single-byte characters only.

Related Commands show terminal

terminal clear

Purpose Use this command to clear the screen of all text entries. It is also useful in writing scripts.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`terminal clear`**

Default(s) None

Guidelines Use this command when you want to clear the screen and start over. It is also useful for capturing screen output, scripting, and pausing a CLI operation.

Sample `bstnA>` **`terminal clear`**

clears the screen of all entries and returns you to the CLI prompt (in the same mode) at the top of the screen.

Related Commands show terminal
reload

terminal confirmation

Purpose Some sensitive configuration commands require user confirmation (for example, "Are you sure? [y/n]"). For non-interactive scripts, you can disable the confirmations with no terminal confirmation.

Use the affirmative form, terminal confirmation, to make the CLI start prompting for confirmations again.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
terminal confirmation
no terminal confirmation
```

Default(s) Enabled

Guidelines Use show terminal to find the current setting for terminal confirmation.

Samples

```
bstnA(cfg)# no terminal confirmation
```

disables terminal-confirmation prompts. For example, if you reboot the ARX the confirmation message ("Are you sure? [y/n]") will be skipped.

```
bstnA(cfg)# terminal confirmation
```

reinstates terminal-confirmation.

Related Commands show terminal

terminal expand-prompt

Purpose Use the terminal expand-prompt command to display the fully expanded CLI command prompt.

Use the no version of the command to shorten the CLI command prompt (mode only).

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
terminal expand-prompt
no terminal expand-prompt
```

Default(s) Enabled

Guidelines The expand-prompt feature shows the name of the current object (such as the namespace or volume) in the prompt. The show terminal command indicates whether or not this feature is enabled.

Samples

```
bstnA(gbl-ns[[nfsNamespace1])#no terminal expand-prompt
```

disables the expanded prompt:

```
bstnA(gbl-ns[nfsNamespace1])# terminal expand-prompt
bstnA(gbl-ns[nfsNamespace1])#
```

reinstates the expanded prompt.

Related Commands show terminal

terminal expert

Purpose Commands that create new configuration objects (such as namespaces and global servers) prompt for confirmation before creating the object. Use the terminal expert command to remove these confirmation prompts.

Use the no version of the command to bring back the confirmation prompts.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
terminal expert
no terminal expert
```

Default(s) no terminal expert

Guidelines This command only applies to confirmations for new configuration objects; use no terminal confirmation to eliminate all confirmation prompts in the CLI.

The CLI prompts for all new configuration objects by default. This is a feature for safeguarding against typing mistakes.

Samples

```
bstnA(gbl-ns[nfsNamespace1])# terminal expert
```

disables prompts for new objects (like namespaces).

```
bstnA(gbl-ns[[nfsNamespace1])# no terminal expert
```

re-enables the safe prompts.

Related Commands terminal confirmation
show terminal

terminal history

Purpose

The CLI keeps a configurable number of CLI commands in a list, so that you can display the list with the show history command. Use the terminal history command to set the number of CLI commands to keep.

Use the no form of this command stop saving any CLI history.

Mode

(any)

Security Role(s)

operator

Syntax

```
terminal history number
terminal history script
no terminal history [script]
```

number (0-255) is the number of CLI commands to keep in the terminal-history list.

script (optional) enables logging for CLI scripts. If the run command invokes a script, the CLI keeps terminal history for the commands in the script, too.

Default(s)

10 lines

script history is not kept

Guidelines

Use show history to view the terminal history. You can also use the up-arrow key to view the history one command at a time.

Samples

```
bstnA(cfg)# terminal history 100
```

keeps 100 CLI commands in the terminal history.

```
bstnA(cfg)# terminal history script
```

keeps terminal history for commands in CLI scripts as well as those entered by an administrator.

```
bstnA(cfg)# no terminal history
```

disables terminal history.

Related Commands

show history
show terminal

terminal length

Purpose

By default, the CLI pipes all of its output through the *more* paging program. Each page of *more* output has a configurable number of lines; use the terminal length command to set the number of lines in each CLI page.

Use no terminal length to return to the default length and enable the *more* feature.

Mode

(any)

Security Role(s)

crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

`terminal length` ***`number`***
`no terminal length`

number (0-255) is the number of lines in each page of CLI output. 0 (zero) means "infinite;" this effectively disables the CLI *more* feature.

Default(s)

24 (and see the *Guidelines* below)

Guidelines

The default terminal length applies to serial connections to the Console port as well as remote connections to one of the IP-based management ports. However, if a remote connection (such as a Telnet session) has a length setting, it overrides the default. If this command sets a specific length, it overrides both.

Use terminal length 0 to disable the *more* feature. The no terminal length command enables the *more* feature and returns the length to the default.

Use terminal width to determine the number of characters in each line of CLI output. Use show terminal to find the current terminal length, and to find whether or not *more* is disabled.

Samples

```
bstnA(cfg)# terminal length 200
```
sets each page of CLI output to 200 lines.

```
bstnA(cfg)# terminal length 0
```
allows for unlimited output. This effectively disables the *more* paging feature.

```
bstnA(cfg)# no terminal length
```
reverts the page length to the default.

Related Commands

terminal width
show terminal

terminal logging

Purpose Use the terminal logging command to capture output from CLI show commands to a report file. This starts a session where the output of all subsequent CLI commands is recorded in the report file.

Use the no form of this command to stop the current terminal-logging session.

Mode (any)

Security Role(s) operator (any, including backup-operator)

Syntax

```
terminal logging [report-prefix]
no terminal logging
```

report-prefix (optional, 1-1024 characters) is the prefix for a report file. The CLI logs its output to a report file named as follows:

report-prefix_yyyymmddHHMM.rpt, where *report-prefix* is chosen here, *yyyy* is the current year, *mm* is the current month, *dd* is the day, *HH* is the hour, and *MM* is the minute.

The default prefix is "cli_."

Default(s) Disabled

report-prefix -"cli_."

Guidelines The CLI displays the name of the report file when you enter the command. Use show reports type CLI to view all CLI-log reports, or show reports *report-name* to view a specific CLI-log file.

The show terminal command indicates whether or not terminal logging is enabled.

Samples

```
bstnA(cfg)# terminal logging showoutput

  Generating CLI output report: 'showoutput_200903301531.rpt'..
```

enables terminal logging and captures show command output to the named log file.

```
bstnA(cfg)# no terminal logging
```

disables terminal logging.

Related Commands show reports
show terminal

terminal stop-on-error

Purpose

In CLI scripts, use terminal stop-on-error to stop script processing if the CLI responds with an error.

Use no terminal stop-on-error to continue processing the CLI script through all errors.

Mode

(any)

Security Role(s)

crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
terminal stop-on-error
no terminal stop-on-error
```

Default(s)

no terminal stop-on-error

Guidelines

Use show terminal to find the current setting (enabled or disabled) for stop-on-error.
Use copy ... scripts *script-name* to download a CLI script onto the switch.

Samples

```
bstnA(cfg)# terminal stop-on-error
```

stops processing on a CLI error.

```
bstnA(cfg)# no terminal stop-on-error
```

allows a CLI script to continue processing through all errors.

Related Commands

copy ftp
copy scp
copy tftp
copy {nfs|cifs}
show terminal

terminal timeout

Purpose By default, a CLI terminal logs off after 15 minutes of idle time, where idle time is time without any user input. Use no terminal timeout to allow infinite idle time.

Use the affirmative form, terminal timeout, to re-instate time-outs.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
no terminal timeout
terminal timeout
terminal timeout seconds
```

seconds (0-2048) is the number of idle seconds before timing out the current CLI session. If you omit this value with the terminal timeout command, you set the default timeout of 900 (seconds, or 15 minutes).

Default(s) 900 (seconds, or 15 minutes) of idle time

Guidelines Use show terminal to find the current setting (enabled or disabled) for terminal time-out.

Samples

```
bstnA(cfg)# no terminal timeout
```

allows the terminal to be idle for infinite time without logging off.

```
bstnA(cfg)# terminal timeout
```

reinstates terminal-timeout at its default, 900 seconds.

```
bstnA(cfg)# terminal timeout 1800
```

sets the terminal-timeout 1800 seconds, or 30 minutes.

Related Commands show terminal

terminal width

Purpose Use the terminal width command to set the maximum number of characters in each line of CLI output.

Use no terminal width to return to the default width.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
terminal width number
no terminal width
```

number (0-255) is the maximum number of characters in each line of CLI output.

Default(s) 80 (and see the *Guidelines*, below)

Guidelines The default terminal width applies to serial connections to the Console port as well as remote connections to one of the IP-based management ports. However, if a remote connection (such as a Telnet session) has a width setting, it overrides the default. If this command sets a specific width, it overrides both.

Use terminal length to determine the number of lines in each page of CLI output. Use show terminal to find the current terminal width.

Samples

```
bstnA(cfg)# terminal width 120
```

sets each line of CLI output to 120 characters.

```
bstnA(cfg)# no terminal width
```

returns to the default line width.

Related Commands terminal length
show terminal

4

Administrative Users

group

Purpose

An administrative *group* defines the access privileges for a list of administrative users. Use the group command to add a new group to the ARX, or to edit an existing group. You can also use this command to re-create or edit a Windows group, so that you can assign administrative privileges to the group's Windows users.

Use the no form of this command to remove a group.

Mode

gbl

Security Role(s)

crypto-officer

Syntax

group *name*
no group *name*

name (1-64 characters) is a name that you choose for the group. Surround this with quotation marks ("") if it contains any spaces.

Default(s)

Several default groups are defined, each with different access privileges: crypto-officer, storage-engineer, network-engineer, network-technician, backup-operator, and operator.

Guidelines

This command puts you into gbl-group mode. From gbl-group mode, you can use the role command to select a *role* for all the group's users. Each role is associated with a set of CLI commands; you can use the role command multiple times to assign multiple roles to a group. You must set at least one role for the group to function. Default groups (such as network-engineer) already have their roles configured.

For a group of Windows administrators, defined externally in your Active Directory, you can choose a group ***name*** that is the same as an existing Windows group. For example, you could create a group named "Domain Admins." Then use the windows-domain (gbl-group) command to specify one or more domains where the group is allowed access; Domain Admins in "medarch.org" may be allowed to access the CLI, but Domain Admins in "competitor.com" may not. Finally, use the authentication command to allow Active-Directory authentications at the CLI and/or GUI. Windows users in the group/domain can then log into the CLI or GUI with their Windows username and password, and they get the access privileges assigned by the role command. This type of group does not require any users; all users are established externally, on your Windows Domain Controllers.

For a locally-defined group of administrators, use the user (gbl-group) command to add each administrator to the group.

Use the show group all command for a list of existing groups.

Samples

```
bstnA(gbl)# group superusers
bstnA(gbl-group[superusers])#
```

creates the group, "superusers."

```
bstnA(gbl)# group "Domain Users"
```

creates another group, "Domain Users."

Related Commands

role
windows-domain (gbl-group)
authentication
user (gbl-group)
show group all

group (gbl-user)

Purpose Use the group command to add the current user to an additional group. Use no group to remove the current user from a group.

Mode gbl-user

Security Role(s) crypto-officer

Syntax

```
group name
no group name
```

name (1-64 characters) is a name of an existing group.

Default(s) None

Guidelines A user can belong to multiple groups, where each group is associated with a role. The group's role determines the access privileges for its users. Every command in this manual is labeled with the role or roles that are permitted to use the command. See *Security Role(s)*, above, as an example: this command can only be used by users who belong to groups with the "crypto-officer" role.

The ARX is shipped with several pre-defined groups. Use the show group all command to show all groups, and use show group roles to show the role associated with each group. You can add new groups with the group command.

Samples

```
bstnA(gbl-user[newuser])# group storage-engineer
```

adds the "newuser" account to a group, "storage-engineer."

```
bstnA(gbl-user[newuser])# no group testgroup2
```

removes the "newuser" account from the "testgroup2" group.

Related Commands

group
show group all
show group roles

password

Purpose Use this command to change the password for your administrative-user account, the one used for your current login session.

Mode priv-exec

Security Role(s) crypto-officer, storage-engineer, network-engineer, or network-technician

Syntax **`password`**

Default(s) None

Guidelines The CLI first authenticates you by prompting for the old password. Then it prompts for the new password that you want, followed by a prompt to re-enter the new password. The maximum number of characters is 28. For maximum security, choose a password with at least one of each of the following:

- lower-case letter (a-z),
- upper-case letter (A-Z),
- number (0-9), and
- non alpha-numeric symbol ($, #, *, @, and so forth). You cannot use any of the following characters: **| & ; ( ) < > ` -**.

An administrator with the crypto-officer role can use the password (gbl-user) command to change the password for any user account.

Sample

```
bstnA# password
Old Password: acopia
New Password: myNewPa$$wd
Validate Password: myNewPa$$wd
bstnA#
```

changes the password for the current account.

Related Commands password (gbl-user)

password (gbl-user)

Purpose Use the password command to change the password for an administrative-user account.

Mode gbl-user

Security Role(s) crypto-officer

Syntax **`password`**

Default(s) None

Guidelines This command changes the password for an administrative-user account that is not necessarily your own. A crypto officer can use this command to manage administrative accounts on the ARX. To change the password on the current account, any administrator can use the password command from priv-exec mode.

This only applies to locally-defined accounts. It has no effect on user accounts defined externally in the Windows Active Directory. (See windows-domain (gbl-group) for information about allowing externally-defined users to access the ARX as administrators.)

The CLI prompts for a new password for the account, then prompts you to re-enter the same password. The maximum number of characters is 28. For maximum security, choose a password with at least one of each of the following:

- lower-case letter (a-z),
- upper-case letter (A-Z),
- number (0-9), and
- non alpha-numeric symbol ($, #, *, @, and so forth). You cannot use any of the following characters: **| & ; ( ) < > ` -**.

Use the show users command to show all administrative users.

Sample

```
bstnA(gbl-user[newuser])# password
Password: n3wcrypt1cPa$$wd
Validate Password: n3wcrypt1cPa$$wd
bstnA(gbl-user[newuser])#
```

changes the password for the “newuser” account.

Related Commands user
show users

role

Purpose

A group's *role* determines the accessible CLI commands for its administrative users. Each group can have multiple roles. Use the role command to set a role for the current group.

Use the no role command to remove a role from the group.

Mode

gbl-group

Security Role(s)

crypto-officer

Syntax

```
role {operator | backup-operator | network-technician |
network-engineer | storage-engineer | crypto-officer}
no role {operator | backup-operator | network-technician |
network-engineer | storage-engineer | crypto-officer}
```

Select one of the following roles:

operator is a clerical administrator,

backup-operator runs backup and restore operations on volumes (see restore data),

network-technician configures layer-2 and IP networks under the guidance of a network-engineer,

network-engineer designs network topologies,

storage-engineer designs and configures network storage, and

crypto-officer keeps passwords and manages network security.

Default(s)

None

Guidelines

Each role is associated with a set of CLI commands. Administrators with the storage-engineer role, for example, have access to commands that are associated with storage management. Use show group roles to show all configured groups and roles.

Each CLI command has one or more *Security Roles* that are listed in this manual. If an administrator's group has one of the roles that can access a command, the administrator can use the command. For example, an administrator with the **network-technician** role can see a command that is assigned to **network-engineer** and **network-technician**, but cannot see a command that allows only **storage-engineer**s.

Administrators with the crypto-officer role can log into a booting ARX sooner than administrators with other roles. Other roles wait until gbl mode is accessible before allowing a login, whereas a crypto officer can log in as soon as cfg mode is available.

Samples

```
bstnA(gbl-group[Domain Admins])# role crypto-officer
```

adds the "crypto-officer" role to the "Domain Admins" group.

```
bstnA(gbl-group[superusers])# no role network-technician
```

removes the "network-technician" role from the "superusers" group.

Related Commands

group
show group roles

ssh-key

Purpose

Use the ssh-key command to paste a public SSH key into the current administrative account.

Use the no ssh-key command to remove one or more SSH keys from the current account.

Mode

gbl-user

Security Role(s)

crypto-officer

Syntax

```
ssh-key {dsa | rsa | rsa1} [“public-key”]
```

dsa | rsa | rsa1 is a required choice, which selects the encryption type and SSH version for the key:

dsa is DSA over SSHv2,

rsa is RSA over SSHv2, and

rsa1 is RSA over SSHv1. By default, SSHv1 is not supported, use ssh-v1 enable to enable SSHv1 support.

public-key (optional, 1-2500 characters) is a public SSH key, pasted from the client. Use quotation marks around this string, as it invariably contains spaces. Take care that the copy/paste operation does not add any <Return> or <Line-Feed> characters to break up the string. If you omit the ***public-key***, the CLI prompts for it on the next line.

Syntax: No Form

```
no ssh-key id key-id
no ssh-key fingerprint fingerprint
no ssh-key {dsa | rsa | rsa1}
no ssh-key all
```

key-id (1-2,147,483,647) identifies the SSH key by an ID number assigned at the switch. Use show ssh-user to see all key IDs for all SSH keys.

fingerprint (1-50) identifies the SSH key by its fingerprint. Use show ssh-user to see all fingerprints.

dsa | rsa | rsa1 identifies the type of key to remove. This removes all SSH keys of the given type from the current administrative account.

all removes all SSH keys from the current administrative account.

Default(s)

None

Guidelines

When administrators access the CLI through SSH, they are typically challenged for the account password. The SSH protocol supports *public-key authentication*, which skips this challenge. When an administrator accesses SSH on the switch, the switch’s SSH server attempts to use the administrator’s public-key first. If the public key is configured properly for this administrator, he or she never sees a password challenge.

An administrator’s SSH keys are unique to every client machine that they use. You can therefore configure multiple SSH keys for each administrative user account.

The show ssh-user command shows all administrative accounts with SSH public keys.

Samples

```
bstnA(gbl-user[su])# ssh-key dsa “ssh-dss
AAAAB3NzaC1kc3MAAACBAPqSVxs6Soxs5D9G7Ul8dQrf7Eo7vNdTawaH0K7DsyV2ND0Rqx
ttRtNpw/fdIcm5cHOrYW4OYL6HJesMeJPguAzY8hbTkwsz+uRJLFnmRTy236DXDFiTc38E
r6UQCoa1On9VrKWhoEGNe1YCn+cIsb3S+s44QPOx9GPFSVN1hqdVAAAAFQC8x+2VKzUH16
xrAMKuvVh50c53lwAAAIEAvOgRX8Ek2e/uCCJXlme0n7EsL3+yTEsOP7C9Bsl05KoCAgCS
YP8G/1rc372Vy0xF3PGL9QsI/bj+48SEAuJJTpJR1eB9MLpwmraVa/IsX16Xhr34eLDwH3
NwtlwqRH9fhkjnWwhEoLRC7Bf/g493HoXPD2dNjbKvqiMgq+s7CBEAAACAcAF+a+S/0OUN
fpuv6QPV+SX9WoaazJthtUiP8pI4yl6sVAhp3Op5LxWT58Xl4ed+F0vUR2cfdjAF23YGYR
wK2c2h4FjnoBjLuoodhXJ+xAC/DPb4EvwEcBtqlPnpWzsPlAFX/I1pPA4fUyUOOifCrP12
etsoZ9mnxawLRAAEa+A= juser@clientLinux”
```

adds an SSH key to the “su” user account.

```
bstnA(gbl-user[su])# ssh-key dsa
Enter user's public key: “ssh-dss
AAAAB3NzaC1kc3MAAACBAPqSVxs6Soxs5D9G7Ul8dQrf7Eo7vNdTawaH0K7DsyV2ND0Rqx
ttRtNpw/fdIcm5cHOrYW4OYL6HJesMeJPguAzY8hbTkwsz+uRJLFnmRTy236DXDFiTc38E
r6UQCoa1On9VrKWhoEGNe1YCn+cIsb3S+s44QPOx9GPFSVN1hqdVAAAAFQC8x+2VKzUH16
xrAMKuvVh50c53lwAAAIEAvOgRX8Ek2e/uCCJXlme0n7EsL3+yTEsOP7C9Bsl05KoCAgCS
YP8G/1rc372Vy0xF3PGL9QsI/bj+48SEAuJJTpJR1eB9MLpwmraVa/IsX16Xhr34eLDwH3
NwtlwqRH9fhkjnWwhEoLRC7Bf/g493HoXPD2dNjbKvqiMgq+s7CBEAAACAcAF+a+S/0OUN
fpuv6QPV+SX9WoaazJthtUiP8pI4yl6sVAhp3Op5LxWT58Xl4ed+F0vUR2cfdjAF23YGYR
wK2c2h4FjnoBjLuoodhXJ+xAC/DPb4EvwEcBtqlPnpWzsPlAFX/I1pPA4fUyUOOifCrP12
etsoZ9mnxawLRAAEa+A= juser@clientLinux”
```

adds the same SSH key by responding to a CLI prompt.

```
bstnA(gbl-user[su])# no ssh-key rsa
```

removes all RSA/SSHv2 keys from the “su” user account.

```
bstnA(gbl-user[admin])# no ssh-key id 7
```

removes SSH-key 7 from the “admin” user account.

Related Commands

user
show ssh-user
ssh-v1 enable

show group all

Purpose Use the show group all command to display all administrative groups configured for the ARX.

Mode (any)

Security Role(s) crypto-officer

Syntax **`show group all`**

Guidelines This shows all administrative groups. Use the show group users command to find the locally-defined administrative users in each group.

Use the show group roles command to find the administrative role assigned to each group. This shows the access privileges for the group's users.

Related Commands group
show group users
show group roles

Figure 4.1 *Sample Output: show group all*

```
bstnA(gbl-group[superusers])# show group all

Configured Groups

     Group Name
     -------------
     backup-operator
     operator
     network-technician
     storage-engineer
     network-engineer
     crypto-officer
     admins
     Domain Users
     Backup Operators
     Administrators
     Domain Admins
     Enterprise Admins
```

show group roles

Purpose Each group has a *role* which defines CLI-access privileges for the group's users. Use the show group roles command to show all administrative groups and their roles.

Mode (any)

Security Role(s) crypto-officer

Syntax **`show group roles`**

Guidelines This shows a table with two columns: administrative groups in the left column and their roles in the right column.

Use the group command to configure a new group.

Use the role command to set a group's role.

Use the show group users command to find the administrative users in each locally-defined group.

Sample

```
bstnA(gbl-group[admins])# show group roles
Group Roles
-----------

Group                          Roles
-----------------------------  -----------------------------
Administrators                 backup-operator
Administrators                 operator
Administrators                 network-engineer
Administrators                 storage-engineer
Administrators                 crypto-officer
admins                         storage-engineer
Backup Operators               backup-operator
Backup Operators               operator
backup-operator                backup-operator
crypto-officer                 crypto-officer
Domain Admins                  backup-operator
Domain Admins                  operator
Domain Admins                  network-engineer
Domain Admins                  storage-engineer
Domain Admins                  crypto-officer
Domain Users                   operator
Enterprise Admins              backup-operator
Enterprise Admins              operator
Enterprise Admins              network-engineer
Enterprise Admins              storage-engineer
Enterprise Admins              crypto-officer
network-engineer               network-engineer
network-technician             network-technician
operator                       operator
storage-engineer               storage-engineer
```

Related Commands

group
role
show group users

show group users

Purpose Use the show group users command to cross-reference the switch's administrative groups and their locally-defined users.

Mode (any)

Security Role(s) crypto-officer

Syntax **`show group users`**

Guidelines This shows a table with two columns: administrative groups in the left column and their users in the right column. This only shows groups with locally-defined users; it does not include groups that are defined externally in your Active Directory.

Use the user command to configure a new administrative user.

Use the group command to create (or edit) a group. Use the user (gbl-group) command to add a user to the group.

Use the show group roles command to find the administrative role assigned to each group.

Samples `bstnA(gbl-group[admins])# show group users`
lists all users in their groups. See Figure 4.2 for sample output.

Related Commands user
group
user (gbl-group)
show group roles

Figure 4.2 *Sample Output: show group users*

```
bstnA(gbl-group[admins])# show group users
Group Users
-----------

Group                          User
-----------------------------  -----------------------------
admins                         adm1
admins                         adm12
crypto-officer                 admin
crypto-officer                 newadmin
operator                       newadmin
operator                       admin
```

show ssh-user

Purpose Use the show ssh-user command to show the SSH public keys entered for administrative users, if any.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax `show ssh-user [account-name]`

account-name (optional, 1-32 characters) identifies a particular administrative account to show. If you omit this, the SSH keys are shown for all administrative accounts.

Guidelines Each administrative account can store multiple SSH public keys. Each public key corresponds to one administrator on one SSH-client machine; if the administrator's public key is copied from the client machine to the administrative account, the administrator can access the account through SSH without providing a password.

This shows a table with one row per SSH public key. Each row has the following information:

User is the name of the administrative-user account, set with the user command.

KeyId is an internally-assigned ID for this public key. You can use this with no ssh-key id to remove the SSH key from the account.

Type is "dsa" (DSA encryption over SSHv2), "rsa" (RSA encryption over SSHv2), or "rsa1" (RSA encryption over SSHv1).

Fingerprint is used by SSH as a shorter equivalent to the public key. This is a unique identifier for a particular user at a particular host. You can use this with no ssh-key fingerprint to remove the SSH key from the account.

Samples

`bstnA# show ssh-user`

lists the SSH public keys configured for all administrative accounts. See Figure 4.3 for sample output.

`bstnA# show ssh-user admin`

lists the SSH public keys for one account only, "admin." See Figure 4.4.

Related Commands user
ssh-key

Figure 4.3 *Sample Output: show ssh-user*

```
bstnA# show ssh-user

User            KeyId      Type  Fingerprint
--------------- ---------- ----- -----------------------------------------------
newadmin        1          dsa   1e:ca:4b:9f:0d:51:97:98:1f:d8:26:81:8e:ab:3d:9b
adm12           2          dsa   8d:45:67:07:53:5f:61:9c:54:47:a8:76:4c:9a:93:05
admin           3          dsa   0b:fa:85:93:15:f9:ff:7e:8a:4c:1b:16:ba:4a:de:e7
admin           4          dsa   1e:ca:4b:9f:0d:51:97:98:1f:d8:26:81:8e:ab:3d:9c
bstnA#
```

Figure 4.4 *Sample Output: show ssh-user admin*

```
bstnA# show ssh-user admin

User            KeyId      Type  Fingerprint
--------------- ---------- ----- ------------------------------------------------
admin           3          dsa   0b:fa:85:93:15:f9:ff:7e:8a:4c:1b:16:ba:4a:de:e7
admin           4          dsa   1e:ca:4b:9f:0d:51:97:98:1f:d8:26:81:8e:ab:3d:9b
bstnA#
```

show users

Purpose Use the show users command to display all administrative users that have been locally configured for the ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show users`**

Guidelines The output contains two tables.

Configured Users is a list of all local administrative users. Use the user command to configure a new local user. This does not show any users defined externally in the Windows Active Directory.

Current User shows login name used for the current administrative session. The first set of rows under the user name is the group(s) to which the current user belongs. Each of these rows is labeled "group." If you logged in through Active Directory, using your Windows credentials, there is another set of rows labeled "role:" these are the administrative roles, or privileges, assigned to you.

Use the show group all command to show all administrative groups.

Sample

```
bstnA(gbl)# show users

Users
    Configured Users
    ----------------
    adm1
    adm12
    admin
    newadmin


    Current User
    ------------
    medarch\administrator
        group   Domain Admins
        group   Enterprise Admins
        group   Domain Users
        role    operator
        role    network-engineer
        role    storage-engineer
        role    crypto-officer
        role    backup-operator
```

Related Commands user
show group all

user

Purpose

Use the user command to add a new, local administrative user to the ARX.

Use the no form of this command to remove a user account.

Mode

gbl

Security Role(s)

crypto-officer

Syntax

`user` ***`username`***
`no user` ***`username`***

username (1-32 characters) is a username that you choose.

Default(s)

operator is the default group for a new user

Guidelines

If the user account is new, the CLI prompts for a password. The maximum number of characters is 28. For maximum security, choose a password with at least one of each of the following:

- lower-case letter (a-z),
- upper-case letter (A-Z),
- number (0-9), and
- non alpha-numeric symbol ($, #, *, @, and so forth). Avoid the characters listed above for *username*.

If you use this command to edit an existing account, there is no password prompt.

This puts you into gbl-user mode, which has commands for editing the user account. Use the password command to change the password for the account. For administrators that use SSH to access this account, you can use the ssh-key command to add their public key to the account; if they log in from a management station with the same public key, they do not have to enter the account password.

A user's *group* determines its CLI-access privileges. A new user's default group, operator, has minimal access privileges. Each user can belong to multiple groups, thereby expanding his access privileges. After you create the user account with this command, you can use the group (gbl-user) command to add this user to another group.

Use the show users command to show all administrative users.

Samples

```
bstnA(gbl)# user newuser
Password: crypt1cPa$$wd
Validate Password: crypt1cPa$$wd
bstnA(gbl-user[newuser])#
```

creates the user, "newuser," with the password, "crypt1cPa$$wd."

Related Commands

password
ssh-key
group (gbl-user)
show users

user (gbl-group)

Purpose

Use the gbl-group user command to add a local administrative user to the current group.

Use no user to remove a local user from the current group.

Mode

gbl-group

Security Role(s)

crypto-officer

Syntax

`user` ***`username`***
`no user` ***`username`***

username (1-64 characters) identifies an administrative user account on this ARX. Use show group users for a list of all available user accounts.

Default(s)

None

Guidelines

This command adds a user to the current group. The no form of the command removes a user, thus revoking the group privileges for that user. These administrative users are locally defined on the ARX.

This command is unnecessary for a Windows group that is defined in your Active Directory (AD). For a group defined in the AD, you can use the windows-domain (gbl-group) command to specify the domain(s) where the group's users can gain access to the ARX. Any valid Windows user in the group and domain can use their Windows username and password to gain access. The users are defined externally, on your Windows Domain Controllers.

Sample

```
bstnA(gbl-group[superusers])# user newuser
```

adds "newuser" to the current group.

Related Commands

group
show group users

windows-domain (gbl-group)

Purpose

An administrative group can be equivalent to one or more Windows groups in your Active Directory. If the group name is the same as a valid Windows-group name, you can allow the group's Windows users to log into the CLI and/or GUI with their Windows credentials. This command declares a valid Windows domain for the current Windows group. You can re-issue the command multiple times to allow access to this group in multiple trusted domains.

Use the no form of the command to remove a Windows domain from the group configuration. This prevents the Windows users from the given domain/group from logging in with their Windows credentials.

Mode

gbl-group

Security Role(s)

crypto-officer

Syntax

```
windows-domain domain-name
no windows-domain domain-name
no windows-domain all
```

domain-name (1-256 characters) is the name of the Windows domain for this Windows group. This must be a Fully-Qualified-Domain Name (FQDN) so Windows users can log in with it. This makes it possible for Windows users to authenticate with Kerberos.

all (optional with the no form) removes all of the Windows domains that have been associated with this group.

Default(s)

None

Guidelines

This command establishes a Windows domain where the group has administrative privileges. This implies that the group name exists in the Active-Directory configuration (outside the ARX), that the group is valid in this Windows domain, and that one or more of the group's users should have administrative access to the ARX.

The role command establishes the administrative permissions for members of this group. To see the current roles for the group, use show group roles.

You must also configure the GUI and/or the CLI to allow Active-Directory (AD) authentication. To open a given access point (such as HTTP or HTTPS for the GUI) for AD authentication, use the authentication command.

Sample

```
bstnA(gbl-group[Domain Users])# windows-domain medarch.org
```

declares that Windows clients from the "Domain Users" group in "medarch.org" can gain administrative access to the ARX.

Related Commands

group
role
show group roles
show group all
authentication

5

Chassis Management

This chapter contains an alphabetical list of commands for managing the chassis of the ARX.

Slot Locations

Some of these commands require module identification by slot ID.

The ARX-4000 has two “slots,” 1 and 2. The smaller ARX-500, and ARX-2000 platforms have a single “slot,” slot 1. The virtual ARX-VE is said to also have a single, virtual slot: also slot 1.

clear metalog usage

Purpose Namespace software and networking software record important log information, called *metalog* data, to be used by a redundant peer in the event of a failover. The ARX-1500 and ARX-2500 store their metalog data on internal disk partitions. You can use the show metalog usage command to view the usage of these metalog partitions over time. Use this command to clear the metalog-usage statistics for the current ARX.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `clear metalog usage`

Default(s) None

Valid Platforms ARX-1500 and ARX-2500

Guidelines The show metalog usage command shows up-to-date usage statistics for all metalog partitions in the ARX. Use this command to clear those usage statistics for the entire system.

The CLI prompts for confirmation before clearing any statistics; enter **yes** to continue.

The show metalog usage command is unavailable for a minute after you issue this command.

Sample

```
stoweA# clear metalog usage
Clear metalog statistics? [yes/no] yes
Metalog driver usage statistics will not be available for a minute.
```

clears the metalog-usage statistics for the ARX named "stoweA."

Related Commands show metalog usage

clear nvr

Purpose The Non-Volatile RAM (NVRAM) is battery-backed memory on the ACM. The NVRAM stores all database transactions for the namespace processes; if the processes fail, they can recover by replaying these transactions. These transactions are called *metalog* data. If the NVRAM gets corrupted somehow, use the clear nvr command to erase it and reboot the chassis.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `clear nvr [reload]`

reload (optional) causes the switch to come back online after halting. If you omit this option, the switch halts until you manually restore power.

Default(s) None

Valid Platforms ARX-500, ARX-2000, or ARX-4000

Guidelines The CLI prompts you with a warning before it clears the NVRAM. Enter **yes** to continue, then enter your username at the next prompt. The switch then halts. If you omitted the **reload** option, you must turn the power back on to restore the ARX to service.

Important

This causes namespace corruption in a running switch, in addition to a service interruption. Only use this command on the advice of F5 personnel.

The following events are appropriate for clearing the NVRAM:

- When the system is first booted. This brings the NVRAM battery online for the first time, so there may some garbled data and pre-startup ECC errors.
- After a hardware failure where the NVRAM got corrupted with ECC errors. Use the show chassis [nvram] command to check for non-recoverable ECC errors.
- (Optional) After clearing all namespaces with delete startup-config and rebooting. This eliminates stale namespace data in the NVRAM.

Sample

```
bstnA# clear nvr reload

WARNING WARNING WARNING WARNING WARNING WARNING WARNING WARNING

This will erase the non-volatile RAM (NVRAM).

The effects are:
    All pending namespace transactions will be lost.
    Any data that is in transition may become corrupted.
    The chassis will reload immediately after clearing the NVRAM.
This should only be done on a switch that has been completely
unconfigured and removed from the network,
or on the advice of technical support.

Are you sure you want to clear non-volatile RAM and reload the system?
[yes/no] yes
Please enter your userid to confirm clearing the non-volatile RAM:
admin

% INFO: Success clearing non-volatile RAM. Rebooting system....


System is resetting.
...
```

Related Commands delete startup-config

clock set

Purpose Use the clock set command to set the time and/or date at the ARX.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
clock set HH:MM:SS
clock set HH:MM:SS mm/dd/yyyy
```

HH:MM:SS is the time (for example, 06:00:00). Enter the local time, not the UTC time.

mm/dd/yyyy (optional) sets the date (for example, **12/31/2003** for the 31st of December, 2003).

Default(s) If you omit the date specification, the date is unchanged.

Guidelines Several processes, including those for RON, redundancy, and policy, have clock-dependant algorithms that can be confused by a major change in time. In general, you should avoid this command and use an NTP server instead (see below). For this reason, the CLI prompts for confirmation before changing the time; enter **yes** to proceed.

This command is unnecessary if you use the ntp server command to synchronize the time to an accurate time source. An NTP server overrides the time you set with this command. For the reason stated above, NTP is strongly recommended.

Use the show clock command to view the current time and date setting.

Samples

```
bstnA# clock set 11:00:00
Changing the time may have an adverse impact on the switch.

Are you sure? [yes/no] yes
```

sets the clock to 11 AM, local time.

```
bstnA# clock set 23:30:00
Changing the time may have an adverse impact on the switch.

Are you sure? [yes/no] yes
```

sets the clock to 11:30 PM.

```
bstnA# clock set 23:30:00 10/24/2004
Changing the time may have an adverse impact on the switch.

Are you sure? [yes/no] yes
```

sets the clock to 11:30 PM on October 24, 2004.

Related Commands

ntp server
show clock

clock timezone

Purpose Use the clock timezone command to set the time zone at the ARX.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
clock timezone region city
clock timezone time-zone-name
clock timezone offset
```

region (1-64 characters) is the continent or ocean of the closet major city. Use **<Tab>** to see a list of possible options.

city (1-64 characters) is the closet major city. As above, use **<Tab>** to see a list of possible options.

time-zone-name (1-64 characters) is the name of the local time zone, if known (for example, EDT, EST, or CDT). The CLI uses this to look up the region and city. Many time-zone names are ambiguous; for example, CST maps to Central Standard Time, China Standard Time, and Australian Central Standard Time. The CLI chooses a region and city based on the current customer base. You can use show clock to see the chosen region and city.

offset (optional; -1400 to 1200) is the offset from Coordinated Universal Time (UTC), in hours and minutes (for example, five hours back is -0500 and 7.5 hours forward is 0730). As with the time-zone-name, the CLI uses this to look up the region and city. Many offsets are ambiguous, too; they essentially choose a longitude, and many longitudes cross through multiple cities with different time-keeping rules. The CLI chooses a region and city based on the current customer base. After you set this number, you can use show clock to see the time zone that corresponds to the offset.

Default(s) None

Guidelines A drastic clock change may have an adverse effect on policy or a number of other time-dependent applications. The CLI issues a warning before changing the time; enter **yes** to continue.

Use the show clock command to view the current time, date, and time zone settings.

Samples

```
bstnA# clock timezone America New_York
Changing the time may have an adverse impact on the switch.

Are you sure? [yes/no] yes
```

sets the time zone to that of New York City in the U.S.

```
bstnA# clock timezone EDT
...
```

sets the time zone to Eastern Daylight Time.

```
bstnA# clock timezone +0900
...
```

sets the time zone to 9 hours East of the Prime Meridian.

Related Commands show clock

dual-reboot

Purpose Use the dual-reboot command to simultaneously reboot both peers in a redundant pair. Use this command only on the advice of F5 Support.

Mode priv-exec

Security Role(s) network-engineer, storage-engineer, or crypto-officer

Syntax **`dual-reboot`**

Default(s) None

Guidelines

Important

This causes a service outage.

The CLI prompts for confirmation before rebooting both peers. Enter **yes** to proceed.

This command does not work on a standalone switch, where you can use reload to reboot.

Sample

```
prtlndB# dual-reboot

This command initiates a simultaneous reboot of both chassis in the
redundant pair.

Do you wish to proceed? [yes/no] yes

06-16 16:02:36 Rebooting: Dual reboot request was issued
...
```

Related Commands reload

hostname

Purpose A hostname is set for the ARX as part of the initial boot process. Use the hostname command to rename the ARX.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax **`hostname`** ***`name`***

name (up to 32 characters) is the hostname that you choose for the ARX. Use only alpha-numeric characters (0-9, a-z, A-Z), hyphens (-), and/or periods (.), as specified in RFC 1035.

Default(s) None

Guidelines The hostname appears in the CLI prompt; consequently, this command changes the prompt. All ARXes in your RON should have unique hostnames. The hostname is case-sensitive: "BSTNA" is a different hostname than "bstnA."

This is *not* the fully-qualified domain name (FQDN) for the ARX; use the ip domain-list command to set one or more domains for the switch.

To create a login message for administrators, use the login-banner command.

Sample

```
bstnA(cfg)# hostname usaax11
usaax11(cfg)#
```

changes the name of the current switch to "usaax11."

Related Commands

show hostname
ip domain-list
login-banner

login-banner

Purpose

After an administrator accesses the CLI or the GUI, a configurable text string appears. You can use the login-banner command to create or edit this text string.

The no login-banner command removes the text string.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`login-banner post-auth {`*`message`*` | configs `*`msg-file`*`}`
`no login-banner`

message (1-2000 characters) is the string to show after a successful authentication. Quote this string if it contains any spaces.

configs *msg-file* (1-255 characters) selects a file in the configs directory and uses its text. Use the show configs command for a list of files in the configs directory, and you can use the copy command to download a file to that directory. This file should be no more than 768 bytes.

Default(s)

None

Guidelines

The login banner can contain access rules and/or a broadcast message to all ARX administrators.

In a redundant pair of ARX peers, the login banner is unique for each chassis; use this command on both peers if you want a login banner in each.

Samples

```
bstnA(cfg)# login-banner post-auth “Running-config last saved 1/7 by J. User”
```

uses an informational message as a login banner.

```
bstnA(cfg)# login-banner post-auth configs banner.txt
```

uses a file in the configs directory, “banner.txt.”

Related Commands

hostname

probe metalog latency

Purpose The ARX-VE installation requires a “datastore” where the ARX stores important files, including its metalog data. Namespace software records important information in the metalog, possibly to be used later for a failure recovery. Namespace performance may suffer unless the latency to this metalog datastore is 400 micro-seconds or less. Use the probe metalog latency command to measure the latency between the ARX-VE and its datastore.

Mode priv-exec

Security Role(s) crypto-officer or storage-engineer

Syntax **probe metalog latency**

Default(s) None

Valid Platforms ARX-VE only

Guidelines Use this command after installing the ARX-VE, to confirm that the latency between the ARX-VE is sufficiently small to ensure good performance. This runs a series of I/O operations and measures the latency from each, then it shows the following values:

I/O Count is the number of I/O operations that the CLI ran.

Latency (usecs) is the average latency for all of those I/O operations, in micro-seconds.

If the latency is consistently greater than 400 micro-seconds, we recommend re-installing the ARX-VE and choosing another datastore. 250 micro-seconds or less is recommended for best performance, if that latency is feasible. The ARX-VE Installation Guide contains instructions for installing the ARX-VE.

The show redundancy metalog command shows similar statistics, but it applies to all platforms except the ARX-VE. The ARX-1500 and ARX-2500 store their metalog data on their internal disks. On those platforms, you can use the show metalog usage command to see the usage statistics for the internal metalog driver.

On all platforms, the namespace software keeps metalog read/write statistics. You can use the show statistics metalog command to see these metalog-usage statistics from a namespace-software perspective.

Sample

```
stkbrgA# probe metalog latency

 Statistics:

 I/O Count    Latency (usecs)
 ---------    ---------------
 1000         340
```

measures and displays the latency between the current ARX-VE and its metalog datastore. This latency is acceptable.

Related Commands show redundancy metalog
show metalog usage
show statistics metalog

reload

Purpose

Use the reload command to reboot the ARX. The ARX then comes up with the current "armed" software release. (By default, the "armed" release is the currently-running release; you can change it to a new release with the boot system command.)

Mode

priv-exec

Security Role(s)

network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

`reload [collect-diags] [reason]`

collect-diags (optional) causes a diagnostic-collection process before the reload. This is useful for capturing the state of the ARX in case of a failure that requires further diagnosis.

reason (optional, 1-255 characters) is a comment you can enter to appear in syslog. This is useful when reviewing the logs during testing and diagnostics.

Default(s)

None

Guidelines

The following prompt appears for confirmation:

```
Reload the entire chassis? [yes/no/diags]
```

Enter **yes** to reboot.

Enter **diags** if you want to collect diagnostic information and then reboot. This is equivalent to using the **collect-diags** option in the command. The collected information is a subset of the information captured by the standard collect state command. After the ARX reloads, use collect diags to gather other relevant information (along with the data collected before reboot) and upload the full diagnostic package to F5 Support.

To power off the system and prepare it for a planned power outage, you can use the shutdown command.

Guidelines: Other Applications

This command can be the final step for upgrading software on the switch. Use reload to restart all modules after you use boot system to arm the switch with a new release file. You can use show boot or show version to verify that the switch is armed with the desired release file.

If someone deleted the full configuration (with delete startup-config) before running this command, all configuration changes are lost after the ARX reloads. The CLI warns you of this and prompts for confirmation; you can enter **no** and run restore startup-config to preserve your configuration. To remove all configuration parameters, enter **yes**.

Samples

```
bstnA# reload

Reload the entire chassis? [yes/no/diags] yes
...
```

reloads the ARX, perhaps to complete a software upgrade.

```
bstnA# reload collect-diags

Reload the entire chassis? [yes/no] yes
...
```

collects diagnostic-state information and then reloads the ARX. After the reboot, this next command uploads that state information along with additional diagnostics:

```
bstnA# collect diag-info ftp://jpublic:jpwd@ftp.wwmed.com/diags

Collect diagnostic information? [yes/no] yes
...
```

Related Commands

collect
boot system
show boot
shutdown
restore startup-config

resource-profile

Purpose

The *resource profile* of an ARX is the allocation of its processes amongst its CPU cores, where each CPU core is dedicated to one process type: system processes, fastpath processes, or volume-group processes. Software Release 6.2.0 introduced an optimized resource profile for the ARX-2500. An ARX-2500 with an earlier software release retains the legacy profile. The new, optimized profile is recommended. Use the no resource-profile legacy command to upgrade the resource profile on the current chassis.

On the advice of F5 personnel only, you can use resource-profile legacy to return to the legacy profile.

Mode

cfg

Security Role(s)

crypto-officer or network-engineer

Syntax

```
no resource-profile legacy
resource-profile legacy
```

Default(s)

`resource-profile legacy`

after an upgrade from a pre-6.2.0 release.

`no resource-profile legacy`

for any ARX-2500 shipped with 6.2.0 or a later release.

Valid Platforms

ARX-2500 only

Guidelines

After you use this command, you must reload the ARX-2500 for the change to take effect. If you have a redundant pair of ARX-2500 devices (see redundancy), run this command on both peers and then use the dual-reboot command to reboot both of them at once.

The above is also true if you replay a running-config script with the resource-profile legacy setting. (One method of replaying a running-config is to save the file on the ARX-2500 and use the run command.) After replaying the config script, you must reload the ARX-2500 for resource-profile legacy to take effect.

The optimized resource profile assigns 3 cores, each running on separate hardware, to fastpath processes. It assigns 3 more to volume-group processing, and the remaining 2 cores to system processes.

The legacy resource profile assigns 4 cores to fastpath processes, 2 on one hardware module and 2 on another. The cores share their hardware resources in this configuration. The legacy profile also assigns 2 more to volume-group processes and the final 2 cores to system processes.

You can use the show processors command to see the current resource profile on your chassis.

Sample

```
stoweA(cfg)# no resource-profile legacy
stoweA(cfg)# end
stoweA# reload

Reload the entire chassis? [yes/no/diags] yes
...
```

upgrades the current ARX-2500, “stoweA,” to an optimized resource profile. This resource profile is generally recommended.

Related Commands

reload
dual-reboot
show processors

show baudrate

Purpose Use the show baudrate command to show the baud rate for the Console (serial) interface.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show baudrate`

Guidelines The serial interface is the port labeled “Console” on the front panel.

Sample

```
bstnA# show baudrate
The serial port baud rate is 9600.
```

Related Commands

show chassis

Purpose Use the show chassis command to display chassis, disk, slot, module, port, and/or temperature information.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show chassis [chassinfo | private-subnet | diskuse | slotinfo |
moduleinfo | temperature | metalog | nvr | summary]
```

chassinfo | ... | nvr is an optional choice to focus on one group of chassis tables. If you omit this, the command shows all groups.

chassinfo - specifies chassis-related tables (as opposed to slot- or module-related) with high-level environment, disk, and power supply information.

private-subnet - displays the private subnet that the switch uses for inter-process communication in the chassis. Every chassis in the RON (see ron tunnel) must have a unique private subnet.

diskuse - displays disk usage and directories on the system drive.

slotinfo - focuses on the slots.

moduleinfo - shows model information and firmware versions for each module.

temperature - specifies a table of sensor temperatures.

metalog - is only available on the ARX-1500 and ARX-2500. This option shows the statistics for metalog-driver software on this chassis. The metalog driver records important recovery data for namespace software, and (in a redundant pair) copies the data to the peer ARX.

nvr - is only available on the ARX-500, ARX-2000, and ARX-4000. This option shows the NVRAM-battery state and whether or not there are any ECC errors. The ARX-1500 and ARX-2500 store their metalog data on the RAID instead of the NVRAM; use the **metalog** option (above) to monitor the metalog-protection mechanism for those device types.

summary - shows a summary of the chassis state.

Guidelines The show chassis command displays several tables. You can focus on a smaller group of tables by using one of the optional keywords. Each group of tables is described in its own section, below.

These tables only appear if they are relevant to the current chassis type.

Guidelines: chassinfo

The chassinfo section contains the following tables:

Identification — the hostname and the Universally-Unique ID (UUID) for the chassis. Use hostname to change the hostname. You can only change the UUID through the initial-startup script.

Chassis shows the chassis type, a model number, hardware version (where applicable), and the serial number.

Chassis Environment contains the base MAC address for all modules (for a chassis with multiple modules, the individual modules' MACs appear in moduleinfo), power status, fan status, and chassis temperature:

- Power can be "Online," "Online Partial" (one power supply is working, the other is absent), or "Failed Partial" (one is working, the other failed).
- Fan(setting) shows the fan status and its speed setting. The fan status can be "Online," "Not Present," "Fault" (the fan tray failed), or "Ctr Fail" (the control to the fan-tray failed). The speed setting is in parentheses, and is "high," "medium," "low," or "Unk" (unknown).
- The temperature setting(s) next are customized for certain chassis types:
  - System Temp., CPU Temp. and CPU show the ambient chassis temperature, the temperature of the CPU chip, and the current CPU speed. This appears only for the platforms that support both ambient and CPU temperature readings (such as the ARX-2500). The System Temp can be "Normal," "Too high," or "Failed" (temperature monitoring failed), followed by both temperatures in Celsius.
    The CPU speed is relevant because these chassis types automatically reduce the CPU speed if the temperature is too high. This shows the current CPU speed followed by an indication of whether or not the speed has been reduced; "(Normal)" means that the CPU is running at 100%, a percentage (such as "85%") indicates that the CPU has been throttled back.
  - Temperature appears for other chassis types. This is the same as the System Temp. described above.

Power Details appears for the ARX-1500, ARX-2000, ARX-2500, and ARX-4000, which have redundant power supplies. This has one row for each power supply showing the power-supply state ("Online," "Absent," or "Failed").

The ARX-1500 and ARX-2500 each have redundant power supplies, A and B. They are accessible from the back of the chassis, on the left side. Each power supply has a plug for its power cable.

The ARX-2000 also has redundant power supplies, 1/1 (on top) and 1/2 (under 1/1). They are accessible from the back of the chassis, on the right side. Each power supply has a plug for its power cable.

The ARX-4000 has redundant power supplies in both the control plane and the dataplane; these are numbered 1/1 and 1/2 for the control plane (in the top half of the chassis) and 2/1 and 2/2 for the data plane (on the bottom half). All power supplies are accessible from the back of the chassis, on the right side. Each power supply has a plug for its power cable.

Guidelines: private-subnet

This shows the VLAN (where applicable; for example, the ARX-VE uses only VLAN 1 and the ARX-1500 and ARX-2500 have no private VLAN, so this does not appear for those platforms), Subnet, and Subnet Mask for the private subnet on this chassis. You set this when you install the switch. You can reset it with ip private vlan or ip private subnet reassign. This subnet must be unique for all switches in a Resilient Overlay Network (RON); see ron tunnel for more information about a RON.

Guidelines: diskuse

The first three tables only appear for a chassis with replaceable disks (any platform except the ARX-500 or the ARX-VE). The SATA Drive Details table appears for the ARX-500. The Disk Usage table appears for all chassis types.

Logical Disk Details shows the status of the RAID as a single, logical disk.

Disk is always 1. This represents the single, logical disk comprised of the disks in the table below.

Status is "Optimal" (both disks are working), "Degraded" (one disk failed or is degraded), "Verifying *n*%" (someone issued the raid verify command to verify disk integrity; the percentage shows the progress of the verification test), or "Failed."

Verification Mode is "Manual" or "Automatic," as set by the raid verification-mode command.

Verification Rate is the percentage of CPU that the RAID can use for verifying a disk. You can change this for manual verifications with raid verification-rate.

Disk Details shows the location and size of each disk, disk state, the data-transfer rate, and the model number.

Disk indicates the location of the physical disk: Bay 1 is on top of Bay 2 in all chassis types with replaceable disks.

Size is the full capacity of the disk.

State is "Online," "Degraded," (the disk may fail soon), "Rebuild *n*%" (someone used raid rebuild to add the disk to the RAID; the percentage shows the progress of the rebuild), "Not Present," "Failed," or "Unknown."

Transfer Rate is the current throughput on the disk's bus channel. This is negotiated at startup between the disk controller and the disks themselves. The next table shows the maximum rate allowed by the controller.

Model is the specific model number of the disk drive.

RAID Controller Details displays the current RAID settings.

Rebuild Rate is the percentage of CPU that the RAID can use for rebuilding a disk. You can change this with raid rebuild-rate.

Max Transfer Rate is the maximum throughput allowed on the disk's bus channel.

Firmware is firmware version running on the disk controller.

RAID Alarm is "Enabled" or "Disabled." You can disable the alarm with the raid silence command.

SATA Drive Details appears for an ARX-500 with a Serial ATA (SATA) disk drive.

Firmware is firmware version running on the disk.

Model and

Serial identify the disk drive.

Disk Usage shows the switch's software directories and their total space (in MB), used space (MB), free space (MB), and used percentage (%).

Guidelines: slotinfo

Slot Environment, the only table in the slotinfo group, shows each slot's contents and status:

- Slot is the location of the module in the chassis. This is1 (the control plane on the top half) or 2 (the data plane on the bottom) for the ARX-4000. It is always 1 for the ARX-500 or ARX-2000.
- Type is "ACM" (for the control plane on top) or "NSM" (for the data plane on the bottom) on the ARX-4000. For the ARX-500 or ARX-2000, the single slot contains an "ACM."
- State is the current state of the module. This is one of the module states documented for the show version command; see Guidelines: Module States, on page 5-54.
- Power can be "Online," "Degraded," or "Failed."
- Temperature is the aggregate of all temperature-sensor readings in the module. This is either "Normal" or "Too High," along with the current temperature range on the module.
- NVR Battery appears for the NSM on the ARX-4000 or the ACM on the ARX-500 or ARX-2000. This is the battery status for Non-Volatile RAM (NVRAM). The possible values in this field are "Good," "Missing," "Charging," "Testing," "Degraded," or "Failed." The NVRAM stores transaction logs from the ACM, which are critical for failure recovery. The battery provides extra protection for these logs. Contact F5 Support if the status is "Missing," "Degraded," or "Failed;" you may need a battery upgrade.
- Drive appears only for the ACM. This shows the drive type (LSI, IDE, or SATA) and status for the internal disk(s). The possible status values are "Good," "Degraded," "Error," or "Unknown." If the status is "Degraded" or "Error," replace the drive (for an ARX-1500, ARX-2000, ARX-2500, or ARX-4000) and/or contact F5 personnel (for an ARX-500). An "Unknown" status would likely be caused by an error in the drive-detection software.

The ARX-VE has no hardware, so this option generates no output when you run it on that platform. This option also generates no output when you run it on the ARX-1500 or the ARX-2500.

Guidelines: Port Media Details Table

Port Media Details only appears in the summary view (without any options) for the platforms that support 10G optics. These are the ARX-2500 and the ARX-4000. Each of these platforms supports up to two transceivers. Each row shows the following details about one of the transceivers:

Slot/Port, also identified on the front panel, identify the transceiver's port,

Type, is the particular transceiver model,

Vendor shows the manufacturer of the transceiver, and

Status, which is Good, Bad, Present, Not Present, Absent, or Unknown.

Guidelines: moduleinfo

The VM Information table only appears for the ARX-VE, which is a Virtual Appliance (VA, which is similar to a VM):

vMAC Address is the virtual MAC address assigned to the ARX-VE.

Adapter describes the network controller used by the ARX-VE.

Module is the heading for several tables with details of each module.

The first table describes the hardware in each module:

Slot is the module's slot number.

Ports is the number of ports on the module.

Procs is the number of processors.

Card is the card type: SCM_40, NSM, NSM_TX/FX, ASM/ASM_FC, or ACM. This does not appear on the ARX-VE.

Hardware only appears for the ARX-VE. This describes the CPU and memory resources reserved for the ARX-VE.

Xeon is the CPU speed for all Xeon processors (used in every module except the NSM) and memory that is allocated to each Xeon processor. This does not appear on the ARX-VE.

Sibyte is the CPU speed for all SiByte processors (used on the NSM and standalone ACMs) and the memory allocated to each SiByte complex. Each SiByte complex has two MIPS-processor cores. This does not appear for the ARX-VE, ARX-1500, or ARX-2500; none of these systems contain SiByte processors.

Serial is the module's serial number. This does not appear for the ARX-VE, ARX-1500, or ARX-2500; for those systems, refer to the overall chassis serial number in the show chassis chassinfo output.

The second table shows the range of MAC addresses assigned to the module and its current hardware revision (the major-revision letter and minor-revision number, along with any Rework or Deviation on the module). For the ARX-1500 and ARX-2500, this shows the BIOS version. This table does not appear for the ARX-VE.

The final tables do not appear for the ARX-VE, ARX-1500, or ARX-2500.

The third table displays the revisions for the Complex Programmable Logic Devices (CPLDs) on the module, the Reset CPLD, the Keeper CPLD, the Power CPLD (ARX-4000 only), and the MUX CPLD (also ARX-4000 only):

- The Reset CPLD coordinates hardware resets when the chassis powers up.
- The Keeper CPLD, in conjunction with an FPGA, ensures that all client-transaction logs are stored safely in NVRAM before the chassis powers down.
- The Power CPLD (ARX-4000 only) coordinates the distribution of power from multiple power supplies to the chassis components.
- The MUX CPLD (ARX-4000 only) multiplexes status signals from various internal components and sends a unified signal to the NVR FPGA (described below). This also activates/deactivates various LEDs on the external panels.
- The BIOS Version appears for the control plane (or ACM) on the ARX-500 and the ARX-4000.

Guidelines: moduleinfo (cont.)

The next table shows versions for bootstrap, diagnostics, and boot-loader firmware. These components manage the module's boot process: the bootstrap software starts the hardware, then the diagnostics run to verify the hardware functions, and finally the boot-loader loads and starts all of the system software. These come bundled with each software release; to install the versions that came with the latest release, use the firmware upgrade command.

The final table presents FPGA information, which varies from platform to platform:

ARX-4000 and ARX-2000 - the table contains the versions of the LBA FPGA and the NVR FPGA. The LBA (Load-Balancing Algorithm) FPGA merges inbound packets from the external interfaces to the NSM cores. The NVR FPGA manages the NVR, monitors chassis status through the MUX CPLD, and runs a watchdog process that reboots the chassis in the event of a serious failure.

ARX-500 - the versions for one FPGA and the boot-related firmware on the NSM processors. The FPGA is a hybrid of the NVR and MTL FPGAs.

Guidelines: temperature

Temperature Details shows temperature-sensor information from each module.

For the ARX-1500 and ARX-2500, the table shows the following fields:

Slot is always "1."

Module is always "ACM."

System Temp. shows the ambient chassis temperature in Celsius.

CPU Temp. shows the temperature of the CPU chip in Celsius.

CPU shows the current CPU speed. These chassis types automatically reduce the CPU speed if the temperature is too high. This field shows the current CPU speed followed by an indication of whether or not the speed has been reduced; "(Normal)" means that the CPU is running at 100%, a percentage (such as "85%") indicates that the CPU has been throttled back.

Fan Status shows the fan status and its speed setting. The fan status can be "Online," "Not Present," "Fault" (the fan tray failed), or "Ctr Fail" (the control to the fan-tray failed). The speed setting is in parentheses, and is "high," "medium," "low," or "Unk" (unknown).

For all other hardware platforms, each processor has one temperature sensor positioned near the processor. The sensor takes its own "local" reading for the ambient temperature of the processor, and it takes a "remote" reading from the processor itself. The remote reading is the temperature of the processor's silica.

Guidelines: metalog

Metalog Usage shows the statistics for metalog usage on this chassis.

Status is the status of the metalog driver software, which writes metalog packets to an internal NFS export and duplicates them to a twin driver in the redundant peer (if there is one). The possible options are "Standalone" (there is no redundant peer), "Active" (the driver is duplicating all metalog packets to the redundant peer), and "Standby" (the driver is receiving all metalog packets from the Active peer).

Statistics is a table with metalog-usage statistics:

- I/O Count is the number of metalog packets sent to local storage (and remote storage, if this is the Active chassis in a redundant pair).
- Retransmit Count only increments on the Active peer in a redundant pair. This is the number of metalog transmissions to the Standby peer that required a retry.
- Hourly Latency (usecs) is the average latency in micro-seconds between the ARX and its local metalog storage. This number is updated once per hour. On the "Active" peer, this is the longer of two events that occur in parallel: writing the data to the local partition and writing the same data to the Standby peer's partition.

You can use show metalog usage to show all of the hourly metalog statistics since the last ARX reboot. The clear metalog usage command clears these statistics and restarts them at 0 (zero).

This output only appears on a chassis with an internal metalog driver, but namespace software keeps metalog read/write statistics on all chassis types. On any platform, you can use the show statistics metalog command to see these read/write statistics from a namespace perspective.

Guidelines: NVR

The NVR table shows the following information about the state of the hardware that protects Non-Volatile RAM:

- NVR Battery is the battery status for Non-Volatile RAM (NVR or NVRAM). The possible values in this field are "Good," "Missing," "Charging," "Testing," "Degraded," or "Failed." The NVRAM stores transaction logs from the control plane, called *metalog* data, which are critical for failure recovery. The battery provides extra protection for these logs in the event of a power outage.
- ECC State is the state of the Error-Correction Circuitry (ECC). The ECC checks all data as it is read from NVRAM, and generates errors if it detects any corruption. This can be "No Error," "Pending," "Non-Correctable Error," or "Unknown." A "Non-Correctable Error" state is serious; contact F5 personnel if you see this.
- NVR Size (MB) is the size of the NVRAM region, in MegaBytes.

The ARX-VE has no NVRAM hardware, so this option generates no output when you run it on that platform. The **nvr** option does not even appear on the ARX-1500 and ARX-2500, which also do not support NVRAM hardware.

Samples

Figure 5.1 on page 5-26, Figure 5.2 on page 5-27, Figure 5.3 on page 5-30, Figure 5.4 on page 5-31, Figure 5.5 on page 5-31, and Figure 5.6 on page 5-33 show sample output for the show chassis command on the ARX-4000, ARX-2500, ARX-1500, ARX-VE, ARX-2000, and ARX-500 respectively.

Related Commands

show processors

Figure 5.1 *Sample Output: show chassis (ARX-4000)*

```
bstnA> show chassis

 Identification:
 Hostname                              UUID
 ------------------------------------ --------------------------------------
 bstnA                                 d9bdece8-9866-11d8-91e3-f48e42637d58

 Chassis:
 Chassis Type  Model Number                          HW Ver.  Serial Number
 ------------  ------------------------------------  -------  -------------
 ARX-4000      SR2500ALLXR                                    BZDS80701617

 Private Subnet:
 VLAN   Subnet             Subnet Mask
 ----   ----------------   ----------------
 1010   169.254.66.0       255.255.255.0

 Chassis Environment:
 Base MAC Address    Power           Fan(setting)   System Temp.   CPU Temp.
 -----------------   --------------  -------------  -------------  ------------
 00:0a:49:17:9e:00   Online Partial  Online (Unk)   Normal(<62 C)

 Power Details:
 Supply   State
 ------   -----
 1/1      Failed
 1/2      Online
 2/1      Failed
 2/2      Online

 Logical Disk Details:
 Disk    Status                Verification Mode  Verification Rate
 ------  --------------------  -----------------  -----------------
 1       Degraded              Manual             50 %

 Disk Details:
 Disk       Size          State               Transfer Rate   Model
 --------   -----------   -----------------   -------------   -----------------
  Bay 1     136.73G       Failed              3.0Gb/sec       ST3146855SS
  Bay 2     136.73G       Online              3.0Gb/sec       ST3146855SS

 RAID Controller Details:
 Rebuild Rate  Max Transfer Rate  Firmware       RAID Alarm
 ------------  -----------------  -------------  ------------
 85 %          3.0Gb/sec          7.0.1-0075     Enabled

 Slot Environment:
 Slot  Type   State             Power    Temperature     NVR Battery  Drive
 ----  -----  ----------------  -------  --------------  -----------  ---------
  1    ACM    Online            Degraded Normal          N/A          LSI Degraded
  2    NSM    Online            Degraded Normal          Good

 Module:
 Slot Ports Procs Card      Xeon               Sibyte             Serial
 ---- ----- ----- --------  -----------------  -----------------  ------------
  1   0     1     ACM       2.6 GHz 16128 MB   N/A                BZDS80701617
  2   14    12    NSM       N/A                900 MHz 4096 MB    006046

 Slot MAC Address                   HW Version  Rework  Deviation
 ---- ----------------------------  ----------  ------  --------------------
  1   000A49179E00 to 000A49179E3F
```

```
 2   000A49179E40 to 000A49179E9F  B 0

Slot Reset   Keeper   Power    Mux   BIOS Version
     CPLD    CPLD     CPLD     CPLD
---- ------- ------- ------- ------- -----------------------------------
 1   N/A     N/A     N/A     N/A     S5000.86B.10.00.0094.101320081858
 2   14      3       5       4

Slot Boot Version      Diag Version      BootLdr Version
---- ----------------  ----------------  ----------------
 1   N/A               N/A               N/A
 2                     5.02.000.12627    5.02.000.12627

            FPGA Version
Slot       LBA          NVR
---- -------------- ---------------
 1   N/A             N/A
 2   65              71

Port Media Details:
Slot/Port   Type                                Vendor            Status
---------   --------------------------------    ---------------   ----------------------
 2/1        10GBASE-SR X2                       Intel             Good
 2/2        10GBASE-SR X2                       Intel             Bad

Disk Usage:
Name                             Total MB    Used MB     Free MB     Used%
-------------------------------  ----------  ----------  ----------  -----
System                           2331        1754        459         80%
Releases                         5285        3402        1614        68%
Logs                             54951       125         52033       1%
Cores; DiagInfo; Lists           21133       342         19717       2%
Scripts                          3172        67          2944        3%
Reports                          8458        36          7991        1%

Temperature Details:

               Sensor 1 (C)   Sensor 2 (C)   Sensor 3 (C)   Sensor 4 (C)
Slot  Module   Local Remote   Local Remote   Local Remote   Local Remote
----  ------   ----- ------   ----- ------   ----- ------   ----- ------
  1   ACM      31    22       N/A   N/A      N/A   N/A      N/A   N/A
  2   ACM      28    46       29    42       39    46       N/A   N/A

NVR:
NVR Battery  ECC State               NVR Size (MB)
-----------  ----------------------  -------------
Good         No Error                2048
```

Figure 5.2 *Sample Output: show chassis (ARX-2500)*

```
stoweA> show chassis

 Identification:
 Hostname                           UUID
 ---------------------------------- ------------------------------------
 stoweA                             05d5a0fa-f2fb-11df-8daf-af50d57e388e

 Chassis:
 Chassis Type  Model Number                          Serial Number
 ------------  -----------------------------------  -------------
 ARX-2500      ARX2500HE-F5                         XX-ABCD-0447
```

```
Private Subnet:
Subnet             Subnet Mask
----------------   ----------------
169.254.96.0       255.255.255.0

Chassis Environment:
Base MAC Address    Power           Fan(setting)   System Temp.   CPU Temp.  CPU
-----------------   --------------  -------------  -------------  ---------  ----------------
00:0a:49:75:44:00   Online          Online (high)  Normal 29 C    36 C       2.6 GHz (Normal)

Power Details:
Supply   State
------   -----
  A      Online
  B      Online

Logical Disk Details:
Disk    Status                 Verification Mode  Verification Rate
------  ---------------------  -----------------  -----------------
1       Optimal                Automatic          10 %

Disk Details:
Disk       Size          State               Transfer Rate   Model
--------   -----------   -----------------   -------------   -----------------
 Bay 1     136.91G       Online              3.0Gb/sec       HUC103014CSS600
 Bay 2     136.91G       Online              3.0Gb/sec       HUC103014CSS600

RAID Controller Details:
Rebuild Rate  Max Transfer Rate  Firmware       RAID Alarm
------------  -----------------  -------------  ------------
90 %          3.0Gb/sec          5.2-0[17945]   Enabled

Module:
Slot Ports Procs Card      Xeon               ECC State
---- ----- ----- --------  -----------------  -----------------
 1   6     1     ACM       2.6 GHz 16128 MB   Good

Slot MAC Address                   BIOS Version
---- ----------------------------  ------------------------------------
 1   000A49754400 to 000A497544FF  080015.20110308.1.3.00004

Port Media Details:
Slot/Port   Type                                Vendor            Status
---------   --------------------------------    ---------------   ----------------------
 2/1        OPT-0016(Short Reach)               Finisar           Present
 2/2        OPT-0016(Short Reach)               Finisar           Present

Disk Usage:
Name                              Total MB    Used MB     Free MB     Used%
-------------------------------   ----------  ----------  ----------  -----
System                            3173        1813        1198        61%
Releases                          6345        3434        2588        58%
Logs                              52838       217         49936       1%
Cores; DiagInfo; Lists            21133       3959        16100       20%
Scripts                           3172        68          2943        3%
Reports                           8458        34          7994        1%

Metalog Usage:

Status: STANDALONE

Statistics:
```

```
I/O Count           Retransmit Count    Hourly Latency (usecs)
----------------    ----------------    ----------------------
7545                0                   36
```

Figure 5.3 *Sample Output: show chassis (ARX-1500)*

```
canbyA> show chassis

 Identification:
 Hostname                            UUID
 ----------------------------------- -------------------------------------
 canbyA                              64a6417e-cc3d-11df-80ca-a73fbeb72ef8

 Chassis:
 Chassis Type  Model Number                         Serial Number
 ------------  ----------------------------------   -------------
 ARX-1500      ARX1500LE-F5                         XX-ABCD-0446

 Private Subnet:
 Subnet             Subnet Mask
 ----------------   ----------------
 169.254.90.0       255.255.255.0

 Chassis Environment:
 Base MAC Address    Power           Fan(setting)   System Temp.   CPU Temp.     CPU.
 -----------------   --------------  -------------  -------------  ------------  ----------------
 00:0a:49:75:41:00   Online          Online (high)  Normal 30 C    39 C         2.6 GHz (Normal)

 Power Details:
 Supply   State
 ------   -----
   A      Online
   B      Online

 Logical Disk Details:
 Disk    Status                Verification Mode  Verification Rate
 ------  --------------------  -----------------  -----------------
 1       Optimal               Automatic          10 %

 Disk Details:
 Disk       Size          State               Transfer Rate   Model
 --------   -----------   -----------------   -------------   -----------------
  Bay 1     136.91G       Online              3.0Gb/sec       HUC103014CSS600
  Bay 2     136.91G       Online              3.0Gb/sec       HUC103014CSS600

 RAID Controller Details:
 Rebuild Rate  Max Transfer Rate  Firmware       RAID Alarm
 ------------  -----------------  -------------  ------------
 90 %          3.0Gb/sec          5.2-0[17945]   Enabled

 Module:
 Slot Ports Procs Card      Xeon               ECC State
 ---- ----- ----- --------  -----------------  -----------------
  1   8     1     ACM       2.6 GHz 8192 MB    Good

 Slot MAC Address                  BIOS Version
 ---- ---------------------------  ----------------------------------
  1   000A49754100 to 000A497541FF  080015.20110308.1.3.00004

 Disk Usage:
 Name                              Total MB    Used MB     Free MB     Used%
 --------------------------------  ----------  ----------  ----------  -----
 System                            3173        1813        1199        61%
 Releases                          6345        3523        2500        59%
 Logs                              52838       67          50086       1%
 Cores; DiagInfo; Lists            21133       485         19574       3%
 Scripts                           3172        62          2949        3%
```

```
 Reports                           8458        34          7994        1%

 Metalog Usage:

 Status: STANDALONE

 Statistics:

 I/O Count           Retransmit Count    Hourly Latency (usecs)
 ----------------    ----------------    ----------------------
 77332               0                   37
```

Figure 5.4 *Sample Output: show chassis (ARX-VE)*

```
stkbrgA# show chassis

 Identification:
 Hostname                             UUID
 ------------------------------------ --------------------------------------
 stkbrgA                              e9e786f6-cb13-11df-a230-7f4d2a0b939d

 Chassis:
 Chassis Type  Model                   Serial Number
 ------------  ----------------------  -------------
 ARX-VE        VMware                  7530570

 Private Subnet:
 Subnet             Subnet Mask
 ----------------   ----------------
 169.254.221.0      255.255.255.0

 VM Information:
 vMAC Address        Adapter
 ------------------  ------------------------------
 00:50:56:88:00:6e   VMXNET3 Ethernet Controller

 Module:
 Slot Ports Card      Hardware
 ---- ----- --------  --------------------------------------------------
  1    1    ACM       (1)  Intel(R) Xeon(R) CPU E5520 @ 2.27GHz, 2048 MB

 Disk Usage:
 Name                              Total MB    Used MB     Free MB     Used%
 --------------------------------  ----------  ----------  ----------  -----
 System                            3173        1684        1328        56%
 Releases                          6353        3128        2902        52%
 Logs                              7238        70          6799        2%
 Cores; DiagInfo; Lists            9508        39          8986        1%
 Scripts                           3169        48          2959        2%
 Reports                           6339        33          5983        1%
```

Figure 5.5 *Sample Output: show chassis (ARX-2000)*

```
prtlndA> show chassis

 Identification:
 Hostname                             UUID
 ------------------------------------ --------------------------------------
 prtlndA                              876616f6-79ac-11d8-946f-958fcb4e6e35

 Chassis:
 Chassis Type  Model Number                          HW Ver.  Serial
 ------------  ------------------------------------  -------  -------------
```

```
ARX-2000      SR2500ALLXR                                    BZDS75100014

Private Subnet:
VLAN   Subnet             Subnet Mask
----   ----------------   ----------------
1008   169.254.100.0      255.255.255.0

Chassis Environment:
Base MAC Address    Power           Fan(setting)   System Temp.   CPU Temp.
-----------------   -------------   ------------   ------------   -----------
00:0a:49:27:84:00   Online Partial  Online (Unk)   Normal(<62 C)

Power Details:
Supply   State
------   -----
1/1      Online
1/2      Unknown

Logical Disk Details:
Disk    Status                Verification Mode  Verification Rate
------  --------------------  -----------------  -----------------
1       Optimal               Automatic          10 %

Disk Details:
Disk    Size          State               Transfer Rate   Model
--------   -----------   -----------------   -------------   ------------------
 Bay 1     136.73G       Online              3.0Gb/sec       ST3146855SS
 Bay 2     136.73G       Online              3.0Gb/sec       ST3146855SS

RAID Controller Details:
Rebuild Rate  Max Transfer Rate  Firmware       RAID Alarm
------------  -----------------  -------------  ------------
90 %          3.0Gb/sec          7.0.1-0061     Enabled

Slot Environment:
Slot  Type   State             Power    Temperature     NVR Battery  Drive
----  -----  ----------------  -------  --------------  -----------  ---------
 1    ACM    Online            Degraded Normal          Missing      LSI Good

Module:
Slot Ports Procs Card      Xeon               Sibyte             Serial
---- ----- ----- --------  -----------------  -----------------  ------------
 1    13    5    ACM       2.3 GHz  12032 MB  1.2 GHz  4096 MB   0306063

Slot MAC Address                   HW Version  Rework  Deviation
---- ----------------------------  ----------  ------  --------------------
 1   000A4917AF00 to 000A4917AF9F  B 0

Slot Reset   Keeper   Power    Mux   BIOS Version
     CPLD    CPLD     CPLD     CPLD
---- ------- ------- ------- ------- -----------------------------------
 1   14      3       5       4       S5000.86B.10.00.0094.101320081858

Slot Boot Version      Diag Version      BootLdr Version
---- ----------------  ----------------  ----------------
 1                     5.01.000.11908    5.01.000.11908

            FPGA Version
Slot       LBA          NVR
---- -------------- ---------------
 1   66             71
```

```
 Disk Usage:
 Name                              Total MB    Used MB     Free MB     Used%
 --------------------------------  ----------  ----------  ----------  -----
 System                            2331        1404        808         64%
 Releases                          5285        2673        2343        54%
 Logs                              54951       180         51978       1%
 Cores; DiagInfo; Lists            21133       142         19917       1%
 Scripts                           3172        47          2963        2%
 Reports                           8458        33          7994        1%

 Temperature Details:
                Sensor 1 (C)   Sensor 2 (C)   Sensor 3 (C)   Sensor 4 (C)
 Slot  Module   Local Remote   Local Remote   Local Remote   Local Remote
 ----  ------   ----- ------   ----- ------   ----- ------   ----- ------
   1   ACM      30    23       25    39       N/A   N/A      N/A   N/A

 NVR:
 NVR Battery  ECC State               NVR Size (MB)
 -----------  ----------------------  -------------
 Missing      No Error                1024

 Battery Charger:
     Power          Sensor              Status
 ---------------- -------------- ----------------------
Charging Current  0 mA           Normal ( <1200 mA)
Pack Voltage      0 mV           Normal ( <17200 mV)
Cell Voltage      0 mV           Normal ( <4300 mV)
```

Figure 5.6 *Sample Output: show chassis (ARX-500)*

```
provA> show chassis

 Identification:
 Hostname                          UUID
 --------------------------------- ------------------------------------
 provA                             876616f6-79ac-11d8-946f-958fcb4e6e35

 Chassis:
 Chassis Type  Model Number                          HW Ver.  Serial
 ------------  ----------------------------------   -------  -------------
 ARX-500       SR1500ALR (Alcolu)                            BZDG80502043C

 Private Subnet:
 VLAN   Subnet             Subnet Mask
 ----   ----------------   ----------------
 1006   169.254.245.128    255.255.255.192

 Chassis Environment:
 Base MAC Address    Power
 -----------------   --------------
 00:0a:49:0f:58:00   Online

 SATA Drive Details:
 Firmware   Model                         Serial
 ---------- ---------------------------- --------------------
 04.04V02   WDC_WD1500HLFS-01G6U0         WD-WXL708114721

 Slot Environment:
 Slot  Type   State             Power    Temperature     NVR Battery  Drive
 ----  -----  ----------------  -------  --------------  -----------  ---------
  1    ACM    Online            Online   Normal ( <45 C) Good         SATA Good
```

```
Module:
Slot Ports Procs Card      Xeon               Sibyte             Serial
---- ----- ----- --------  -----------------  -----------------  ------------
 1    2     2    ACM       2.0 GHz  4096 MB   700 MHz  512 MB    3928

Slot MAC Address                   HW Version  Rework  Deviation
---- ----------------------------  ----------  ------  --------------------
 1   000A490F580F to 000A490F5810  C 6         0       0

Slot Reset CPLD Keeper CPLD BIOS Version
---- ---------- ----------- ----------------------------------------
 1   5          5                S5000.86B.10.00.0094.101320081858

Slot FPGA Version     NSM Boot Version  NSM Diag Version  NSM BootLdr Version
---- ---------------  ----------------  ----------------  -------------------
 1   macau      11    5.01.000.11898    5.01.000.11898    5.01.000.11898

Disk Usage:
Name                              Total MB    Used MB     Free MB     Used%
--------------------------------  ----------  ----------  ----------  -----
System                            2331        1408        804         64%
Releases                          5285        2646        2370        53%
Logs                              54951       58          52100       1%
Cores; DiagInfo; Lists            21133       69          19990       1%
Scripts                           3172        101         2910        4%
Reports                           8458        33          7994        1%

Temperature Details:
               Sensor 1 (C)   Sensor 2 (C)   Sensor 3 (C)   Sensor 4 (C)
Slot  Module   Local Remote   Local Remote   Local Remote   Local Remote
----  ------   ----- ------   ----- ------   ----- ------   ----- ------
  1   ACM      0              0              26    44       0

NVR:
NVR Battery  ECC State               NVR Size (MB)
-----------  ----------------------  -------------
Good         No Error                128
```

show clock

Purpose Use the show clock command to see the current time/date setting on the ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show clock`**

Guidelines The ARX has a separate clock in each module, synchronized internally through NTP. The clock display is from the ACM, which has the master clock for all processes in the switch.

Use the clock set command to reset the clock. Use the clock timezone command to reset the time zone. Use the ntp server command to synchronize the clock with an external NTP server.

Sample

```
stoweA# show clock
     Local time:  Mon 03 Nov 2008 01:48:45 AM EST -0500 America New_York
 Universal time:  Mon 03 Nov 2008 06:48:45 AM UTC
```

Related Commands

clock set
clock timezone
ntp server

show hostname

Purpose Use the show hostname command to see the ARX's hostname.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show hostname`**

Guidelines To set the hostname, use the cfg-mode hostname command.

Sample

```
bstnA(cfg)# show hostname
bstnA
bstnA(cfg)#
```

Related Commands hostname

show memory usage

Purpose The show memory usage command shows usage statistics for ARX memory.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show memory usage [report]
```

report (optional) causes the CLI to create a report instead of showing output on the command line. The CLI shows the report name after you type the command. The report is named "memory_usage_*yyyymmddHHMMSS*.rpt," where *yyyymmddHHMMSS* is the date and time when the report was created. Use show reports to list all reports and view the report's contents.

Platforms ARX-1500 and ARX-2500

Guidelines The output contains a table with 3-hour samples of memory usage, where the most-recent sample is at the top. Every row of the table shows the memory usage for one 3-hour sample:

Date
MM/DD is the month and day of the 3-hour sample.

Time Interval is the start time and end time for the 3-hour sample.

%Memory shows the low, average (Avg), and high percentages of memory used during the sample period.

Use the show processors usage command to show similar statistics for all processors on the switch; for other platform types, such as the ARX-2000, this also includes memory-usage statistics. The show system tasks command shows the currently-running tasks on one or more processors.

Samples

```
stoweA# show memory usage
```

shows the memory-usage statistics for all processors on an ARX-2500. See Figure 5.7 on page 5-38 for sample output.

```
canbyA# show memory usage report

Scheduling report: memory_usage_201103220344.rpt
canbyA#
```

sends the processor-usage statistics for an ARX-1500 to a report. See Figure 5.8 on page 5-38 for a sample report.

Related Commands show processors usage
show system tasks

Figure 5.7 *Sample Output: show memory usage*

```
stoweA# show memory usage

Memory Usage:

Date                     %Memory
MM/DD  Time Interval   Low  Avg  High
-----  -------------   --------------
03/22  01:00 - 03:27    11   12   12
03/21  22:00 - 01:00    11   11   11
```

Figure 5.8 *Sample Report: memory_usage...*

```
canbyA# show reports memory_usage_201103220344.rpt
**** Memory Usage Report: Started at Tue Mar 22 03:44:23 2011 ****
**** Software Version: 6.00.000.13561 (Mar 20 2011 22:55:54) [nbuilds]
**** Hardware Platform: ARX-1500
**** Report Destination:

Memory Usage:

Date                     %Memory
MM/DD  Time Interval   Low  Avg  High
-----  -------------   --------------
03/22  01:00 - 03:42    18   19   20
03/21  22:00 - 01:00    18   18   18

**** Total processed:              2
**** Elapsed time:           00:00:00
**** Memory Usage Report: DONE at Tue Mar 22 03:44:23 2011 ****
```

show metalog usage

Purpose Namespace software and network software each record important log information, called *metalog* data, to be shared with the ARX device's redundant peer. This metalog data facilitates a failover from the active ARX to its backup ARX. The ARX-1500 and ARX-2500 each store this data on internal disk partitions, one partition per volume group plus one partition per fastpath processor. The show metalog usage command shows usage statistics for these metalog-storage partitions.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show metalog usage [report]`

report (optional) causes the CLI to create a report instead of showing output on the command line. The CLI shows the report name after you type the command. The report is named "metalog_usage_*yyyymmddHHMMSS*.rpt," where *yyyymmddHHMMSS* is the date and time when the report was created. Use show reports to list all reports and view the report's contents.

Valid Platforms ARX-1500 and ARX-2500

Guidelines The output contains one table per metalog partition. At the top of each table is the following label:

`Metalog Usage for Disk Partition id used by {Volume-Group n | Fastpath slot.proc}`

Where

- Disk Partition *id* shows the ID of the internal-disk partition to store this metalog data.
- Volume-Group *n* is the volume group that stores its metalog information on partition *id*. The show volume-group command lists all volume groups on the chassis.
- Fastpath *slot.proc* indicates that the fastpath (or network) processor at *slot.proc* is using this internal partition to store its metalog data. The show processors command lists all processors on the chassis and shows which of them run the Fastpath (network) processes.

Guidelines (Cont.)

Each table shows 1-hour samples of memory usage, where the oldest sample is at the top. Every row of the table shows the metalog-partition usage for a single one-hour sample:

Date is the month and day of the sample.

Time Interval is the start time and end time for the sample.

I/O Count is the number of metalog packets sent to the local partition. If this ARX is in a redundant pair, this is the sum of the packets sent to the local partition and the duplicate packets sent to the partition on the backup peer. You can use the show redundancy command to check the redundancy status of this chassis.

Retransmit Count only increments on the active peer in a redundant pair. This is the number of metalog transmissions to the backup peer that required a retry.

Latency (usecs) is the minimum, maximum, and average latency in micro-seconds between the ARX and its local metalog partition. These numbers are updated once per hour. On the active peer in a redundant pair, this is the longer of two events that occur in parallel: writing the data to the local partition and writing the same data to the backup peer's partition.

The ARX keeps these statistics from the most-recent time it has booted. You can use the clear metalog usage command to clear these statistics immediately.

The show redundancy metalog command shows similar statistics, but it applies to all platforms (except the ARX-VE). To find similar statistics on the ARX-VE platform, use the probe metalog latency command.

On all platforms, the namespace software keeps metalog read/write statistics. You can use the show statistics metalog command to see these metalog-usage statistics from a namespace-software perspective.

Samples

```
stoweA# show metalog usage
```

shows the metalog-usage statistics for all processors on an ARX-2500. See Figure 5.9 on page 5-41 for sample output.

```
canbyA# show metalog usage report

Scheduling report: metalog_usage_201104110152.rpt
canbyA#
```

sends the metalog-usage statistics for an ARX-1500 to a report. See Figure 5.10 on page 5-41 for a sample report.

Related Commands

clear metalog usage
show processors
probe metalog latency
show statistics metalog
show redundancy metalog

Figure 5.9 *Sample Output: show metalog usage*

```
stoweA# show metalog usage

 Metalog Usage for Disk Partition 24 used by Fastpath 1.3

 Date  Time Interval I/O Count     Retransmit     Latency (usecs)
                                   Count          Min       Max       Avg
 ----- ------------- ------------- -------------  --------  --------  --------
 03/04 00:15-01:15   3             0              32        46        38
 03/04 01:15-02:15   0             0              0         0         0
 03/04 02:15-03:15   0             0              0         0         0
 03/04 03:15-03:57   0             0              0         0         0

 Metalog Usage for Disk Partition 25 used by Fastpath 1.4

 Date  Time Interval I/O Count     Retransmit     Latency (usecs)
                                   Count          Min       Max       Avg
 ----- ------------- ------------- -------------  --------  --------  --------
 03/04 00:15-01:15   3             0              31        42        35
 03/04 01:15-02:15   0             0              0         0         0
 03/04 02:15-03:15   0             0              0         0         0
 03/04 03:15-03:57   0             0              0         0         0

 Metalog Usage for Disk Partition 26 used by Fastpath 1.5

 Date  Time Interval I/O Count     Retransmit     Latency (usecs)
                                   Count          Min       Max       Avg
 ----- ------------- ------------- -------------  --------  --------  --------
 03/04 00:15-01:15   3             0              34        57        42
 03/04 01:15-02:15   0             0              0         0         0
 03/04 02:15-03:15   0             0              0         0         0
 03/04 03:15-03:57   0             0              0         0         0

 Metalog Usage for Disk Partition 27 used by Fastpath 1.6

 Date  Time Interval I/O Count     Retransmit     Latency (usecs)
                                   Count          Min       Max       Avg
 ----- ------------- ------------- -------------  --------  --------  --------
 03/04 00:15-01:15   3             0              32        46        36
 03/04 01:15-02:15   0             0              0         0         0
 03/04 02:15-03:15   0             0              0         0         0
 03/04 03:15-03:57   0             0              0         0         0
```

Figure 5.10 *Sample Report: metalog_usage...*

```
canbyA# show reports metalog_usage_201104110152.rpt
**** Metalog Usage Report: Started at Mon Apr 11 01:52:27 2011 ****
**** Software Version: 6.00.000.13571 (Apr  6 2011 20:15:54) [nbuilds]
**** Hardware Platform: ARX-1500
**** Report Destination:

 Metalog Usage for Disk Partition 0 used by Volume-Group 1

 Date  Time Interval I/O Count     Retransmit     Latency (usecs)
                                   Count          Min       Max       Avg
 ----- ------------- ------------- -------------  --------  --------  --------
 03/11 00:21-01:21   0             0              0         0         0
 03/11 01:21-01:51   59759         0              28        96547     1

 Metalog Usage for Disk Partition 24 used by Fastpath 1.3

 Date  Time Interval I/O Count     Retransmit     Latency (usecs)
```

```
                                  Count          Min       Max       Avg
 ----- ------------ ------------- -------------  --------  --------  --------
 03/11 00:21-01:21   3             0              33        50        39
 03/11 01:21-01:51   12            0              36        43        0

 Metalog Usage for Disk Partition 25 used by Fastpath 1.4

 Date  Time Interval I/O Count     Retransmit     Latency (usecs)
                                  Count          Min       Max       Avg
 ----- ------------ ------------- -------------  --------  --------  --------
 03/11 00:21-01:21   3             0              31        51        38
 03/11 01:21-01:51   0             0              0         0         0
**** Total processed:               0
**** Elapsed time:           00:00:00
**** Metalog Usage Report: DONE at Mon Apr 11 01:52:27 2011 ****
```

show processors

Purpose Use the show processors command to list the processors (CPUs) installed on the ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show processors
```

Guidelines: Output on the ARX-1500, ARX-2500, and ARX-VE

The ARX-1500 and ARX-2500 use a multi-core processor, where each core is primarily dedicated to either network (fastpath) processing, storage (volume-group) processing, or CLI/GUI (system-management) processing. The output for this command shows two tables to describe the current state of these processor cores.

The first table shows the role and CPU usage of each core, where each core appears on a separate row. The ARX-1500 uses fewer cores than the ARX-2500, and the ARX-VE has one:

Proc is in *processor.core* format. The *processor* is always 1, and the *core* number identifies the particular core.

Role is either "Fastpath" (network processing), "Volume-Group" (storage-related processing), or "System" (CLI, GUI, and other management processing). This indicates the type of process that primarily runs on this core.

CPU1M is the average CPU usage over the last 60 seconds. This number is a percentage.

CPU5M is the average CPU usage over the last 5 minutes. As above, this is a percentage.

The second table shows a summary for the entire processor. This table only appears on the ARX-1500 and ARX-2500:

Up Time is time since the last reboot. You can use the reload command to manually reboot. For the ARX-VE, you can also reboot from your VM-client console.

Memory (MB) is a heading for the following processor-memory measures, in MegaBytes:

- Total is the processor's total memory.
- Free is the processor's available memory.

Swap (MB) is a heading for similar swap-space measures, also in MegaBytes. This is space on the internal hard disk that is used as a memory region when free memory is low:

- Total is the processor's total swap space.
- Free is the processor's available swap space.

Guidelines: Output on Other Platforms

For the other platforms (not the ARX-1500, ARX-2500, or ARX-VE), the output is a table with one row per processor. Each row contains the following fields:

Proc shows the slot location of the processor in *slot.processor* format. Processor 1.1 is the management processor for all platforms.

Module is the processor's module type (ACM only on the ARX-500 or ARX-2000; ACM or NSM on the ARX-4000).

State is the current processor state:

- Up
- Reset - should appear very briefly, just before Boot.
- Boot - the processor is running its diagnostic tests, invoked during boot.
- Init - the processor passed its diagnostic tests and is being provisioned.
- Waiting - the processor is waiting for configuration parameters before it can begin processing. If all NSM and/or ACM processors are in the "Waiting" state, they require proper configuration for the interfaces. Minimally, each interface must be started with no shutdown (cfg-if-gig), and each NSM processor requires a proxy-IP address (see ip proxy-address).
- Downloading - if (during "Init") the processor discovers that it needs new software, it enters this state to fetch and install the software.
- Unknown - the CLI cannot ascertain the processor's state.
- Standby - appears for an NSM processor that failed over to a peer processor and then came back online. The peer processor is now handling this processor's traffic, and this processor is in a hot-standby state. If the peer processor fails, this processor takes control. The nsm recovery command configures the NSM processors to go into this state if they experience a failure. NSM processors also go into this state after a full-chassis failover (you prepare an ARX for failovers with the redundancy command).
- Failed
- FW Upgrade - indicates that the module is upgrading its firmware. You can start a firmware upgrade after installing a software release with new firmware; see the documentation for the firmware upgrade command.
- FW Upgrade Failed - means that the module failed an attempted firmware upgrade. Call F5 Support if you see this state.

Up Time is time since the last reboot. Use the reload command to reboot.

Memory (MB) is a heading for the following processor-memory measures, in MegaBytes:

- Total is the processor's total memory.
- Free is the processor's available memory.

CPU1M is the current CPU usage in a 1-minute period.

CPU5M is the current CPU usage in a 5-minute period.

Guidelines: Related Commands

Use the show system tasks command to view the tasks that are currently running on one or more processors. For statistics on CPU and memory usage over time, use the show processors usage command.

Samples

`stoweA> show processors`
shows the processors on an ARX-2500. See Figure 5.11 on page 5-45 for sample output.

`stkbrgA# show processors`
shows the processors on an ARX-VE. See Figure 5.12 on page 5-45 for sample output.

`bstnA# show processors`
shows the processors on an ARX-4000. See Figure 5.13 on page 5-46 for sample output.

`prtlndA# show processors`
shows the processors on an ARX-2000. See Figure 5.14 on page 5-46.

Related Commands

show chassis
show system tasks
show processors usage
nsm recovery
resource-profile

Figure 5.11 *Sample Output: show processors (ARX-2500)*

```
stoweA> show processors

Proc  Role                  CPU 1m  CPU 5m
----  --------------------  ------  ------
1.1   System                     3       1
1.2   System                     3       1
1.3   FastPath                   2       1
1.4   FastPath                   2       1
1.5   FastPath                   2       2
1.6   Volume-Group               2       3
1.7   Volume-Group               2       2
1.8   Volume-Group               3       1

Summary:

                   Memory (MB)      Swap (MB)
Up Time            Total    Free    Total    Free
-----------------  -------  ------  -------  ------
0 days, 00:20:44     16041   13540    32773   32773
```

Figure 5.12 *Sample Output: show processors (ARX-VE)*

```
stkbrgA# show processors

Proc  Role                  CPU 1m  CPU 5m
----  --------------------  ------  ------
1.1   Volume-Group              25      19
```

Figure 5.13 *Sample Output: show processors (ARX-4000)*

```
bstnA# show processors

                                          Memory (MB)
Proc  Module  State       Up Time             Total    Free    CPU 1m  CPU 5m
----  ------  ----------  ------------------  -------  ------  ------  ------
1.1   ACM     Up          0 days, 01:21:34      16030   15121       4       3
2.1   NSM     Up          0 days, 01:17:35       2650    2093       0       0
2.2   NSM     Up          0 days, 01:17:35       2650    2093       0       0
2.3   NSM     Up          0 days, 01:17:35       2650    2093       0       0
2.4   NSM     Up          0 days, 01:17:35       2650    2093       1       1
2.5   NSM     Up          0 days, 01:17:30       2650    2093       0       0
2.6   NSM     Up          0 days, 01:17:30       2650    2093       0       0
2.7   NSM     Up          0 days, 01:17:30       2650    2093       0       0
2.8   NSM     Up          0 days, 01:17:30       2650    2093       1       1
2.9   NSM     Up          0 days, 01:17:30       2650    2091       0       0
2.10  NSM     Up          0 days, 01:17:30       2650    2091       0       0
2.11  NSM     Up          0 days, 01:17:30       2650    2091       0       0
2.12  NSM     Up          0 days, 01:17:30       2650    2091       1       1
```

Figure 5.14 *Sample Output: show processors (ARX-2000)*

```
prtlndA# show processors

                                          Memory (MB)
Proc  Module  State       Up Time             Total    Free    CPU 1m  CPU 5m
----  ------  ----------  ------------------  -------  ------  ------  ------
1.1   ACM     Up          0 days, 02:29:46      11940   11211       2       3
1.2   ACM     Up          0 days, 02:20:40       2650    2093       0       0
1.3   ACM     Up          0 days, 02:20:40       2650    2093       1       3
1.4   ACM     Up          0 days, 02:20:40       2650    2093       0       0
1.5   ACM     Up          0 days, 02:20:40       2650    2093       0       0
```

show processors usage

Purpose The show processors usage command shows usage statistics for the processors (CPUs) installed on the ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show processors usage
show processors usage {slot.processor | all} [report]
```

slot.processor (optional) specifies one processor.

slot is the slot number.

processor is the processor number. Use the show processors command to show all processors and their associated *slot.processor* IDs.

all (optional) specifies all processors.

report (optional) is only an option if you select one processor or use the **all** keyword. This causes the CLI to create a report instead of showing output on the command line. The CLI shows the report name after you type the command. The report is named "processor_usage_*yyyymmddHHMMSS*.rpt," where *yyyymmddHHMMSS* is the date and time when the report was created. Use show reports to list all reports and view the report's contents.

Guidelines The output contains a table for each chosen processor. Each table shows 3-hour samples of CPU usage, with the most-recent sample at the top. The tables are labeled as follows:

Processor Usage for *slot.proc* appears for each processor on most platforms,

and

Processor Usage for *slot.proc role* appears instead on the ARX-1500 and ARX-2500. The *role* describes the types of processes that used the CPU cycles in the table:

- FastPath is for network-related processes, similar to those that run on the NSM in other platforms. These are also called *data plane* processes.
- Volume-Group is for storage-related processes, such as managed-volume processes and the policy engine.
- System is for the CLI and GUI (or manager) processes, and other processes related to system administration.

Guidelines (Cont.) Every row of the table shows the CPU usage for one 3-hour sample:

Date is the month and day of the 3-hour sample.

Time Interval is the start time and end time for the 3-hour sample.

%CPU shows the low, average (Avg), and high percentages of CPU cycles used during the sample period.

%Memory shows the same measures for memory. This does not appear in the output for the ARX-1500 or ARX-2500; you can use show memory usage to get these measures on either of those platforms.

%Swap shows the percentages for swap-space usage. An ACM only uses swap space when a large portion of standard memory is used up. On the ARX-1500 and ARX-2500, this command omits the %Swap table.

Use the show processors command to show the current state of all processors on the switch. The show system tasks command shows the currently-running tasks on one or more processors.

Samples

```
bstnA# show processors usage
```

shows the processor-usage statistics for all processors on an ARX-4000. See Figure 5.15 on page 5-48 for sample output.

```
canbyA# show processors usage
```

shows the processor-usage statistics for all processors on an ARX-1500. See Figure 5.16 on page 5-50 for sample output.

```
bstnA# show processors usage 1.1 report

Scheduling report: processor_usage_201104010144.rpt
```

sends the processor-usage statistics to a report. See Figure 5.17 on page 5-51 for a sample report.

Related Commands show processors
show system tasks
show memory usage

Figure 5.15 *Sample Output: show processors usage*

```
bstnA# show processors usage
Processor Usage for  1.1

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    1    2    3       2    4    5       0    0    0

Processor Usage for  2.1

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    0    0    0      21   21   22       0    0    0

Processor Usage for  2.2

Date                      %CPU             %Memory          %Swap
```

```
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    0    0    0      21   21   22       0    0    0

Processor Usage for  2.3

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    0    0    4      21   21   22       0    0    0

Processor Usage for  2.4

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    0    1    1      21   21   22       0    0    0

Processor Usage for  2.5

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    0    0    0      21   21   21       0    0    0

Processor Usage for  2.6

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    0    0    0      21   21   21       0    0    0

Processor Usage for  2.7

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    0    0    0      21   21   21       0    0    0

Processor Usage for  2.8

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    0    1    1      21   21   21       0    0    0

Processor Usage for  2.9

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    0    0    0      21   21   22       0    0    0

Processor Usage for  2.10

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  -------------    -------------    -------------
02/23  23:00 - 01:56    0    0    0      21   21   22       0    0    0

Processor Usage for  2.11
```

```
Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  --------------    --------------    --------------
02/23  23:00 - 01:56    0    0    0      21   21   22       0    0    0

Processor Usage for  2.12

Date                      %CPU             %Memory          %Swap
MM/DD  Time Interval  Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------  --------------    --------------    --------------
02/23  23:00 - 01:56    0    1    1      21   21   22       0    0    0
```

Figure 5.16 *Sample Output: show processors usage (ARX-1500)*

```
canbyA# show processors usage

Processor Usage for 1.1 System

Date                      %CPU
MM/DD  Time Interval   Low  Avg  High
-----  -------------   --------------
04/01  02:00 - 04:09   0    0    1
03/31  23:00 - 02:00   0    0    2


Processor Usage for 1.2 System

Date                      %CPU
MM/DD  Time Interval   Low  Avg  High
-----  -------------   --------------
04/01  02:00 - 04:09   0    2    3
03/31  23:00 - 02:00   0    4    3


Processor Usage for 1.3 FastPath

Date                      %CPU
MM/DD  Time Interval   Low  Avg  High
-----  -------------   --------------
04/01  02:00 - 04:09   2    2    2
03/31  23:00 - 02:00   1    2    2


Processor Usage for 1.4 FastPath

Date                      %CPU
MM/DD  Time Interval   Low  Avg  High
-----  -------------   --------------
04/01  02:00 - 04:09   2    2    2
03/31  23:00 - 02:00   1    2    2


Processor Usage for 1.5 Volume-Group

Date                      %CPU
MM/DD  Time Interval   Low  Avg  High
-----  -------------   --------------
04/01  02:00 - 04:09   0    0    2
03/31  23:00 - 02:00   0    0    2


Processor Usage for 1.6 Volume-Group
```

```
Date                      %CPU
MM/DD  Time Interval   Low  Avg  High
-----  -------------   --------------
04/01  02:00 - 04:09   0    0    1
03/31  23:00 - 02:00   0    1    1
```

Figure 5.17 *Sample Report: processor_usage_...*

```
bstnA# show reports processor_usage_201104010144.rpt
**** Processor Usage Report: Started at Fri Apr  1 01:44:24 2011 ****
**** Software Version: 6.00.000.13568 (Mar 30 2011 20:17:43) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

Processor Usage for 1.1

Date                      %CPU             %Memory           %Swap
MM/DD  Time Interval   Low  Avg  High    Low  Avg  High    Low  Avg  High
-----  -------------   --------------    --------------    --------------
03/31  23:00 - 01:39   1    2    4       2    4    6       0    0    0
**** Total processed:              1
**** Elapsed time:           00:00:00
**** Processor Usage Report: DONE at Fri Apr  1 01:44:24 2011 ****
```

show uptime

Purpose Use the show uptime command to see how long the system has been "up" since the last reboot.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show uptime`**

Guidelines This command displays the length of time since the last reboot in weeks, days, hours, and minutes, as shown in the sample below.

Sample

```
bstnA> show uptime
 bstnA uptime is 0 weeks, 0 days, 0 hours, 27 minutes.
```

Related Commands show processors
show version

show version

Purpose Use the show version command to display the installed-software versions and a summary of the chassis' configuration and state.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show version`

Guidelines: Release Table Running Release is the software version that is currently running. This appears in the following format:

release-file : Version *version-num* (*build-date*) *username*

where *release-file* is in the local "releases" directory, which you can see with the show releases command.

Armed Release is the software to be loaded on the next reboot, if any. This appears in the same format as the Running Release. Use boot system to arm the switch with a new release file. Use reload to reboot and put the new release file into service.

Backup Release is the software release that was running prior to the current release, if any. F5 personnel can roll the switch back to this release if needed.

Guidelines: System Configuration Table System Configuration is a table of module information for the hardware platforms, one row per module.

Slot is only 1 for all chassis except the ARX-4000, which has a second slot devoted to data-plane processes.

Admin is "enabled."

ModuleType is ACM (which runs management processes, control-plane processes, and possibly data-plane processes) or NSM (which exclusively runs data-plane processes).

ModuleState values appear in *Guidelines: Module States*, below.

FW Upgrade is "enabled" or "disabled." If enabled, a chassis-software upgrade automatically upgrades the module's firmware. To enable this feature, contact F5 personnel. To install firmware manually, after a software upgrade, use the firmware upgrade command.

For the ARX-VE, ARX-1500, and ARX-2500, the System Configuration table contains a single row showing how long it has been since the last reboot.

Guidelines: Resource Table The Resource table only appears for the ARX-500, ARX-2000, or ARX-4000 platforms. This table shows the high-level state of the switch, and whether it supports packet forwarding (like a MAC bridge) or not (like an end station).

State is "Up" or an error message. If there is an error, use show chassis for details.

Forwarding is "Enabled" (meaning the ARX can forward packets from one client/server port to another, performing the functions of a MAC bridge) or "Disabled." You can set this with the switch-forwarding enable command.

Guidelines: Module States

The ModuleState field can contain any of the following values:

- Online
- Empty
- Removed - the slot previously had a module, but it has been removed.
- Discovery - the chassis is finding all module states, very early in the boot process.
- Boot - the module is running its diagnostic tests, invoked during boot.
- Init - the module passed its diagnostic tests and is provisioning its internal processors.
- Downloading - if (during "Init") the module discovers that it needs new software, it enters this state to fetch and install the software.
- Online Partial - at least one processor is up, but not all of them are online yet. If a processor does not come up after 5 minutes, this changes to "Failed Partial" state.
- Failed Partial - at least one processor is up, but at least one processor failed.
- Failed
- Offline - applies to the network processors. This indicates that the ACM has not provided the network processors with important configuration parameters, so the processors cannot come online. Any of the following issues can result in this state on an ARX-500, ARX-2000, or ARX-4000:
  - The NVRAM battery failed on the ACM. Use show chassis nvr to confirm this; if the issue persists, contact F5 Support.
  - The NVRAM is not up yet. As above, you should contact F5 Support if the NVRAM issue persists.
- Poweroff - the slot has no power.
- FW Upgrade - the module is installing new firmware from the installed software release. The firmware upgrade command starts a firmware installation.
- FW Upgrade Failed - a firmware upgrade failed on this module.

Related Commands

boot system
show releases
reload

Figure 5.18 *Sample Output: show version (ARX-4000)*

```
bstnA> show version
  Copyright (c) 2002-2011 by F5 Networks, Inc. All rights reserved.
Running Release
test1.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Armed Release
test1.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Backup Release
test3.rel :  Version 6.00.000.13541 (Feb 8 2011 14:29:34) [nbuilds]

System Configuration:  Version 600000.29

  bstnA uptime is 0 weeks, 0 days, 1 hours, 3 minutes.
```

```
  Slot  Admin     ModuleType     ModuleState        FW Upgrade
  ----  -------   ------------   ----------------   ----------
    1   Enabled   ACM            Online             Disabled
    2   Enabled   NSM            Online             Disabled


 Resource         State                             Forwarding
 --------         -------------------------------   ----------
  Switch           Up                               Disabled
```

Figure 5.19 *Sample Output: show version (ARX-2500)*

```
stoweA> show version
  Copyright (c) 2002-2011 by F5 Networks, Inc. All rights reserved.
Running Release
test3.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Armed Release
test3.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Backup Release
test2.rel :  Version 6.00.000.13542 (Feb 10 2011 20:57:05) [nbuilds]

System Configuration:  Version 600000.29

  stoweA uptime is 0 weeks, 0 days, 4 hours, 39 minutes.
```

Figure 5.20 *Sample Output: show version (ARX-1500)*

```
canbyA> show version
  Copyright (c) 2002-2011 by F5 Networks, Inc. All rights reserved.
Running Release
test1.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Armed Release
test1.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Backup Release
test3.rel :  Version 6.00.000.13542 (Feb 10 2011 20:09:40) [nbuilds]

System Configuration:  Version 600000.29

  canbyA uptime is 0 weeks, 0 days, 4 hours, 53 minutes.
```

Figure 5.21 *Sample Output: show version (ARX-2000)*

```
prtlndA> show version
  Copyright (c) 2002-2011 by F5 Networks, Inc. All rights reserved.
Running Release
test3.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Armed Release
test3.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Backup Release
test1.rel :  Version 5.02.000.12642 (Feb 11 2011 16:00:53) [nbuilds]

System Configuration:  Version 600000.29

  prtlndA uptime is 0 weeks, 0 days, 0 hours, 19 minutes.

  Slot  Admin     ModuleType     ModuleState        FW Upgrade
  ----  -------   ------------   ----------------   ----------
```

```
    1   Enabled   ACM            Online             Disabled


 Resource         State                            Forwarding
 --------         -------------------------------  ----------
  Switch           Up                              Disabled
```

Figure 5.22 *Sample Output: show version (ARX-500)*

```
provA> show version
  Copyright (c) 2002-2011 by F5 Networks, Inc. All rights reserved.
Running Release
test1.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Armed Release
test1.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Backup Release
test2.rel :  Version 6.00.000.13541 (Feb 8 2011 09:00:16) [nbuilds]

System Configuration:  Version 600000.29

  provA uptime is 0 weeks, 0 days, 3 hours, 16 minutes.

  Slot  Admin     ModuleType     ModuleState        FW Upgrade
  ----  -------   ------------   ----------------   ----------
    1   Enabled   ACM            Online             Disabled


 Resource         State                            Forwarding
 --------         -------------------------------  ----------
  Switch           Up                              Disabled
```

Figure 5.23 *Sample Output: show version (ARX-VE)*

```
stkbrgA# show version
  Copyright (c) 2002-2011 by F5 Networks, Inc. All rights reserved.
Running Release
test3.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Armed Release
test3.rel :  Version 6.00.000.13543 (Feb 11 2011 20:13:05) [nbuilds]

Backup Release
test1.rel :  Version 6.00.000.13541 (Feb 8 2011 11:11:19) [nbuilds]

System Configuration:  Version 600000.29

  stkbrgA uptime is 0 weeks, 0 days, 0 hours, 6 minutes.
```

shutdown

Purpose Use the shutdown command to halt the ARX, perhaps from a remote location. This gracefully stops all power on the system.

Mode priv-exec

Security Role(s) crypto-officer, storage-engineer, network-engineer, or network-technician

Syntax **`shutdown`**

Default(s) None.

Platforms ARX-500, ARX-2000, ARX-4000, or ARX-VE

Guidelines This command prepares the ARX for a planned power outage. After you run this command on one of the hardware platforms, you have up to 72 hours before the battery drains for the Non-Volatile RAM (NVRAM). The NVRAM is memory on the ACM that stores all database transactions for the namespace processes; if the processes fail, they can recover by replaying these transactions. This is called *metalog* data. A managed volume requires its metalog data when it starts up.

You must manually power on the chassis to restore it to service. If you restore power after the NVRAM battery drains, all of the managed volumes on the ARX must re-import.

The CLI warns you that you must manually restore power to the chassis; enter **yes** to proceed with the power-off sequence. You can restore power on site by flipping the power switch or pushing the power button. Refer to the Hardware Installation Guide for your chassis to find the location of the power button or switch. The ARX-VE is a software-only platform that runs as a Virtual Appliance (VA, similar to a VM); use your VM client to restart the ARX-VE.

The ARX-VE has no NVRAM hardware, and stores its metalog data in one of the hypervisor's disk partitions. A re-import of managed volumes is therefore unnecessary for the ARX-VE, no matter how long it is shut down. (The same is true of the ARX-1500 and ARX-2500, which store their metalog data on their internal disk drives.)

You can use the reload command to reboot the ARX. To clear the NVRAM (in case of an unrecoverable corruption), use clear nvr.

Sample

```
provB# shutdown

This command turns off the chassis and powers it down.
You will need to manually restore power to return the chassis to
service.
Are you sure? [yes/no] yes
```

shuts down an ARX named "provB."

Related Commands reload
clear nvr

6

RAID Management

This chapter contains an alphabetical list of commands for managing the redundant array of independent disks (RAID) in the ARX chassis.

raid offline

Purpose Use the raid offline command to set a disk drive offline in the switch chassis.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`raid offline { disk1 | disk2 }`**

Valid Platforms ARX-2000 and ARX-4000

Guidelines After issuing the command, enter the show chassis diskuse command to verify changes and to view disk status.

Note

The following sample output shows the disk and raid details from the show chassis diskuse *command. See* Figure 5.1 on page 5-26 *and* Figure 5.6 on page 5-33 *for complete* show chassis *output samples.*

Sample

```
bstnA# raid offline disk1
Set disk drive in Bay 1 offline ? [yes/no] yes
bstnA# show chassis diskuse

 Logical Disk Details:
 Disk    Status                Verification Mode  Verification Rate
 ------  --------------------  -----------------  -----------------
 1       Optimal               Manual             50 %

 Disk Details:
 Disk    Size          State               Transfer Rate   Model
 --------   -----------   -----------------   -------------   -----------------
  Bay 1     136.73G       Online              3.0Gb/sec       ST3146855SS
  Bay 2     136.73G       Offline             3.0Gb/sec       ST3146855SS

 RAID Controller Details:
 Rebuild Rate  Max Transfer Rate  Firmware       RAID Alarm
 ------------  -----------------  -------------  ------------
 85 %          3.0Gb/sec          7.0.1-0061     Enabled

 Disk Usage:
 Name                             Total MB    Used MB     Free MB     Used%
 -------------------------------  ----------  ----------  ----------  -----
 System                           2331        1408        804         64%
 Releases                         5285        2687        2328        54%
 Logs                             54951       87          52071       1%
 Cores; DiagInfo; Lists           21133       219         19840       2%
 Scripts                          3172        59          2951        2%
 Reports                          8458        35          7992        1%
```

Related Commands raid rebuild
show chassis

raid rebuild

Purpose Use this command to rebuild a RAID configuration on a disk in the switch chassis.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`raid rebuild { disk1 | disk2 }`**

disk1 is the drive in Bay 1

disk2 is the drive in Bay 2

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines This command is required after installing a disk, so that it can rejoin the RAID.

This process slows down various internal processes, so we recommend running it on the backup peer in a redundant pair (see redundancy for information on redundant pairs).

The ARX-1500 and ARX-2500 write their metalog data to the RAID, and managed-volume performance is affected if these writes are slowed. Those platforms write all of their metalog data both to the active peer and the backup peer at the same time, so a raid rebuild on the backup peer still affects managed-volume processing on the active peer. For the ARX-1500 and ARX-2500, we recommend performing the rebuild during off hours only. You can use the show metalog usage command to view usage statistics for the metalog driver.

After issuing the command, enter the show chassis diskuse command to verify the change and to view the rebuild progress.

Sample

```
bstnA# raid rebuild disk1
bstnA# show chassis diskuse

 Logical Disk Details:
 Disk    Status                Verification Mode  Verification Rate
 ------  --------------------  -----------------  -----------------
 1       Optimal               Manual             50 %

 Disk Details:
 Disk    Size          State               Transfer Rate   Model
 --------   -----------   -----------------   -------------   -----------------
  Bay 1     136.73G       Rebuild 22%         3.0Gb/sec       ST3146855SS
  Bay 2     136.73G       Online              3.0Gb/sec       ST3146855SS

 RAID Controller Details:
 Rebuild Rate  Max Transfer Rate  Firmware       RAID Alarm
 ------------  -----------------  -------------  ------------
 85 %          3.0Gb/sec          7.0.1-0061     Enabled

 Disk Usage:
 Name                             Total MB    Used MB     Free MB     Used%
 -------------------------------  ----------  ----------  ----------  -----
 System                           2331        1408        804         64%
...
```

Related Commands

raid offline
raid rebuild-rate
show chassis

raid rebuild-rate

Purpose This command specifies how much of the system's RAID-controller resources to use in rebuilding the RAID.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax **`raid rebuild-rate`** ***`rate`***

rate (1-99) is the percentage of the system's RAID controller, an internal CPU used for managing access to the internal disks. A higher number ensures a faster rebuild, but slower disk access during the rebuild.

Default(s) 90

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines The CLI and the GUI access the internal disks for many operations, as do several processes that create reports. Whenever a raid rebuild occurs, this command determines the division of RAID-controller cycles between these disk-access operations and the RAID-rebuild operation. After issuing the command, enter the show chassis diskuse command to verify changes and to view the rebuild rate.

Sample `bstnA(cfg)#` **`raid rebuild-rate 85`**

sets the RAID-rebuild rate to 85% of the RAID-controller's time. This leaves 15% of the controller's cycles to manage disk I/O operations, such as recording the running configuration or writing reports.

Related Commands raid rebuild
show chassis **`diskuse`**

raid silence

Purpose Use this command to silence the audible RAID alarm. Use the no form to re-activate the alarm.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
raid silence
no raid silence
```

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines The RAID subsystem sounds an alarm on disk failure. Use show chassis to find the disk drive that failed. Replace the failed disk as instructed in the ARX®–4000 Hardware Installation Guide, the ARX®–2000 Hardware Installation Guide, or the instructions that come with the replacement disk.

Use this command to silence the alarm while you replace the drive. Once done, re-enable the alarm with no raid silence.

Samples

```
bstnA# raid silence
```

silences the alarm.

```
bstnA# no raid silence
```

reactivates a silenced alarm.

Related Commands show chassis

raid verification-mode

Purpose Use the raid verification-mode command to determine whether or not the RAID verification test runs automatically.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
raid verification-mode manual
raid verification-mode automatic [at start every
  {num days | day-of-week}]
```

at *start* every {*num* days | *day-of-week*}] (optional) is only available on the ARX-1500 and ARX-2500.

start is the scheduled start time for verifications to run each day. The time is in the following format:

***HH*:*MM*[:00]**

where the only possible value for seconds is "00."

every *num* days (optional) chooses some number of days between automatic verifications. We recommend a ***num*** of 1, so that the verifications run daily.

every *day-of-week* (optional) chooses a single day of the week to run automatic verifications. The options are "sunday," "monday," "tuesday," and so on. This makes the automatic verifications run weekly instead of daily; it is not generally recommended.

Default(s)

```
automatic
at 23:00 every 1 days
```

on ARX-1500 and ARX-2500 only. On all other platforms, automatic verification is continuous.

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines The verification test checks the integrity of every block on each disk. Whenever a bad block is found, the test attempts to replace it with its counterpart on the other disk. The verification fails if any block is corrupted on both disks. It also checks the disk controller for errors.

If the verification mode is manual, you can use the raid verification-rate command to determine the CPU cycles used by the verification process. If the verification mode is automatic, the verification rate is set by a system default: on the ARX-1500 and ARX-2500, it runs once at 11PM (see below); on all other chassis types, it runs continuously, every 5 minutes.

Whether or not the verification mode is manual, you can use the raid verify command to run the verification manually.

After issuing the command, you can use the show chassis diskuse command to verify the change.

Guidelines: Use Manual Verification for the ARX-1500 and ARX-2500

The ARX-1500 and ARX-2500 write their metalog data to the RAID, and managed-volume performance is affected if these writes are slowed. Constant, automatic RAID verification may have an adverse affect on storage processing for those platforms. For the ARX-1500 and ARX-2500, we recommend using the at option with raid verification-mode automatic, and choose a schedule that runs verifications during slow business hours. If you choose automatic on these platforms, the verification runs daily at 11PM by default.

You can use the show metalog usage command to monitor the internal metalog driver's usage of the RAID.

Samples

```
bstnA# raid verification-mode manual
bstnA# show chassis diskuse

 Logical Disk Details:
 Disk    Status                Verification Mode  Verification Rate
 ------  --------------------  -----------------  -----------------
 1       Optimal               Manual             50 %

 Disk Details:
 Disk    Size          State               Transfer Rate   Model
 --------   -----------   -----------------   -------------   -----------------
  Bay 1     136.73G       Online              3.0Gb/sec       ST3146855SS
  Bay 2     136.73G       Online              3.0Gb/sec       ST3146855SS

 RAID Controller Details:
 Rebuild Rate  Max Transfer Rate  Firmware       RAID Alarm
 ------------  -----------------  -------------  ------------
 85 %          3.0Gb/sec          7.0.1-0061     Enabled

 Disk Usage:
 Name                             Total MB    Used MB     Free MB     Used%
 -------------------------------  ----------  ----------  ----------  -----
 System                           2331        1408        804         64%
 Releases                         5285        2687        2328        54%
 Logs                             54951       87          52071       1%
 Cores; DiagInfo; Lists           21133       219         19840       2%
 Scripts                          3172        59          2951        2%
 Reports                          8458        35          7992        1%
```

sets manual verification in the "bstnA" ARX.

```
stoweA# raid verification-mode automatic at 02:00:00 every 1 days
```

sets automatic verification in the "stoweA" ARX, an ARX-2500, and schedules the verification for 2 AM every morning. This is a best practice for the ARX-1500 and ARX-2500, where the RAID verifications may slow the performance of managed volumes.

Related Commands

raid verify
raid verification-rate
show chassis

raid verification-rate

Purpose This command specifies how much of the system's RAID-controller resources to use in verifying the integrity of the RAID.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax `raid verification-rate` ***rate***

rate (1-99) is the percentage of the system's RAID-controller resources; the higher the number, the heavier the burden on the RAID controller.

Default(s) 10

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines The RAID controller is a separate CPU from the CPUs used for data-plane (network) processing or control-plane processing. This controller is used for the internal disks, and runs when the internal disks are used. The system uses the internal disks for the various copy commands (such as copy ftp) and move commands (such as move ... ftp). This CPU is also required when internal processes write to their log files (see show logs), create reports (see show reports), write packet-capture files (see capture session), collect diagnostic information and put it into a local file, or otherwise write to the internal disks.

You cannot change the default for automatic verification tests. These run continuously on the system; a higher rate would degrade system performance. This command changes the RAID-controller percentage used for *manual* (user-invoked) verification tests. You can stop the test from running automatically with the raid verification-mode manual command, and then use the raid verify command to manually run the test.

The verification test checks the integrity of every block on each disk. Whenever it finds a bad block, it attempts to replace the block with its twin block on the other disk. The verification fails if any block is corrupted on both disks. The test also checks the RAID controller for errors.

After issuing this command, enter the show chassis diskuse command to verify the change.

Sample

```
bstnA# raid verification-rate 50
bstnA# show chassis diskuse

 Logical Disk Details:
 Disk    Status                Verification Mode  Verification Rate
 ------  --------------------  -----------------  -----------------
 1       Optimal               Manual             50 %

 Disk Details:
 Disk    Size          State               Transfer Rate   Model
 --------   -----------   -----------------   -------------   -----------------
  Bay 1     136.73G       Online              3.0Gb/sec       ST3146855SS
  Bay 2     136.73G       Online              3.0Gb/sec       ST3146855SS

 RAID Controller Details:
 Rebuild Rate  Max Transfer Rate  Firmware       RAID Alarm
 ------------  -----------------  -------------  ------------
 85 %          3.0Gb/sec          7.0.1-0061     Enabled

 Disk Usage:
 Name                              Total MB    Used MB     Free MB     Used%
 --------------------------------  ----------  ----------  ----------  -----
 System                            2331        1408        804         64%
 Releases                          5285        2687        2328        54%
 Logs                              54951       87          52071       1%
 Cores; DiagInfo; Lists            21133       219         19840       2%
 Scripts                           3172        59          2951        2%
 Reports                           8458        35          7992        1%
```

Related Commands raid verify
show chassis

raid verify

Purpose This command invokes a verification test for both disks in the RAID. Use the no form of the command, no raid verify, to stop a RAID verification that is currently in progress.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
raid verify
no raid verify
```

Default(s) None.

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines You cannot use this command while another RAID-verify operation is running, and another RAID-verify operation is frequently running when the raid verification-mode command is at its default (automatic). The show chassis diskuse command shows the current mode, manual or automatic. Use raid verification-mode manual to allow manual verification at any time, and guarantee that no RAID-verify process is running in the background. You can also use no raid verify to cancel the current RAID-verify operation, whether it was invoked automatically or manually.

The verification test checks the integrity of every block on each disk. Whenever it finds a bad block, it attempts to replace the block with its twin block on the other disk. The verification fails if any block is corrupted on both disks. The test also checks the disk controller for errors.

After issuing the command, enter the show chassis diskuse command to view the progress of the test, and its final results. Refer to the Status of the Logical Disk (that is, the full RAID), in the first table of the output. The test takes more than one hour. If the verification fails, you can replace one or both disks as instructed in the ARX®–4000 Hardware Installation Guide, the ARX®–2000 Hardware Installation Guide, or in the instructions that come with the replacement disk.

Sample

```
bstnA# raid verify

% INFO: RAID verification has started.

bstnA# show chassis diskuse

 Logical Disk Details:
 Disk    Status                Verification Mode  Verification Rate
 ------  --------------------  -----------------  -----------------
 1       Verifying 17%         Manual             50 %

 Disk Details:
 Disk    Size          State               Transfer Rate   Model
 --------   -----------   -----------------   -------------   -----------------
  Bay 1     136.73G       Online              3.0Gb/sec       ST3146855SS
  Bay 2     136.73G       Online              3.0Gb/sec       ST3146855SS

 RAID Controller Details:
 Rebuild Rate  Max Transfer Rate  Firmware       RAID Alarm
 ------------  -----------------  -------------  ------------
 85 %          3.0Gb/sec          7.0.1-0061     Enabled

 Disk Usage:
 Name                              Total MB    Used MB     Free MB     Used%
 --------------------------------  ----------  ----------  ----------  -----
 System                            2331        1408        804         64%
 Releases                          5285        2687        2328        54%
 Logs                              54951       87          52071       1%
 Cores; DiagInfo; Lists            21133       219         19840       2%
 Scripts                           3172        59          2951        2%
 Reports                           8458        35          7992        1%
```

Related Commands raid rebuild
show chassis **chassinfo**

7

File Management

at

Purpose Use the at command to schedule a CLI command or script to run at a later time.

Mode cfg

Security Role(s) crypto-officer

Syntax `at start [every interval] do action [report report-prefix]`

is the high-level syntax for this command.

start is the scheduled start time. This cannot be in the past unless you specify the optional **date** (below) and use **every *interval*** to run the job regularly. It breaks down into the following format:

[date *mm/dd/yyyy*] *HH*:*MM*[:00]

- ***mm/dd/yyyy*** (optional) specifies a date (for example, 01/07/2005 for January 7, 2005).
- ***HH*:*MM*[:00]** is the hours and minutes on a 24-hour clock (for example, 04:00 means 4AM). The only possible value for seconds is "00."

every *interval* (optional) creates a regular interval at which the CLI jobs run. This breaks down into one of two formats:

every *count* {minutes | hours | days | weeks | months}

- ***count*** (1-4,294,967,295) is the number of minutes, hours, days, etc. between CLI runs. For example, **every 10 minutes** or **every 6 months**.
- **minutes | ... months** is a required choice.

every {sunday | monday | tuesday | ... | saturday}

creates a weekly interval that runs on the given day of the week.

do *action* is required. The ***action*** is either a CLI command or a directive to run a script:

do *cli-command* | do run *cli-script*

- ***cli-command*** (1-255 characters) is a valid CLI command. Surround this with quotation marks if it contains any spaces. For example, **do "show cifs-service user-sessions all"**.
- ***cli-script*** (1-255 characters) identifies a script of CLI commands in the scripts directory. Use the show scripts command for a full listing of available scripts.

report *report-prefix* (optional, 1-255 characters) creates a report for the scheduled job. The report is named as follows:

report-prefix_yyyymmddHHMM.rpt, where *report-prefix* is chosen here, *yyyy* is the year of the scheduled run, *mm* is the month, *dd* is the day, *HH* is the hour, and *MM* is the minute. Each run of the CLI job generates a report (if you specify an interval with the **every** option), so the date makes it possible to differentiate multiple reports for the same job.

Default(s) **every *interval*** - none; the CLI command or script runs only once.

report *report-prefix* - none; the at command does not generate a report (though the command itself may generate one).

Guidelines The CLI displays the scheduled execution time after you enter this command. Use show at to view all pending CLI jobs. You can remove a job from the schedule with the clear at command.

To run a script immediately, you can use the run command. To download a CLI-script file from an FTP, SCP, or TFTP site, use the copy command. You can use expect monitor to repeat any show command until a certain string appears in its output.

Samples

```
bstnA(cfg)# at 01:06 every 5 minutes do "show sessions" report
adminSessions

  The scheduled execution time for AT job ID '1' is: 8/26/10 1:06 AM
```

runs the show sessions command every 5 minutes. Each run generates a report with the prefix, "adminSessions." See Figure 7.1 for a sample report.

```
bstnA(cfg)# at 10:30:00 do "expect show firewall timeout 30"

  The scheduled execution time for AT job ID '2' is: 12/6/05 10:30 AM.
```

runs the expect show firewall command at 10:30 in the morning, once only.

```
bstnA(cfg)# at date 05/05/2011 03:45:00 every 1 days do
"active-directory update forest medarch.org proxy-user acoProxy2"

  The scheduled execution time for AT job ID '3' is: 5/5/11 3:45 AM.
```

runs the active-directory update forest command at 3:45 AM, daily.

Related Commands show at
clear at
copy ftp, copy scp, copy {nfs|cifs}, and copy tftp
show scripts

Figure 7.1 *Sample Report: adminSessions_....rpt*

```
bstnA# show reports adminSessions_201008260431.rpt
bstnA- terminal character-set unicode-utf-8
bstnA- remark notice "Cli commands from AT command scheduler"
bstnA- no terminal confirmation
bstnA- enable
bstnA- show sessions
Connected Sessions
------------------

Session ID:                  14465
Username:                    admin
Access:                      ssh
Connect Time:                0 days, 00:01:32
Source IP:                   172.16.100.183

bstnA- remark notice "Cli commands complete AT command scheduler"
```

clear at

Purpose Every time someone uses the at command, one CLI job is added to the command schedule. Use the clear at command to remove one CLI job from this schedule, or to remove all of them.

Mode priv-exec

Security Role(s) crypto-officer

Syntax `clear at [job-id]`

job-id (optional, 1-2,147,483,647) identifies a specific CLI job to clear. If you omit this, the command clears all CLI jobs. Use the show at command for a list of all scheduled CLI jobs, with their job IDs.

Default(s) clear all scheduled CLI jobs if no ***job-id*** is used.

Guidelines If you clear all scheduled jobs, the CLI prompts for confirmation. Enter **yes** to proceed.

Use the at command to schedule a CLI command (or script) to run in the future, once or periodically. Use show at to view all pending CLI jobs.

Samples

```
bstnA# clear at
Remove all 2 jobs from the command scheduler? [yes/no] yes
bstnA#
```

clears all scheduled CLI jobs.

```
bstnA# clear at 3
```

removes one scheduled CLI job, with job ID 3.

Related Commands at
show at

copy ftp

Purpose

Use the copy ftp command to transfer a file (such as a software-release file or a log file) to or from the ARX via FTP.

Mode

priv-exec

Security Role(s)

network-technician, network-engineer, storage-engineer, or crypto-officer

(A backup-operator can copy reports.)

Syntax: Download to the Switch

```
copy ftp://[user[:password]@]server/source-file dir file
```

ftp://[*user*[:***password*]@]*server/source-file*** (1-1024 characters) is the URL for the source file:

> ***user*[:*password*]@** (optional) are the credentials for FTP access. If you omit them, they default to the credentials set by the ip ftp-user command. If you omit the password, the CLI prompts for one.
>
> ***server*** is the IP address or hostname for the FTP server.
>
> ***source-file*** is the source-file path. Lead with an extra slash ("/") if the path is absolute (for example, "ftp://10.1.1.5//var/rels/aco4665.rel"). Use only one slash if the path is local to the home directory for ***user*** (for example, "ftp://10.1.1.5/basic_ns.scr"). This conforms with the specification for FTP URLs in RFC 1738. You cannot use wildcards (such as *).

dir is the destination directory. Choose one of the following: **configs**, **scripts**, **license**, or **releases**.

file (1-255 characters) is the name you choose for the copy.

Syntax: Upload

```
copy directory file ftp://[user:password@]server/dest-path
  [format {text|report|csv|xml}]
```

directory is one of the following: **logs**, **stats-logs**, **cores**, **configs**, **replicated-configs**, **reports**, **diag-info**, **software**, **scripts**, **capture**, or **license**.

file (1-255 characters) identifies the source file from the above directory. You can use wildcards (such as *) to select multiple files.

ftp://[*user*:*password*@]*server/dest-path* (1-255 characters) is the URL for the destination directory or file (see above).

format {text|report|csv|xml} (optional) applies to the "reports" directory only. This chooses the output format for the external copy of the report. The "xml" option converts the report to XML format, "csv" converts the report to a Comma-Separated-Value format (useful for spreadsheets), and the "text" and "report" options leave the reports in plain-text format.

Default(s)

user:*password* from the URL - the default set by ip ftp-user.

Guidelines

If you enter a user name without a password, the CLI prompts for the password before continuing. Enter the password for the *user* you included in the URL.

Use copy scp to upload or download over a secure connection, using SCP. The copy tftp command uses TFTP (Trivial FTP) as a transport. To send out a file as an E-mail attachment, use copy smtp. The copy {nfs|cifs} command transfers files between the maintenance directories on the ARX and its client-accessible volumes.

For a download to the local disks, you can use show directories *directory-name* to verify that the copy was successful.

To move a file from the local disk to an FTP server, thereby removing it from the local disk, you can use the move ... ftp command.

To manage local files, use delete and/or rename. To view ASCII files, use show directory file-name, tail, and/or grep.

Guidelines: Reformatting Reports to XML or CSV

For copies from the reports directory, you can reformat the copy to XML or CSV (comma-separated value) format. Use a .xml or .csv extension for the *destination-file-name*, and the copy assumes the chosen format. A .txt or .rpt extension leaves the destination file in plain-text format.

Guidelines: Wildcards

You can use several wildcard characters to select multiple source-file names for an upload:

- * matches any string, including an empty string.
- ? matches any single character. This is also a special meta character in the CLI, so you must quote the file-path string to use it; for example, copy logs traplog.? ftp ... would fail, but copy logs “traplog.?” ftp ... would succeed.
- [] surrounds a class of characters (for example, [abcz]), and matches any one of the characters inside the square brackets (a, b, c, or z).
- [a-z] matches any single character from the range a through z (all lower-case).

Samples

```
bstnA# copy ftp://jpublic@arxftp.f5.com/rel6.0.rel releases rel6.0.rel
Password: jpassword
```

downloads a release file, rel6.0.rel, from the FTP server at arxftp.f5.com.

```
bstnA# copy logs syslog ftp://noc.phredco.com/syslog-switch10
```

uploads the syslog to the FTP server at noc.phredco.com. This example uses the default username and password, set by ip ftp-user.

Related Commands

ip ftp-user
copy scp
copy tftp
copy smtp
copy {nfs|cifs}
move ... ftp
delete
rename
grep
tail
show directories

copy {nfs|cifs}

Purpose Use the copy {nfs|cifs} command to transfer a file (such as a software-release file or a log file) between the ARX and one of its managed volumes.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

(A backup-operator can copy reports.)

Syntax: From Volume to Switch

```
copy {nfs|cifs} namespace vol-path source-file directory file
```

nfs | cifs is a required choice. This chooses the protocol for the file transfer.

namespace (1-30 characters) identifies the namespace that holds the file.

vol (1-1024 characters) is the name of the volume that holds the source file.

source-file (1-255 characters) is path to the source file, starting at the root of the above volume. You can use wildcards (such as *, ?, or [a-z]) to select multiple files; see the *Guidelines: Wildcards* below.

directory is the destination directory. Choose one of the following: **configs**, **scripts**, **license**, or **releases**.

file (1-255 characters) is the name you choose for the copy. If you copy a file into the **releases** directory, its extension must be ".rel" (for example, "lastestRelease.rel").

Syntax: From Switch to Volume

```
copy directory file {nfs|cifs} namespace vol [dest-file-path]
  [format {text|report|csv|xml}]
```

directory is one of the following: **logs**, **stats-logs**, **cores**, **configs**, **replicated-configs**, **reports**, **diag-info**, **software**, **scripts**, **capture**, or **license**.

file (1-255 characters) identifies the source file from the above directory. As above, you can use wildcards (such as *) to specify multiple files.

nfs | cifs is a required choice. This chooses the protocol for the file transfer.

namespace (1-30 characters) identifies the namespace to hold the copy.

vol (1-1024 characters) is the volume to hold the copy.

dest-file-path (optional, 1-255 characters) is the path to the destination file, starting at the volume root. If you omit this, the destination file appears in the volume root, and has the same name as the source file.

format {text|report|csv|xml} (optional) applies to the "reports" directory only. This chooses the output format for the external copy of the report. The "xml" option converts the report to XML format, "csv" converts the report to a Comma-Separated-Value format (useful for spreadsheets), and the "text" and "report" options leave the reports in plain-text format.

Default(s) None

Guidelines

This command is useful on a site that prohibits FTP access to the datacenter, or otherwise restricts network connections between the ARX and the Internet. You can go to a host in the client network to access the Internet, download a desired file to an ARX volume, and then use this copy command to copy the file into an ARX directory. For example, you could use this method to download a new release file from https://downloads.f5.com/esd/productlines.jsp and then copy it to the "releases" directory.

The CLI chooses a user identity based on the transfer protocol. For NFS, the CLI uses *root* as its identity. For CIFS, the CLI uses the proxy-user (gbl-ns) for the chosen namespace. These identities typically ensure that there are no permissions problems during the copy operation.

There are several alternative options for the copy command that access the network directly. The copy ftp command copies files over an FTP connection to or from a remote server. Use copy scp to upload or download over a secure connection, using SCP. The copy tftp command uses TFTP (Trivial FTP) as a transport. To send out a file as an E-mail attachment, use copy smtp.

For a download to the local disks, you can use show directories *directory-name* to verify that the copy was successful.

To move a file from the local disks to a ARX volume, thereby deleting it from the local disk, you can use the move ... {nfs|cifs} command.

To manage local files, use delete and/or rename. To view ASCII files, use show directory *file-name*, tail, and/or grep.

Guidelines: Wildcards

You can use several wildcard characters to select multiple file names:

- * matches any string, including an empty string.
- ? matches any single character. This is also a special meta character in the CLI, so you must quote the file-path string to use it; for example, copy logs traplog.? cifs insur /claims would fail, but copy logs "traplog.?" cifs insur /claims would succeed.
- [] surrounds a class of characters (for example, [abcz]), and matches any one of the characters inside the square brackets (a, b, c, or z).
- [a-z] matches any single character from the range a through z (all lower-case).

Samples

```
bstnA# copy cifs medarcv /rcrds/maint/ 11658.rel releases latest.rel
```

copies a release file, 11658.rel, from a directory in the "medarcv~/rcrds" volume. After the copy is complete, this command sequence could go on to load this release (with boot system and reload).

```
bstnA# copy logs syslog* nfs wwmed /acct/.admin/
```

copies all syslog files from the local "logs" directory to a hidden directory in "wwmed~/acct," an NFS volume.

Related Commands

ip ftp-user
copy scp
copy tftp
copy smtp
move ... {nfs|cifs}
delete
rename
grep
tail
show directories

copy ron

Purpose

Use the copy ron command to securely copy a configuration or script file to another ARX on the current Resilient-Overlay Network (RON).

Mode

priv-exec

Security Role(s)

network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax: Download to the Current Switch

`copy ron source-arx source-directory source-file directory file`

source-arx (1-30 characters) is the hostname of the remote ARX. Use the show ron command for a list of all ARX systems on the current RON.

source-directory is the source directory on the remote ARX. Choose **configs** or **scripts**.

source-file (1-255 characters) is the name of the remote file.

directory is the destination directory on the local ARX. You can choose **configs** or **scripts**.

file (1-255 characters) is the name you choose for the local copy.

Syntax: Upload to a Remote Switch

`copy directory file ron dest-arx dest-directory [dest-file]`

directory is **configs** or **scripts**.

file (1-255 characters) identifies the source file from the above directory. You can use wildcards (such as *, ?, or [a-z]) to select multiple files; see the *Guidelines: Wildcards* below.

dest-arx (1-30 characters) is the hostname of the remote ARX. Use the show ron command for a list of all ARX systems on the current RON.

dest-directory is the destination directory on the remote ARX. Choose **configs** or **scripts**.

dest-file (optional, 1-255 characters) is the name of the remote file. The CLI ignores this if you used wildcards to select multiple files.

Default(s)

None

Guidelines

This copy operation uses the ip ron-user credentials at the remote switch, or the current user credentials if no such RON user is defined. If the administrative account is not defined at the remote switch, or if it is defined with a different password, the copy operation fails.

The file transfer is encrypted and sent with the Secure Copy (SCP) protocol.

There are several other copy commands for copying files to or from other servers, or for copying to or from a client-accessible volume on the current switch. Use copy ftp or copy tftp to upload or download using FTP or TFTP. To copy a file to or from an SCP server in your network, use copy scp. To send out a file as an E-mail attachment, use copy smtp. The copy {nfs|cifs} command transfers files between the maintenance directories on the ARX and its client-accessible volumes.

After a download to the local disks, you can use show directories *directory-name* to verify that the copy was successful.

To manage local files, use delete and/or rename. To view ASCII files, use show directory file-name, tail, and/or grep.

Guidelines: Wildcards

You can use several wildcard characters to select multiple source-file names for an upload:

- * matches any string, including an empty string.
- ? matches any single character. This is also a special meta character in the CLI, so you must quote the file-path string to use it; for example, copy scripts homeDir.? ron ... would fail, but copy scripts “homeDir.?” ron ... would succeed.
- [] surrounds a class of characters (for example, [abcz]), and matches any one of the characters inside the square brackets (a, b, c, or z).
- [a-z] matches any single character from the range a through z (all lower-case).

Samples

```
bstnA# copy ron prtlndA scripts setupUsrShr.scr scripts setup.scr
```

downloads a script, setupUsrShr.scrl, from a remote ARX named “prtlndA.”

```
provA# copy configs test.rcfg ron newptA configs test2.rcfg
```

sends a replicated-config file to a remote ARX named “newptA.”

Related Commands

ip ron-user
copy ftp
copy tftp
copy scp
copy smtp
copy {nfs|cifs}
delete
rename
grep
tail
show directories

copy scp

Purpose Use the copy scp command to securely copy a file (such as a software-release file or a log file) to or from the ARX over SCP.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

(A backup-operator can copy reports.)

Syntax: Download to the Switch

```
copy scp://[user@]server:source-file directory file
[accept-host-key]
```

scp://[*user*@]*server*:*source-file* (1-1024 characters) is the URL for the source file:

> ***user*@** (optional if someone created an ip scp-user) is the username to present to the other end of the SCP connection. This user must be valid at the remote host. If you omit this and an ip scp-user is defined, it defaults to the username set by that command.
>
> ***server*:** is the IP address or hostname for the SCP host. End with a colon (:).
>
> ***source-file*** is the source-file path. Lead with a slash ("/") if the path is absolute (for example, "scp://root@10.1.1.5:/var/rels/aco4665.rel"). Use no slash if the path is local to the home directory for ***user*** (for example, "scp://root@10.1.1.5:basic_ns.scr").

directory is the destination directory. Choose one of the following: **configs**, **scripts**, **license**, or **releases**.

file (1-255 characters) is the name you choose for the local copy.

accept-host-key (optional) indicates that if the other end of the connection has an unknown SSH host key (that is, if it is new, or if its key has changed since the last time the host was contacted), the ARX should accept the new host key and continue with the download. Otherwise, the ARX stops the download if the host presents an unknown key.

Syntax: Upload

```
copy directory file scp://[user@]server:dest-path
[accept-host-key] [format {text|report|csv|xml}]
```

directory is one of the following: **logs**, **stats-logs**, **cores**, **configs**, **replicated-configs**, **reports**, **diag-info**, **software**, **scripts**, **capture**, or **license**.

file (1-255 characters) identifies the source file from the above directory. You can use wildcards (such as *, ?, or [a-z]) to select multiple files; see the *Guidelines: Wildcards* below.

scp://[*user*@]*server*:*dest-path* (1-255 characters) is the URL for the destination.

accept-host-key (optional) means to accept an unknown SSH host key from the remote host, if it presents one.

format {text|report|csv|xml} (optional) applies to the "reports" directory only. This chooses the output format for the external copy of the report. The "xml" option converts the report to XML format, "csv" converts the report to a Comma-Separated-Value format (useful for spreadsheets), and the "text" and "report" options leave the reports in plain-text format.

Default(s) None

Guidelines If you enter a user name without a password, the CLI prompts you for a password after you issue this command. Enter the password for the *user* in the URL. The file is transferred over SCP, which uses a secure, encrypted connection to the remote host.

Every host identifies itself with an encrypted SSH key. Whenever the ARX transfers a file using SCP, it exchanges keys with the remote host. If the remote-host key is different from one received previously, the host may not be the one you intended; an attacker may be posing as the intended host. For this reason, the ARX does not download or upload to a host with an unknown key unless you raise the **accept-host-key** flag. Note that the first contact with any server will require this flag to succeed.

The ARX deletes all SSH host keys on reboot, so all remote hosts are again unknown when the switch comes back up.

Use copy ftp or copy tftp to upload or download using FTP or TFTP. To send out a file as an E-mail attachment, use copy smtp. The copy {nfs|cifs} command transfers files between the maintenance directories on the ARX and its client-accessible volumes.

After a download to the local disks, you can use show directories directory-name to verify that the copy was successful.

To move a file from the local disk to an SCP server, thereby removing it from the local disk, you can use the move ... scp command.

To manage local files, use delete and/or move. To view ASCII files, use show directory file-name, tail, and/or grep.

Guidelines: Wildcards You can use several wildcard characters to select multiple source-file names for an upload:

- * matches any string, including an empty string.
- ? matches any single character. This is also a special meta character in the CLI, so you must quote the file-path string to use it; for example, copy logs syslog.? scp ... would fail, but copy logs “syslog.?” scp ... would succeed.
- [] surrounds a class of characters (for example, [abcz]), and matches any one of the characters inside the square brackets (a, b, c, or z).
- [a-z] matches any single character from the range a through z (all lower-case).

Samples

```
bstnA# copy scp://jpublic@myserver:rel2.0.rel releases r2.0.rel
Password: password-for-jpublic
```

downloads a release file, rel2.0.rel, from jpublic’s home directory at the host named “myserver.”

```
bstnA# copy logs syslog scp://noc:/var/logs/acolog10 accept-host-key
```

uploads the syslog to a host named “noc,” using a default username and password. If the host key has changed for noc (or this is the first SCP to or from noc), the switch accepts the key and continues the upload.

Related Commands

ip scp-user
copy ftp
copy tftp
copy smtp
copy {nfs|cifs}
move ... scp
delete
move
grep
tail
show directories

copy smtp

Purpose Use the copy smtp command to send a maintenance file (such as a software-release file or a log file) as an E-mail attachment. This syntax only supports uploads from the switch.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

(A backup-operator can copy reports, but no other type of file.)

Syntax

```
copy directory file smtp://[e-mail-address/]destination-file
  [format {text|report|csv|xml}]
```

directory is one of the following: **logs**, **stats-logs**, **cores**, **configs**, **replicated-configs**, **reports**, **diag-info**, **software**, **scripts**, **capture**, or **license**.

file (1-255 characters) identifies the source file from the above directory, such as a log file or a report. You can use wildcards (such as *, ?, or [a-z]) to select multiple files; see the *Guidelines: Wildcards* below. Each matching file results in a separate E-mail message.

smtp://[*e-mail-address*/]*destination-file* (1-255 characters) is an E-mail destination for the file:

> **smtp://** is required. This declares that the destination is an E-mail address.
>
> ***e-mail-address*** (optional) is the recipient of the E-mail in *username@host* format (for example, "jsmith@myco.com"). If you omit this, the CLI uses the default address set by the cfg-smtp to command.
>
> ***destination-file*** is the name of the copy. If you used wildcards to select multiple files, this is the prefix for each copy. Each copy is sent as an attachment to the outbound E-mail message.

format {text|report|csv|xml} (optional) applies to the "reports" directory only. This chooses the output format for the external copy of the report. The "xml" option converts the report to XML format, "csv" converts the report to a Comma-Separated-Value format (useful for spreadsheets), and the "text" and "report" options leave the reports in plain-text format.

Default(s) the destination E-mail address defaults to the one set by the to command

Guidelines DNS lookups must be configured for SMTP to function (see the documentation for ip name-server).

Use copy ftp, copy tftp, or copy scp to upload or download using FTP, TFTP, or SCP. The copy {nfs|cifs} command transfers files between the maintenance directories on the ARX and its client-accessible volumes.

To manage local files, use delete and/or move. To view ASCII files, use show directory file-name, tail, and/or grep.

Guidelines: Wildcards

You can use several wildcard characters to select multiple source-file names for an upload:

- * matches any string, including an empty string.
- ? matches any single character. This is also a special meta character in the CLI, so you must quote the file-path string to use it; for example, copy logs syslog.? smtp ... would fail, but copy logs “syslog.?” smtp ... would succeed.
- [] surrounds a class of characters (for example, [abcz]), and matches any one of the characters inside the square brackets (a, b, c, or z).
- [a-z] matches any single character from the range a through z (all lower-case).

Samples

```
bstnA# copy logs syslog smtp://juser@wwmed.com/syslog
```

sends a copy of the syslog to “juser@wwmed.com.” The ARX sends the syslog as an attachment to a brief E-mail message.

```
bstnA# copy reports metadata_only.17.rpt smtp://juser@wwmed.com/mdo.csv
```

sends a copy of a report to the same E-mail recipient. This command reformats the attachment into CSV before sending it (note the “.csv” extension on the destination file name).

Related Commands

copy ftp
copy tftp
copy scp
copy {nfs|cifs}
delete
move
grep
tail
show directories

copy tftp

Purpose Use the copy tftp command to transfer a file (such as a software-release file or a log file) to or from the ARX via TFTP.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

(A backup-operator can upload reports to a remote host.)

Syntax: Download to the ARX

```
copy tftp://server/source-file directory file
```

tftp://*server*/*source-file* (1-1024 characters) is the URL for the source file:

server is the IP address or hostname for the TFTP server.

source-file is the source-file path. Lead with an extra slash ("/") if the path is absolute (for example, "tftp://10.1.1.5**//**var/rels/aco4665.rel"). Use only one slash if the path is local to the "tftpboot" directory. This conforms with the specification for FTP URLs in RFC 1738.

directory is the destination directory. Choose one of the following: **configs**, **scripts**, **license**, or **releases**.

file (1-255 characters) is the name you choose for the copy.

Syntax: Upload

```
copy directory file tftp://server/dest-file
  [format {text|report|csv|xml}]
```

directory is one of the following: **logs**, **stats-logs**, **cores**, **configs**, **replicated-configs**, **reports**, **diag-info**, **software**, **scripts**, **capture**, or **license**.

file (1-255 characters) identifies the source file from the above directory. You can use wildcards (such as *, ?, or [a-z]) to select multiple files; see the *Guidelines: Wildcards* below.

tftp://*server*/*dest-file* (1-255 characters) is the URL for the destination file (see above).

format {text|report|csv|xml} (optional) applies to the "reports" directory only. This chooses the output format for the external copy of the report. The "xml" option converts the report to XML format, "csv" converts the report to a Comma-Separated-Value format (useful for spreadsheets), and the "text" and "report" options leave the reports in plain-text format.

Default(s) None

Guidelines

Use copy scp to upload or download over a secure connection, using SCP. Use copy ftp to upload or download using FTP. To send out a file as an E-mail attachment, use copy smtp. The copy {nfs|cifs} command transfers files between the maintenance directories on the ARX and its client-accessible volumes.

After a copy to the local disks, you can use show directories directory-name to verify that the copy was successful.

To move a file from the local disk to a TFTP server, thereby removing it from the local disk, you can use the move ... tftp command.

To manage local files, use delete and/or rename. To view ASCII files, use show directory file-name, tail, and/or grep.

Guidelines: Wildcards

You can use several wildcard characters to select multiple source-file names for an upload:

- * matches any string, including an empty string.
- ? matches any single character. This is also a special meta character in the CLI, so you must quote the file-path string to use it; for example, copy capture nas?.cap tftp ... would fail, but copy logs "nas?.cap" tftp ... would succeed.
- [] surrounds a class of characters (for example, [abcz]), and matches any one of the characters inside the square brackets (a, b, c, or z).
- [a-z] matches any single character from the range a through z (all lower-case).

Sample

```
bstnA# copy tftp://tftp.f5.com/rel6.0.rel releases r6.0.rel
```

downloads a release file, rel6.0.rel, from the TFTP server at tftp.f5.com.

Related Commands

copy ftp
copy scp
copy smtp
copy {nfs|cifs}
move ... tftp
delete
rename
grep
tail
show directories

delete

Purpose The ARX contains directories for software-release files, log files, core files, report files, diagnostic information, and other maintenance files. Use the delete command to delete all files from a specified directory or to delete a specified file.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
delete {releases | logs | cores | reports | diag-info | scripts
  | capture} [file-name]
delete {configs | license} file-name
```

releases | ... | capture identifies the directory. This is a required choice.

file-name (1-255 characters) identifies a file to delete from the above directory. This is optional for any directory except **configs** or **license**; if you omit it for one of the other directories, the command deletes all files in the chosen directory. You can also use wildcards for those directories (such as *, ?, or [a-z]); see the *Guidelines: Wildcards* below.

Some files in the configs and license directories are required for system operation and/or reboot recovery, so you cannot delete all files (or use wildcards) in those directories. You must specify a file to delete.

Default(s) None

Guidelines The CLI prompts for confirmation before deleting any files; enter **yes** to delete the file(s).

You cannot delete the running-release, armed-release, or backup-release file. You also cannot delete the "active.license" file in the license directory; removing the active license disables all storage services, so it is not recommended. (You can use clear active-license if you have been advised to remove the active license.) Do not delete core files unless advised to do so by F5 personnel; these files contain valuable data for diagnosing software problems.

The show directories command displays a listing of all directories on the ARX. Use the grep, rename, copy, and delete commands to maintain these files.

Guidelines: Wildcards You can use several wildcard characters to select multiple file names:

- * matches any string, including an empty string.
- ? matches any single character. This is also a special meta character in the CLI, so you must quote the file-path string to use it; for example, delete logs traplog.? would fail, but delete logs "traplog.?" would succeed.
- [] surrounds a class of characters (for example, [abcz]), and matches any one of the characters inside the square brackets (a, b, c, or z).
- [a-z] matches any single character from the range a through z (all lower-case).

Samples

```
bstnA# delete logs syslog5.log
Delete file 'syslog5.log' in directory 'logs'? [yes/no] yes
```

deletes the file named “syslog5.log” from the logs directory.

```
bstnA# delete logs syslog*
Delete file 'syslog' in directory 'logs'? [yes/no/all] yes
Delete file 'syslog.1' in directory 'logs'? [yes/no/all] all
bstnA#
```

deletes all syslog files from the logs directory.

```
bstnA# delete cores
Delete all core files in directory 'cores'? [yes/no] yes
```

deletes all files in the cores directory.

Related Commands

copy ftp
copy scp
copy tftp
copy {nfs|cifs}
grep
rename
show directories

grep

Purpose The ARX contains directories for software-release files, log files, core files, and other maintenance files. Use the grep command to filter one of these files, displaying the lines that match a pattern or text string.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax: One File

```
grep filter {releases | logs | stats-logs | cores | configs |
  replicated-configs | reports | software | scripts | capture |
  license} file-name [ ignore ignore ] [ tail lines ]
```

searches through one file.

filter (1-255 characters) is the pattern or text string to search for in the file. See the *Guidelines: Regular Expression* below for the syntax. The grep displays all lines containing this pattern.

releases | ... | license selects the directory. This is a required choice.

file-name (1-255 characters) identifies the file to read.

ignore (optional) filters out text, and

ignore (1-1024 characters) specifies a text string to ignore.

tail (optional) displays the end of the file, and

lines (1-1024) specifies how many lines to show at the end of the file.

Syntax: Directory

```
grep filter {releases | logs | stats-logs | cores | configs |
  replicated-configs | reports | software | scripts | capture |
  license} [ ignore ignore ] [ tail lines ]
```

searches through all files in a directory.

filter (1-255 characters) is the pattern or text string to search for in the directory. See the *Guidelines: Regular Expression* below for the syntax. The grep displays all lines from all files containing this pattern.

releases | ... | license selects the directory. This is a required choice.

ignore (optional) filters out text, and

ignore (1-1024 characters) specifies a text string to ignore.

tail (optional) displays the end of the file, and

lines (1-1024) specifies the number of lines to show.

Default(s) None

Guidelines: Regular Expression

The regular expression syntax follows grep's "basic" regular-expression syntax:

. matches any single character.

.* matches any string, including the null string.

[...] matches any one of the enclosed characters.

[a-z] matches any character in the sorted range, a through z.

**** matches the next character, even if it has special meaning (for example, **\.** matches a period instead of any character).

[^...] matches any character that is *not* enclosed.

Guidelines: Available Files

The show directories command displays maintenance directories of the hard disks on the ARX. These are the directories and files you can search with the grep command. Use the move, copy, and delete commands to maintain these files.

Samples

```
bstnA# grep :POLICY logs syslog
```

searches for the string ":POLICY" in the syslog. Sample output appears in Figure 7.2 on page 7-24.

```
bstnA# grep :POLICY logs syslog ignore init
```

is the same search, but without any lines containing 'init.' Sample output appears in Figure 7.3 on page 7-24.

```
bstnA# grep :POLICY.*\.\.\. logs syslog ignore init
```

searches for the same strings followed by an ellipses (...). Note the use of "\." Sample output appears in Figure 7.4 on page 7-24.

```
bstnA# grep :POLICY.*\.\.\. logs syslog ignore init tail 5
```

is the same as the above example, but shows only the last 5 lines. Sample output appears in Figure 7.5 on page 7-24.

```
prtlndA# grep "prtlndA.*CLI_COMMAND.*show namespace.*wwmed" logs
```

Searches for certain CLI commands in the logs directory. All log files are searched. Sample output appears in Figure 7.6 on page 7-25.

Related Commands

copy ftp
copy scp
copy tftp
copy {nfs|cifs}
copy smtp
delete
move
move ... ftp
move ... {nfs|cifs}
move ... scp
move ... tftp
show directories

Figure 7.2 *Sample Output: grep*

```
bstnA# grep :POLICY logs syslog
2004-05-12T05:42:11.378-0400:bstnA:1-1-SCM-1435:POLICY_PDP-0-6-MSG6:: afnpdpd: Version 0.12.0.4671
(May 12 2004 14:27:17) [dfeng]
2004-05-12T05:42:11.379-0400:bstnA:1-1-SCM-1435:POLICY_PDP-0-6-MSG6:: afnpdpd: (build 4671)
starting up.
2004-05-12T05:42:11.379-0400:bstnA:1-1-SCM-1435:POLICY_PDP-0-6-MSG6:: afnpdpd: Copyright (c)
2002-2004 by Acopia Networks Inc.  All rights reserved.
2004-05-12T05:42:11.380-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: init_daemon
succeeded
2004-05-12T05:42:11.380-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: waiting for OM
GLOBAL Scope ...
2004-05-12T05:42:17.397-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: pdpd: OM GLOBAL Scope
available.
2004-05-12T05:42:25.417-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: initializing IPC
2004-05-12T05:42:25.752-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: IPC OK
2004-05-12T05:42:25.752-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: create IPC thread
2004-05-12T05:42:25.753-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: IPC thread OK
...
```

Figure 7.3 *Sample Output: grep ... ignore*

```
bstnA# grep :POLICY logs syslog ignore init
2004-05-12T05:42:11.378-0400:bstnA:1-1-SCM-1435:POLICY_PDP-0-6-MSG6:: afnpdpd: Version 0.12.0.4671
(May 12 2004 14:27:17) [dfeng]
2004-05-12T05:42:11.379-0400:bstnA:1-1-SCM-1435:POLICY_PDP-0-6-MSG6:: afnpdpd: (build 4671)
starting up.
2004-05-12T05:42:11.379-0400:bstnA:1-1-SCM-1435:POLICY_PDP-0-6-MSG6:: afnpdpd: Copyright (c)
2002-2004 by Acopia Networks Inc.  All rights reserved.
2004-05-12T05:42:11.380-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: waiting for OM
GLOBAL Scope ...
2004-05-12T05:42:17.397-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: pdpd: OM GLOBAL Scope
available.
2004-05-12T05:42:25.752-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: IPC OK
2004-05-12T05:42:25.752-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: create IPC thread
2004-05-12T05:42:25.753-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: IPC thread OK
...
```

Figure 7.4 *Sample Output: grep Using '\'*

```
bstnA# grep :POLICY.*\.\.\. logs syslog ignore init
2004-05-12T05:42:11.380-0400:bstnA:1-1-SCM-1439:POLICY_PDP-0-7-MSG7:: afnpdpd: waiting for OM
GLOBAL Scope ...
2004-05-12T05:42:27.006-0400:bstnA:1-1-SCM-1552:POLICY_PDP-0-7-MSG7:: afnpdpd: clean rules thread
starting ...
2004-05-12T05:42:28.024-0400:bstnA:1-1-SCM-1552:POLICY_PDP-0-7-MSG7:: afndpdpd : scrubbing rule
table...
2004-05-12T05:47:28.009-0400:bstnA:1-1-SCM-1552:POLICY_PDP-0-7-MSG7:: afndpdpd : scrubbing rule
table...
2004-05-12T05:52:28.008-0400:bstnA:1-1-SCM-1552:POLICY_PDP-0-7-MSG7:: afndpdpd : scrubbing rule
table...
2004-05-12T05:57:28.010-0400:bstnA:1-1-SCM-1552:POLICY_PDP-0-7-MSG7:: afndpdpd : scrubbing rule
table...
...
```

Figure 7.5 *Sample Output: grep ... tail*

```
bstnA# grep :POLICY.*\.\.\. logs syslog ignore init tail 5
2004-05-13T02:12:30.063-0400:bstnA:1-1-SCM-1517:POLICY_PDP-0-7-MSG7:: afndpdpd : scrubbing rule
table...
2004-05-13T02:17:30.028-0400:bstnA:1-1-SCM-1517:POLICY_PDP-0-7-MSG7:: afndpdpd : scrubbing rule
table...
```

```
2004-05-13T02:22:30.025-0400:bstnA:1-1-SCM-1517:POLICY_PDP-0-7-MSG7:: afndpdpd : scrubbing rule
table...
2004-05-13T02:27:30.031-0400:bstnA:1-1-SCM-1517:POLICY_PDP-0-7-MSG7:: afndpdpd : scrubbing rule
table...
2004-05-13T02:32:30.027-0400:bstnA:1-1-SCM-1517:POLICY_PDP-0-7-MSG7:: afndpdpd : scrubbing rule
table...
bstnA#
```

Figure 7.6 *Sample Output: grep (directory)*

```
prtlndA# grep “prtlndA.*CLI_COMMAND.*show namespace.*wwmed” logs
2004-05-07T08:21:29.207-0400:prtlndA:1-1-ACM-3177:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T08:21:39.147-0400:prtlndA:1-1-ACM-3177:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T08:21:49.107-0400:prtlndA:1-1-ACM-3177:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T08:21:59.007-0400:prtlndA:1-1-ACM-3177:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T08:22:08.937-0400:prtlndA:1-1-ACM-3177:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T08:22:09.657-0400:prtlndA:1-1-ACM-3177:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace wwmed
2004-05-07T12:59:45.102-0400:prtlndA:1-1-ACM-3170:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T12:59:55.062-0400:prtlndA:1-1-ACM-3170:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T13:00:05.022-0400:prtlndA:1-1-ACM-3170:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T13:00:14.922-0400:prtlndA:1-1-ACM-3170:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T13:00:15.682-0400:prtlndA:1-1-ACM-3170:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace wwmed
2004-05-07T13:47:38.088-0400:prtlndA:1-1-ACM-3180:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T13:47:48.038-0400:prtlndA:1-1-ACM-3180:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T13:47:57.968-0400:prtlndA:1-1-ACM-3180:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T13:48:07.908-0400:prtlndA:1-1-ACM-3180:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T13:48:17.918-0400:prtlndA:1-1-ACM-3180:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace status wwmed
2004-05-07T13:48:18.668-0400:prtlndA:1-1-ACM-3180:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA(gbl)# show namespace wwmed
2004-05-07T14:21:36.589-0400:prtlndA:1-1-ACM-20578:SCM_CLI-0-6-CLI_COMMAND:: User admin: Command:
prtlndA# grep “prtlndA.*CLI_COMMAND.*show namespace.*wwmed” logs syslog
prtlndA#
```

ip ftp-user

Purpose

When you use the copy or move command to transport files to/from an FTP site, you enter a username and password for the FTP server. Use the ip ftp-user command to set a default username and password.

Use the no form of this command to revert to the "anonymous" default.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
ip ftp-user name
no ip ftp-user
```

name (1-32 characters) is the FTP username.

Default(s)

name - anonymous

password - upgrade-*hostname*, where you can set the hostname with the hostname command.

Guidelines

The CLI prompts twice for a password. See the sample below.

This command makes the copy ftp and move ... ftp commands easier to use; you can omit the username and password from each of these commands, defaulting to this one.

Sample

```
bstnA(cfg)# ip ftp-user jsmith
Password: jpasswd
Validate Password: jpasswd
bstnA(cfg)# ...
```

provides a username and password, jsmith and jpasswd, for transferring files to/from an FTP server.

Related Commands

copy ftp
move ... ftp

ip ron-user

Purpose

When you use the copy ron command to copy a file onto another ARX, you require an administrative username and password for the remote peer. Use the ip ron-user command to set a default username and password.

Use the no form of this command to revert to using the current administrative account as the default.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
ip ron-user name
no ip ron-user
```

name (1-32 characters) is a valid administrative user at a remote ARX.

Default(s)

If this is not set, the copy ron command uses the username and password of the current administrative account.

Guidelines

The CLI prompts twice for a password. See the sample below.

Sample

```
bstnA(cfg)# ip ron-user jsmith
Password: jpasswd
Validate Password: jpasswd
bstnA(cfg)# ...
```

provides a username and password, jsmith and jpasswd, for sending file copies to a remote ARX.

Related Commands

copy ron

ip scp-user

Purpose

When you use the copy or move command to transfer files to/from an SCP server, you enter a username and password for the SCP server. Use the ip scp-user command to set a default username and password for these operations.

Use the no form of this command to remove the default.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
ip scp-user name
no ip scp-user
```

name (1-32 characters) is a Unix username that is valid on the remote SCP server(s).

Default(s)

None.

Guidelines

The CLI prompts twice for a password. See the sample below.

This command makes the copy scp and move ... scp commands easier to use; you can omit the username and password from each command, defaulting to this one.

Sample

```
bstnA(cfg)# ip scp-user jsmith
Password: jpasswd
Validate Password: jpasswd
bstnA(cfg)# ...
```

provides a username and password, jsmith and jpasswd, for copying to/from an SCP server.

Related Commands

copy scp
move ... scp

move

Purpose The ARX contains directories for software release files, log files, report files, and other maintenance files. Use the move command to change the name of a file in one of these directories.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
move {logs | cores | configs | replicated-configs | diag-info |
  software | reports | scripts | capture | license}
  src-file-name dest-file-name
```

logs | ... | license identifies the directory. This is required.

logs selects a log file. Use the show logs command for a full list of all available log files.

cores selects a core-dump file generated by a failing process. The show cores command lists all such files that currently exist, if any.

configs selects a configuration file, such as those generated by copy running-config and copy global-config. Use show configs to list them.

replicated-configs selects a replicated-configuration file, from a remote ARX cluster. You can use these files for a disaster-recovery operation. Use show replicated-configs to list them.

diag-info selects a diagnostics file generated by collect or collect logs. Use show diag-info to see if any are available.

software selects a software-related file or a software-user manual. The show software command lists all of them.

reports selects a report file.

scripts selects a script file. Use show scripts to see all available scripts.

capture selects a packet-capture file created by the capture session command. The show capture command lists all packet-capture files currently on the ARX.

license selects a license file. The show license command lists all license-related files currently on the ARX.

src-file-name (1-255 characters) is the original file name, the file to rename.

dest-file-name (1-255 characters) is the new name you choose for the file.

Default(s) None

Guidelines The show directories command displays all files and directories on the ARX hard disks.

You can also move files to an external servers. The move ... ftp command uploads over FTP, then deletes the source file after a successful copy operation. Use move ... scp to use a secure connection, SCP. The move ... tftp command uses TFTP (Trivial FTP) as a transport. The move ... {nfs|cifs} command transfers files from the maintenance directories on the ARX to its client-accessible volumes.

Use the grep, copy, and delete commands to maintain these files.

Sample

```
bstnA# move diag-info ns_wwmed.tgz wwmed_11_12.tgz
```

renames a diagnostics file.

Related Commands

move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
copy ftp
copy scp
copy tftp
copy smtp
copy {nfs|cifs}
grep
delete
show directories

move ... ftp

Purpose Use the move ... ftp command to transfer a file (such as a core-dump file) off of the ARX via FTP. This command deletes the file after it is successfully copied, or leaves the file at the source location if the copy operation fails.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

(A backup-operator can move reports.)

Syntax

```
move directory file ftp://[user:password@]server/dest-file
  [format {text|report|csv|xml}]
```

directory is one of the following: **logs**, **cores**, **configs**, **replicated-configs**, **reports**, **diag-info**, **software**, **scripts**, **capture**, or **license**.

file (1-255 characters) identifies the source file from the above directory.

ftp://[*user*:*password*@]*server*/*dest-file* (1-255 characters) is the URL for the destination file:

> ***user*[:*password*]@** (optional) are the credentials for FTP access. If you omit them, they default to the credentials set by the ip ftp-user command. If you omit the password, the CLI prompts for one.
>
> ***server*** is the IP address or hostname for the FTP server.
>
> ***dest-file*** is the new path that you want for the file. Lead with an extra slash ("/") if the path is absolute (for example, "ftp://10.1.1.5**//**var/caps/arx0107.cap"). Use only one slash if the path is local to the home directory for ***user*** (for example, "ftp://10.1.1.5**/**gffstnA.tgz"). This conforms with the specification for FTP URLs in RFC 1738.

format {text|report|csv|xml} (optional) applies to the "reports" directory only. This chooses the output format for the report. The "xml" option converts the report to XML format, "csv" converts the report to a Comma-Separated-Value format (useful for spreadsheets), and the "text" and "report" options leave the reports in plain-text format.

Default(s) *user*:*password* from the URL - the default set by ip ftp-user.

Guidelines

If you enter a user name without a password, the CLI prompts for the password before continuing. Enter the password for the *user* you included in the URL.

To copy a file over FTP instead of moving it, you can use the copy ftp command.

You can also move (or copy) files using several additional protocols. Use move ... scp or copy scp to use an SCP connection for the move or copy operation. The move ... tftp and copy tftp commands use TFTP (Trivial FTP) as a transport. To send out a file as an E-mail attachment, use copy smtp. The move ... {nfs|cifs} and copy {nfs|cifs} commands transfer files between the maintenance directories on the ARX and its client-accessible volumes. To rename a file in its current ARX directory, you can use the move command.

The move commands only upload to an external server. The equivalent copy commands can also download. For example, the copy ftp command FTP-copies a file to or from the ARX.

To manage local files, use delete and/or move. To view ASCII files, use show directory file-name, tail, and/or grep.

Guidelines: Reformatting Reports to XML or CSV

For move operations out of the reports directory, you can reformat the file to XML or CSV (comma-separated value) format. Use a .xml or .csv extension for the *destination-file-name*, and the move operation converts the file to the chosen format. A .txt or .rpt extension leaves the file in plain-text format.

Sample

```
bstnA# move capture fsrvr.cap ftp://noc.phredco.com/bstnA_fsrvr.cap
```

moves a packet-capture file off of the ARX and onto the FTP server at noc.phredco.com. This example uses the default username and password, set by ip ftp-user.

Related Commands

ip ftp-user
move
move ... scp
move ... tftp
move ... {nfs|cifs}
copy ftp
copy scp
copy tftp
copy smtp
copy {nfs|cifs}
delete
grep
tail
show directories

move ... {nfs|cifs}

Purpose Use the move ... {nfs|cifs} command to transfer a file (such as a software-release file or a log file) from the ARX to one of its volumes. This command deletes the file after it is successfully copied to the volume, or leaves the file at the source location if the copy operation fails.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

(A backup-operator can move reports.)

Syntax

```
move directory file {nfs|cifs} namespace vol [dest-file-path]
  [format {text|report|csv|xml}]
```

directory is one of the following: **logs**, **cores**, **configs**, **replicated-configs**, **reports**, **diag-info**, **software**, **scripts**, **capture**, or **license**.

file (1-255 characters) identifies the source file from the above directory. You can use wildcards (such as *, ?, or [a-z]) to select multiple files; see the *Guidelines: Wildcards* below.

nfs | cifs is a required choice. This chooses the protocol for the file transfer.

namespace (1-30 characters) identifies the namespace to use as a destination.

vol (1-1024 characters) is the volume where the file will go.

dest-file-path (optional, 1-255 characters) is the destination path for the file, starting at the volume root. If you omit this, the file goes to the volume root and keeps its original name.

format {text|report|csv|xml} (optional) applies to the "reports" directory only. This chooses the output format for the report. The "xml" option converts the report to XML format, "csv" converts the report to a Comma-Separated-Value format (useful for spreadsheets), and the "text" and "report" options leave the reports in plain-text format.

Default(s) None

Guidelines

The CLI chooses a user identity to establish its read/write privileges in the ARX volume. It chooses its identity based on the transfer protocol. For NFS, the CLI uses *root* as its identity. For CIFS, the CLI uses the proxy-user (gbl-ns) for the chosen namespace. These identities typically ensure that there are no permissions problems during the copy operation.

To copy a file to a volume instead of moving it, you can use the copy {nfs|cifs} command.

You can also move (or copy) files using several additional protocols.Use move ... scp or copy scp to use an SCP connection for the move or copy operation. The move ... tftp and copy tftp commands use TFTP (Trivial FTP) as a transport. To send out a file as an E-mail attachment, use copy smtp. The move ... ftp and copy ftp commands use FTP to send files to an external server. To rename a file in its current ARX directory, you can use the move command.

The move commands only upload to an external server. The equivalent copy commands can also download. For example, the copy {nfs|cifs} command copies a file to or from an ARX volume.

To manage local files, use delete and/or move. To view ASCII files, use show directory *file-name*, tail, and/or grep.

Guidelines: Wildcards

You can use several wildcard characters to select multiple file names:

- * matches any string, including an empty string.
- ? matches any single character. This is also a special meta character in the CLI, so you must quote the file-path string to use it; for example, move diag-info collNum?.tgz cifs medarcv /rcrds would fail, but move diag-info “collNum?.tgz” cifs medarcv /rcrds would succeed.
- [] surrounds a class of characters (for example, [abcz]), and matches any one of the characters inside the square brackets (a, b, c, or z).
- [a-z] matches any single character from the range a through z (all lower-case).

Sample

```
bstnA# move cores core-* nfs wwmed /acct/.admin/
```

moves all core-dump files from the local “cores” directory to a hidden directory in “wwmed~/acct,” an NFS volume.

Related Commands

move
move ... ftp
move ... scp
move ... tftp
copy ftp
copy scp
copy tftp
copy smtp
copy {nfs|cifs}
delete
rename
grep
tail
show directories

move ... scp

Purpose Use the move ... scp command to securely copy a file (such as a software-release file or a log file) to or from the ARX over SCP. This command deletes the file after it is successfully transferred, or leaves the file at the source location if the copy operation fails.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

(A backup-operator can move reports.)

Syntax

```
move directory file scp://[user@]server:dest-file
  [accept-host-key] [format {text|report|csv|xml}]
```

directory is one of the following: **logs**, **cores**, **configs**, **replicated-configs**, **reports**, **diag-info**, **software**, **scripts**, **capture**, or **license**.

file (1-255 characters) identifies the file to move out of the above directory.

scp://[*user*@]*server*:*dest-file* (1-255 characters) is the URL for the file's destination:

> ***user*@** (optional if someone created an ip scp-user) is the username to present to the other end of the SCP connection. This user must be valid at the remote host. If you omit this and an ip scp-user is defined, it defaults to the username set by that command.
>
> ***server*:** is the IP address or hostname for the SCP host. End with a colon (:).
>
> ***dest-file*** is the desired path for the file at the remote host. Lead with a slash ("/") if the path is absolute (for example, "scp://root@10.1.1.5:**/var/arxDiags/6_5.tgz**"). Use no slash if the path is local to the home directory for ***user*** (for example, "scp://root@10.1.1.5:**6_5.tgz**").

accept-host-key (optional) means to accept an unknown SSH host key from the remote host, if it presents one.

format {text|report|csv|xml} (optional) applies to the "reports" directory only. This chooses the output format for the report. The "xml" option converts the report to XML format, "csv" converts the report to a Comma-Separated-Value format (useful for spreadsheets), and the "text" and "report" options leave the reports in plain-text format.

Default(s) None

Guidelines If you enter a username in this command, the CLI prompts you for a password. Enter the password for the *user* in the URL. The file is transferred over SCP, which uses a secure, encrypted connection to the remote host.

Every host identifies itself with an encrypted SSH key. Whenever the ARX transfers a file using SCP, it exchanges keys with the remote host. If the remote-host key is different from one received previously, the host may not be the one you intended; an attacker may be posing as the intended host. For this reason, the ARX does not download or upload to a host with an unknown key unless you raise the **accept-host-key** flag. Note that the first contact with any server will require this flag to succeed.

The ARX deletes all SSH host keys on reboot, so all remote hosts are again unknown when the switch comes back up.

Guidelines: Additional Commands for Moving and Copying Files

To copy a file over SCP instead of moving it, you can use the copy scp command.

You can also move (or copy) files using several additional protocols. The move ... tftp and copy tftp commands use TFTP (Trivial FTP) as a transport. To send out a file as an E-mail attachment, use copy smtp. The move ... ftp and copy ftp commands use FTP to send files to an external server. Use move ... {nfs|cifs} or copy {nfs|cifs} to move or copy a file to a volume on the current ARX. To rename a file in its current ARX directory, you can use the move command. The move commands only upload to an external server. The equivalent copy commands can also download. For example, the copy scp command copies a file to or from an external server.

To manage local files, use delete and/or move. To view ASCII files, use show directory *file-name*, tail, and/or grep.

Sample

```
bstnA# move capture ntap.cap scp://root@noc:/var/logs/arx2ntap.cap
accept-host-key
Password: root-password
```

uploads a packet-capture file to a host named "noc." If the host key has changed for noc (or this is the first SCP to or from noc), the switch accepts the key and continues the upload.

(For details on creating a packet-capture file, see the documentation for the capture session command.)

Related Commands

ip scp-user
move
move ... ftp
move ... tftp
move ... {nfs|cifs}
copy scp
copy ftp
copy tftp
copy smtp
copy {nfs|cifs}
delete
grep
tail
show directories

move ... tftp

Purpose Use the move ... tftp command to transfer a file (such as a packet-capture file) off of the ARX via Trivial FTP (TFTP). This command deletes the file after it is successfully transferred, or leaves the file at the source location if the copy operation fails.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

(A backup-operator can upload reports to a remote host.)

Syntax

```
move directory file tftp://server/dest-file
  [format {text|report|csv|xml}]
```

directory is one of the following: **logs**, **cores**, **configs**, **replicated-configs**, **reports**, **diag-info**, **software**, **scripts**, **capture**, or **license**.

file (1-255 characters) identifies the source file from the above directory.

tftp://*server*/*dest-file* (1-255 characters) is the URL for the destination file:

> ***server*** is the IP address or hostname for the TFTP server.
>
> ***dest-file*** is the new path for the file. Lead with an extra slash ("/") if the path is absolute (for example, "tftp://10.1.1.5//var/diags/arxCllct.tgz"). Use only one slash if the path is local to the "tftpboot" directory. This conforms with the specification for FTP URLs in RFC 1738.

format {text|report|csv|xml} (optional) applies to the "reports" directory only. This chooses the output format for the report. The "xml" option converts the report to XML format, "csv" converts the report to a Comma-Separated-Value format (useful for spreadsheets), and the "text" and "report" options leave the report in plain-text format.

Default(s) None

Guidelines To copy a file over TFTP instead of moving it, you can use the copy tftp command.

You can also move (or copy) files using several additional protocols. The move ... scp and copy scp commands use SCP (Secure CoPy) as a transport. To send out a file as an E-mail attachment, use copy smtp. The move ... ftp and copy ftp commands use FTP to send files to an external server. Use move ... {nfs|cifs} or copy {nfs|cifs} to move or copy a file to a volume on the current ARX. To rename a file in its current ARX directory, you can use the move command. The move commands only upload to an external server. The equivalent copy commands can also download. For example, the copy tftp command copies a file to or from an external server.

To manage local files, use delete and/or move. To view ASCII files, use show directory *file-name*, tail, and/or grep.

Sample

```
bstnA# move capture 11_12.cap tftp://tftp.f5.com/11_12_arx.cap
```

sends a packet-capture file, 11_12.cap, to the TFTP server at tftp.f5.com.

(For details on creating a packet-capture file, see the documentation for the capture session command.)

Related Commands

move
move ... ftp
move ... scp
move ... {nfs|cifs}
copy ftp
copy tftp
copy scp
copy smtp
copy {nfs|cifs}
delete
grep
tail
show directories

pause

Purpose This command is useful in writing/running scripts when you want to pause the CLI temporarily.

Mode (any)

Security Role(s) operator

Syntax `pause [seconds]`

seconds (optional; 0-3600) is the number of seconds to pause the CLI operation. If you omit this argument, the **<Enter>** key stops the pause.

Default(s) 0 (indefinite pause, until you press **<Enter>**)

Guidelines You can use this command with terminal clear, which is also useful in writing scripts. For example, you could write a script that lets you view share status during a managed-volume import:

```
terminal clear
show share status
pause 60
terminal clear
show share status
pause 60
terminal clear
show share status
pause 60
terminal clear
show share status
pause 60
terminal clear
show share status
```

Samples

```
bstnA# pause 300
```

pauses the CLI for 300 seconds (five minutes).

```
bstnA# pause
Press <enter> to continue.
```

pauses the CLI until you press **<Enter>**.

Related Commands run
terminal clear

remark

Purpose Use this command to enter a text-string comment into the syslog.

Mode (any)

Security Role(s) operator

Syntax **`remark {critical | error | warning | notice | info | debug}`** ***`comment-string`***

critical | error | warning | notice | info | debug is a required choice. This sets the severity of the message.

comment-string (1-255 characters) is the message text. Insert quotation marks around the string if it contains any spaces.

Default(s) None

Guidelines Use show logging levels to verify the log-level settings for each component.

From any mode, use show logs syslog or grep pattern logs syslog to view the log messages in the syslog file. See the manual, ARX Log Catalog, for a full list of log messages.

Sample

```
bstnA# remark critical "J. Random bringing down NFS"
```

enters a critical message into the syslog.

Related Commands show logging levels
pause
grep

rename

Purpose The ARX contains directories for software release files, log files, report files, and other maintenance files. Use the rename command to change the name of a file in one of these directories.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
rename {releases | logs | cores | reports | scripts | capture |
  license} src-file-name dest-file-name
```

releases | ... | license identifies the directory. This is required.

logs selects a log file.

cores selects a core dump file.

reports selects a report file.

scripts selects a script file.

capture selects a packet-capture file created by the capture session command.

license selects a license file.

src-file-name (1-255 characters) is the original file name, the file to rename.

dest-file-name (1-255 characters) is the new name you choose for the file.

Default(s) None

Guidelines The CLI prompts for confirmation before renaming a file; enter **yes** to proceed with the rename operation.

The show directories command displays all files and directories on the ARX hard disks.

Use the grep, copy, rename, and delete commands to maintain these files.

Sample

```
bstnA# rename releases rel8.rel fallback.rel

Rename file 'rel8.rel' in directory 'releases' to 'fallback.rel'?
[yes/no] yes
```

renames a release file.

Related Commands

grep
copy ftp
copy scp
copy tftp
copy {nfs|cifs}
delete
show directories

run

Purpose This command enables you to run a CLI script.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `run {scripts | configs} script-name`

scripts | configs selects the directory that holds the script. Use show scripts or show configs for file listings in these directories.

script-name (1-255 characters) identifies the CLI script file to run. Use show scripts to display the scripts directory contents.

Default(s) None

Guidelines Use this command to run a CLI script. You can only run a script that is in the scripts or configs directory. The script is comprised of CLI commands, one per line. Comments start with a semi-colon (;).

To download a CLI-script file from an FTP, SCP, TFTP or other network service, use the copy command. If you want to run the script periodically, on a schedule, use the at ... do run command.

There are several CLI commands (such as attach) that do not get committed until you exit the mode with either exit or end. The best practice is to use the end command at the end of any script; you can use this command from any mode under cfg or gbl mode.

If you auto-generate a script (for example, with show running-config or show global-config), you must edit the script before you run it on an earlier release of ARX software. Command syntax may have changed as well as the best practices for command order. Consult the documentation from the earlier release (available from the GUI) and change the script's commands as needed before you use this command to run it.

Sample

```
bstnA# run scripts testdemo1.scr
```

runs the script, "testdemo.scr."

Related Commands

copy ftp
copy scp
copy tftp
copy {nfs|cifs}
show scripts
show configs
show software
show directories

show at

Purpose Use the show at command to display all scheduled CLI commands and/or scripts.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show at`

Guidelines Use the at command to schedule a CLI job to run in the future.

This command displays two rows per scheduled CLI job. Each pair of rows contains the following columns:

Job is a numeric ID for the scheduled job. This is assigned by the CLI when the user issues the at command. You can use this ID with clear at to clear the job.

NextTime is the date and time for the next scheduled job run. This is in the following format: ***mm/dd/yyyy.HH:MM***.

Interval is the time between runs, if there is one. This is the every interval part of the at command.

Command is the exact CLI command (or script) that is invoked at each run.

Report is the prefix for the report generated at each run, if there is one. Each report is named as follows: *report-prefix_yyyymmddHHMM*.rpt. The part after the underscore is the date and time of the run. Use show reports to view the reports, if there are any.

Sample `bstnA# show at`
shows all scheduled commands and scripts. See Figure 7.7 for sample output.

Related Commands at
clear at

Figure 7.7 *Sample Output: show at*

```
bstnA# show at

Job    NextTime         Interval      Command
  Report
-------------------------------------------------------------------------------

1     03/31/2009.00:49 1 Days      nis update

2     03/30/2009.03:30 89 Days     cifs rekey all

4     03/30/2009.01:34 5 Min       show sessions
    adminSessions
5     03/31/2009.00:00 1 Days      copy startup-config
ftp://root:rootpw@172.16.100.183//tmp/acocfg.conf

6     03/30/2009.03:45 1 Days      active-directory update forest MEDARCH.ORG proxy-user acoProxy2

7     03/30/2009.05:00 1 Days      active-directory update forest NY.COM proxy-user ny_admin

8     03/31/2009.01:19 1 Days      copy reports snap*
ftp://ftpuser:ftpuser@172.16.100.183//var/arxSnapRpts/ format xml
```

show capture

Purpose The capture directory stores all files created by capturing IP traffic. Use this command to display the contents of the capture directory and/or the contents of a specified capture file.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show capture [`*`file-name`* `[summary [cifs|non-cifs]]]`

file-name (optional, 1-255 characters) is the name of the file to display. If you omit this, the CLI shows a list of all files in the directory.

summary (optional) reduces the output to a group of summary tables, like those shown with **Tshark -z**.

cifs | non-cifs (optional) is a filter on the summary output. If you choose **cifs**, the summary only counts packets related to CIFS traffic. These are packets to or from UDP/88, TCP/88, UDP/137, UDP/138, TCP/139, and/or TCP/445. The **non-cifs** option summarizes all packets *except* the ones to or from those ports.

Guidelines You can capture IP traffic into a file with the capture session command. This command shows a list of all captured files or shows the contents of one of them.

The contents of the capture file are similar to the output of TShark, a commonly available network-packet analyzer. The TShark program is built on the same Packet Capture (PCap) library as WireShark, a similar program with a graphical interface.

The capture summary is similar to the output of **Tshark -z**. This shows a table of UDP conversations and TCP conversations, followed by RTT statistics for NFSv2, NFSv3, and SMB (or CIFS). The TCP and UDP tables have one row per pairing of IP addresses. The NFS tables do not appear if you use the **cifs** option, and the SMB tables do not appear if you enter the **non-cifs** option.

Use the show directories command to show the contents of all directories on the ARX. To maintain this directory, use the copy, move, grep, and delete commands.

While a capture session is underway, the ARX writes to its output file each time it captures IP packets. If the capture session writes to multiple files, it creates them all with a time stamp in each of their names. Use show capture sessions for a list of capture sessions that are currently underway.

Samples

```
bstnA> show capture
```

```
  capture
      cifsVol.cap             06/29 01:55   220 kB
      clientCap.cap           06/29 01:45   15 MB
      fsrvr.cap               06/29 02:12   224 kB
      nasTraffic.cap          06/29 01:55   212 kB
      ntaps_00001_20110629060427.cap 06/29 02:04   24 B
      proxyTraffic_01620_20110629055529.cap 06/29 01:55   150 kB
      proxyTraffic_01621_20110629055531.cap 06/29 01:55   94 kB
```

lists the directory contents.

```
bstnA> show capture nasTraffic_00001_20060328081103.cap
```

Shows the contents of a capture file. See Figure 7.8 for sample output.

```
bstnA> show capture nasTraffic.cap summary
```

Summarizes the contents of a capture file. See Figure 7.9 on page 7-46 for sample output.

Related Commands

capture session
show capture sessions
copy ftp, copy scp, copy {nfs|cifs}, copy tftp, and copy smtp
move, move ... ftp, move ... scp, move ... {nfs|cifs}, move ... tftp, and move ... scp
delete
show directories

Figure 7.8 *Sample Output: show capture*

```
bstnA# show capture nasTraffic.cap
  1   0.000000  192.168.25.21 -> 192.168.25.141 TCP 2049 > 652 [PSH, ACK] Seq=0 Ack=0 Win=26280
Len=172
  2   1.344955  192.168.25.21 -> 192.168.25.33 UDP Source port: 2049  Destination port: 640
  3   1.346355  192.168.25.21 -> 192.168.25.141 UDP Source port: 2049  Destination port: 640
  4   1.346602  192.168.25.21 -> 192.168.25.141 UDP Source port: 2049  Destination port: 640
  5   1.347700  192.168.25.21 -> 192.168.25.141 UDP Source port: 2049  Destination port: 640
  6   1.347899  192.168.25.21 -> 192.168.25.141 UDP Source port: 2049  Destination port: 640
  7   1.348156  192.168.25.21 -> 192.168.25.141 UDP Source port: 2049  Destination port: 640
  8   1.518532  192.168.25.21 -> 192.168.25.33 UDP Source port: 2049  Destination port: 640
  9   1.523073  192.168.25.21 -> 192.168.25.33 UDP Source port: 2049  Destination port: 640
 10   1.524320  192.168.25.21 -> 192.168.25.141 UDP Source port: 2049  Destination port: 640
 11   1.540995  192.168.25.21 -> 192.168.25.141 UDP Source port: 2049  Destination port: 640
 12   1.554278  192.168.25.21 -> 192.168.25.33 UDP Source port: 2049  Destination port: 640
...
302 137.728365 192.168.25.31 -> 192.168.25.21   NFS V3 FSSTAT Call, FH:0x4f2885a4
303 137.728617 192.168.25.31 -> 192.168.25.21   NFS V3 FSSTAT Call, FH:0x9a53d568
304 137.729335 192.168.25.31 -> 192.168.25.21   SMB Trans2 Request, QUERY_FS_INFO, Query FS Size
Info
305 137.768215 192.168.25.31 -> 192.168.25.21   TCP 10294 > netbios-ssn [ACK] Seq=1192 Ack=1260
Win=14480 Len=0 TSV=202118 TSER=2655943
...
```

Figure 7.9 *Sample Output: show capture ... summary*

```
bstnA# show capture nasTraffic.cap summary
================================================================================
UDP Conversations
Filter:<No Filter>
                                                |       <-      | |       ->      | |     Total     |
                                                | Frames  Bytes | | Frames  Bytes | | Frames  Bytes |
192.168.25.21:2049   <-> 192.168.25.33:640          0         0     894    183604     894    183604
192.168.25.21:2049   <-> 192.168.25.141:640        0         0      21      3978      21      3978
================================================================================
================================================================================
TCP Conversations
Filter:<No Filter>
                                                |       <-      | |       ->      | |     Total     |
                                                | Frames  Bytes | | Frames  Bytes | | Frames  Bytes |
192.168.25.141:15231 <-> 192.168.25.21:445         15     10227       0         0      15     10227
192.168.25.148:51031 <-> 192.168.25.21:445          8      1113       0         0       8      1113
192.168.25.21:2049   <-> 192.168.25.141:641         0         0       1       226       1       226
192.168.25.141:15226 <-> 192.168.25.21:445          1       118       0         0       1       118
192.168.25.21:2049   <-> 192.168.25.141:653         0         0       1        86       1        86
192.168.25.21:2049   <-> 192.168.25.141:654         0         0       1        86       1        86
192.168.25.21:2049   <-> 192.168.25.141:652         0         0       1       226       1       226
================================================================================

====================================================================
NFS Version 2 RTT Statistics:
Filter:
Procedure        Calls   Min RTT   Max RTT   Avg RTT
NULL                 0   0.00000   0.00000   0.00000
GETATTR              0   0.00000   0.00000   0.00000
SETATTR              0   0.00000   0.00000   0.00000
ROOT                 0   0.00000   0.00000   0.00000
LOOKUP               0   0.00000   0.00000   0.00000
READLINK             0   0.00000   0.00000   0.00000
READ                 0   0.00000   0.00000   0.00000
WRITECACHE           0   0.00000   0.00000   0.00000
WRITE                0   0.00000   0.00000   0.00000
CREATE               0   0.00000   0.00000   0.00000
REMOVE               0   0.00000   0.00000   0.00000
RENAME               0   0.00000   0.00000   0.00000
LINK                 0   0.00000   0.00000   0.00000
SYMLINK              0   0.00000   0.00000   0.00000
MKDIR                0   0.00000   0.00000   0.00000
RMDIR                0   0.00000   0.00000   0.00000
READDIR              0   0.00000   0.00000   0.00000
STATFS               0   0.00000   0.00000   0.00000
====================================================================

====================================================================
NFS Version 3 RTT Statistics:
Filter:
Procedure        Calls   Min RTT   Max RTT   Avg RTT
NULL                 0   0.00000   0.00000   0.00000
GETATTR              0   0.00000   0.00000   0.00000
SETATTR              0   0.00000   0.00000   0.00000
LOOKUP               0   0.00000   0.00000   0.00000
ACCESS               0   0.00000   0.00000   0.00000
READLINK             0   0.00000   0.00000   0.00000
READ                 0   0.00000   0.00000   0.00000
WRITE                0   0.00000   0.00000   0.00000
CREATE               0   0.00000   0.00000   0.00000
MKDIR                0   0.00000   0.00000   0.00000
```

```
SYMLINK              0   0.00000   0.00000   0.00000
MKNOD                0   0.00000   0.00000   0.00000
REMOVE               0   0.00000   0.00000   0.00000
RMDIR                0   0.00000   0.00000   0.00000
RENAME               0   0.00000   0.00000   0.00000
LINK                 0   0.00000   0.00000   0.00000
READDIR              0   0.00000   0.00000   0.00000
READDIRPLUS          0   0.00000   0.00000   0.00000
FSSTAT               0   0.00000   0.00000   0.00000
FSINFO               0   0.00000   0.00000   0.00000
PATHCONF             0   0.00000   0.00000   0.00000
COMMIT               0   0.00000   0.00000   0.00000
=====================================================================

=====================================================================
SMB RTT Statistics:
Filter:
Commands                  Calls   Min RTT   Max RTT   Avg RTT

Transaction2 Commands     Calls   Min RTT   Max RTT   Avg RTT

NT Transaction Commands   Calls   Min RTT   Max RTT   Avg RTT
=====================================================================
```

show configs

Purpose Use this command to display the configs directory, which contains the switch's configuration files. You can also use this command to read one file from the directory.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show configs [file-name]`

file-name (optional, 1-255 characters) is the name of the file to display. If you omit this, the CLI shows a list of the directory's files.

Guidelines Use the show directories command to show the contents of all directories on the ARX. To maintain this directory, use the copy, move, grep, and delete commands.

The following files commonly appear in the configs directory:

- acopia.keystore is the SSH key for HTTPS connections to the GUI. This is the GUI's "security certificate" that is discussed in a pop-up when you first connect to the GUI. This file is in binary format, so you cannot read it. If you delete this file and restart the GUI (using gui restart), the ARX generates a new, self-signed, security certificate.
- omDbVersion.info contains the last understood database version. During a software upgrade, the ARX uses this file to determine if it requires a change in the database schema.
- boot-config contains the configuration parameters for setting up the administrative account created in the initial-boot script. This is a text file that can be run as a CLI script. The ARX runs this after it boots if (and only if) the startup-config is missing.
- startup-config is binary file with all configuration parameters on the ARX. The ARX uses this after it boots to recreate its entire configuration. It ignores the boot-config file if this file is present. You cannot view this file directly, though you can use copy startup-config, copy running-config, or copy global-config to copy the startup-config (or its components, the running-config and global-config) to a CLI script.

Sample

```
bstnA# show configs

  configs
      wwmed.keystore          06/16 10:33   1.3 kB
      active.license          06/29 01:08   9.3 kB
      arx.dossier             04/11 22:54   3.0 kB
      arx.regkey              06/29 01:08   34 B
      boot-config             06/29 01:15   1.5 kB
      eeprom.dat              06/29 00:17   119 B
      fbstat.out              06/29 02:05   41 kB
      nlmd_ports.conf         06/29 00:19   52 B
      NSM_WR.cfg              06/23 13:53   10 kB
      oem-menu-log.txt        06/29 00:13   3.1 kB
      omDbVersion.info        06/29 00:16   281 B
      post_login.conf         06/29 01:14   40 B
      startup-config          06/29 02:05   7.5 MB
```

Related Commands

copy ftp
copy scp
copy tftp
copy smtp
copy {nfs|cifs}
copy ron
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
show directories

show cores

Purpose A *core file* contains system-memory dumps from a software failure; this information is very useful in diagnosing software problems. All core files reside in the cores directory. Use the show cores command to display this directory or information from a core file.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show cores
show cores file-name { backtrace | syslog | list | message |
cmdinfo}
```

file-name (optional, 1-255 characters) is the name of the core file to view.

backtrace | syslog | list | message | cmdinfo is a required choice.

backtrace displays information from the current program stack. See Figure 7.10 on page 7-51 for a sample backtrace.

syslog displays contextual lines from the syslog file (many lines before and 20 lines after the core-dump event).

list shows what types of information/files you can extract from the core file. You can use one of these items with show cores file-name list-item syntax, described below.

message displays the syslog-file message that announced the core dump. Use the time stamp in this message to focus on the time of the software failure.

cmdinfo displays the CLI output for the specified file/argument.

Guidelines Use the copy command to copy the file off to an FTP server, an SCP host, a TFTP server, an E-mail recipient, or a namespace. The move command can send the file to one of those destinations and delete the original. Use move (with a local destination) to rename it, or delete to remove it.

Samples

```
stoweA# show cores

  cores
      core-0001.dmp           06/29 03:50   332 kB
      core-0002.dmp           06/29 03:51   335 kB
```

shows two core-dump files.

```
bstnA> show cores core-0012.dmp backtrace
```

shows the backtrace in another core-dump file. See Figure 7.10.

Related Commands

copy ftp
copy scp
copy tftp
copy smtp
copy {nfs|cifs}
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
show directories

Figure 7.10 *Sample Output: show cores ... backtrace*

```
bstnA> show cores core-0012.dmp backtrace
Using host libthread_db library "/lib/libthread_db.so.1".
Core was generated by `'.
Program terminated with signal 6, Aborted.
Reading symbols from /acopia/lib/libunicode.so...done.
Loaded symbols for /acopia/lib/libunicode.so
Reading symbols from /usr/lib/i686/libcrypto.so.0.9.6...done.
Loaded symbols for /usr/lib/i686/libcrypto.so.0.9.6
Reading symbols from /acopia/lib/libncftp.so...done.
Loaded symbols for /acopia/lib/libncftp.so
...
Reading symbols from /acopia/lib/msg/msgScm_cli_1.msg_so...done.
Loaded symbols for /acopia/lib/msg/msgScm_cli_1.msg_so
#0  0x40cd3781 in kill () from /lib/libc.so.6
(gdb) #0  0x40cd3781 in kill () from /lib/libc.so.6
#1  0x40221e5e in pthread_kill () from /lib/libpthread.so.0
#2  0x40222339 in raise () from /lib/libpthread.so.0
#3  0x40cd4be1 in abort () from /lib/libc.so.6
#4  0x082d839f in msgAccess::validateId (id=256) at msgAccess.cc:107
#5  0x082d85b3 in msgAccess::getClass (id=256) at msgAccess.cc:194
#6  0x083ab0cb in Status::raiseErrorWork (this=0x91ad8fc, helper=@0x91ca9b0,
    id=256, argList=0xbffff4ec) at Status.cc:93
#7  0x083a7211 in Operation::raiseErrorArgsExitWork (this=0x91ad8f8,
    errCode=256) at Operation.cc:474
#8  0x085b2a07 in CopyFile::respond (this=0x91ad8f8, helper=@0x91ca9b0)
    at CopyFile.pub_op_cc:185
#9  0x083a878c in Operation::respond (this=0x91ad8f8) at Operation.cc:753
#10 0x083a66a3 in Operation::executePhases (this=0x91ad8f8) at Operation.cc:289
#11 0x083a6e30 in Operation::execute (this=0x91ad8f8) at Operation.cc:342
#12 0x08074f15 in CliOpIf::run (this=0xbffff848) at CliOpIf.cc:241
#13 0x080c0203 in cliPrivexecPlainCopy9 (dataPtr=0x9168700, args=0x90a5680,
    variableName=0x0, outputStr=0x90a56d6 "", userCookie=0x90a5ee4)
    at acopiaobj/as1_SCM_dev/cliGlue.cc:9669
#14 0x0808bf7a in RcParseLine (theConnectionPtr=0x9168880) at RcParse.c:601
#15 0x080905ec in RcFiniteStateMachine (theConnectionPtr=0x9168880)
    at RcCmdLin.c:677
#16 0x0808eadd in HandleConnectionTask (theConnectionPtr=0x9168880)
    at AsMain.c:1730
#17 0x0808e977 in AllegroMainTask (theTaskDataPtr=0x9168700,
    theHttpTasks=0xbffffb14, theTcpTasks=0xbffffb18) at AsMain.c:1337
#18 0x0808a8e1 in instantiate_task (scriptname=0x0) at RpTask.c:240
#19 0x0808a79f in main (argc=3, argv=<incomplete type>) at RpTask.c:150
```

show diag-info

Purpose Use this command to display the contents of the diag-info directory.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show diag-info`**

Guidelines If someone issues a collect command with a local file as the target, a large diagnostics file is written to the diag-info directory. Use this command to prove that the file was created. Use the copy command to copy the file off to an FTP server, an SCP host, a TFTP server, an E-mail recipient, or a namespace. The move command can send the file to one of those destinations and delete the original. Use move (with a local destination) to rename it, or delete to remove it.

To conserve disk space, the system only stores one diag-info file at a time.

Sample

```
bstnA> show diag-info

  diag-info
      juser-ns-wwmed.tgz      06/29 08:17  6.3M
```

Related Commands

collect
copy ftp
copy scp
copy tftp
copy smtp
copy {nfs|cifs}
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
show directories

show directories

Purpose The ARX contains directories on its hard disks where you can store release files, configuration scripts, and other switch-maintenance files. Use the show directories command to display all directories or enter show directory-name to display the contents for one directory only.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show directories
show {releases | license | logs | stats-logs | cores | configs
| reports | diag-info | auto-diagnostic | software | scripts }
```

directories shows all directories and their contents

releases | logs | ... | scripts focuses on one directory.

Guidelines The show directories command displays the maintenance files on the ARX's local hard disks. To maintain these directories, use the copy, move, grep, and delete commands.

Samples

`bstnA> show directories`

shows all directories and their contents. See the sample output in Figure 7.11, below.

`bstnA> show releases`

shows the contents of the releases directory.

Related Commands

copy ftp
copy scp
copy tftp
copy {nfs|cifs}
copy smtp
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
grep

Figure 7.11 *Sample Output: show directories*

```
bstnA# show directories

  releases
      test1.rel              06/24 10:23   1.0 GB        V6.01.000.14033
  R A test2.rel              06/29 00:18   1.0 GB        V6.01.000.14032
  B   test3.rel              06/28 21:06   1.0 GB        V6.01.000.14036
```

```
configs
    wwmed.keystore          06/16 10:33   1.3 kB
    active.license          06/29 01:08   9.3 kB
    arx.dossier             04/11 22:54   3.0 kB
    arx.regkey              06/29 01:08   34 B
    boot-config             06/29 01:15   1.5 kB
    boot-config.err         05/25 15:49   1.3 kB
    debug.dossier           06/29 01:08   1.7 kB
    eeprom.dat              06/29 00:17   119 B
    fbstat.out              06/29 02:05   41 kB
    nlmd_ports.conf         06/29 00:19   52 B
    NSM_WR.cfg              06/23 13:53   10 kB
    oem-menu-log.txt        06/29 00:13   3.1 kB
    omDbVersion.info        06/29 00:16   281 B
    post_login.conf         06/29 01:14   40 B
    startup-config          06/29 02:05   7.5 MB

scripts
    analyze-failover        06/20 20:00   95 B
    change-mfg-date.scr     06/20 20:00   631 B
    check-global.scr        06/20 20:00   1.9 kB
    check-hw.scr            06/20 20:00   1.8 kB
    check-run.scr           06/20 20:00   3.0 kB
    clean_reports           06/20 20:00   1.3 kB
    cli_script_testdir.scr  11/16/2010     0 B
    cli_script.scr          11/16/2010     0 B
    global                  06/29 02:05   24 kB
    import_rate             06/20 20:00   16 kB
    jiltdump                06/20 20:00   2.1 kB
    logging.scr             07/10/2010    218 B
    monitor                 06/20 20:00   1.3 kB
    nslookup                06/20 20:00   773 B
    power                   06/20 20:00   1.6 kB
    run_cfg.scr             07/08/2010    10 kB
    running                 06/29 02:05   12 kB
    schemadump.sql          06/08 14:34   696 B
    share_status            06/20 20:00   4.1 kB
    show                    06/20 20:00   660 B
    show_chassis.scr        05/18 10:37   70 B
    start_conf              06/29 02:05   36 kB
    test-global-config_orig.scr 06/15 17:27   24 kB
    test-global-config.scr  06/15 17:38   24 kB
    test-running-config_orig.scr 06/15 17:27   10 kB
    test-running-config.scr 06/15 17:38   10 kB
    traceroute              06/20 20:00   922 B
    ttcp                    06/20 20:00   2.1 kB
    unh                     06/20 20:00   7.1 kB

logs
    error.log               06/29 02:03   185 kB
    fastpath                06/29 01:55   33 kB
    ha-reboot-history.log   06/29 00:13   322 B
    syslog                  06/29 02:05   5.2 MB
    traplog                 06/29 02:04   594 kB

cores

reports
  Codes: AbCh=Access-based Enum Changes,
         Act=DR activate configuration CLI output,
         AdUp=Active Directory Forest Update, ArC=Show Archive Contents,
         At=Command Scheduler, CLI=CLI Log, Diag=Collect Diag-Info,
```

```
          Dstg=Destage, ExMp=Export Mapping, FDR=File and Directory Tracking,
          Fs=Fileset, Imp=Import, Inc=Inconsistencies, iPl=Inline Place Rule,
          Load=DR load configuration CLI output, MdO=Metadata Only,
          MdU=Metadata Upgrade, Mem=Metalog Usage Statistics,
          Mem=Memory Usage Statistics, MgMd=Migrate Metadata, NIS=NIS Update,
          Plc=Place Rule, Proc=Processor Usage Statistics,
          PrSu=Promote Subshares, Rbld=Rebuild, RDbg=Rule Debug,
          Repl=DR Replicate configuration, Rm=Remove, RmNs=Remove Namespace,
          RmSh=Remove Share, RsD=Restore Data, RsSu=Remove Storage Subshares,
          RSSySu=Replica Snapshot Sync Subshares, SCp=Shadow Copy,
          Snapshot=Snapshot, SuCa=Subshare cache contents., Sum=Summary,
          SuNS=Subshare Sync New Storage, SuSv=Subshare Sync from Service,
          SuVo=Subshare Sync from Volume, SymL=Symlinks, Sync=Sync Files/Dirs,
          VPH=Virtual Path History
      active-directory-MEDARCH.ORG.rpt            06/29 01:14   13 kB        AdUp DONE: 18 in
00:00:00
      active-directory-wells.me.org.rpt 06/29 01:15   3.4 kB        AdUp DONE: 0 in 00:00:00
      active-directory-vt.com.rpt 06/29 01:14   6.6 kB        AdUp DONE: 9 in 00:00:01
      adminSessions_201106290129.rpt 06/29 01:29   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201106290134.rpt 06/29 01:34   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201106290139.rpt 06/29 01:39   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201106290144.rpt 06/29 01:44   1.3 kB        At   DONE: 10 in 00:00:00
...
      syncSshrVolToService_20110629012633.rpt 06/29 01:26   2.4 kB        SuVo DONE: 9 in 00:00:00
      syncSshrVolToService_20110629012636.rpt 06/29 01:26   2.3 kB        SuVo DONE: 9 in 00:00:00
      wwmed_bills2.rpt        06/29 02:02   1.8 kB        Inc  DONE: 4432 in 00:00:03
      wwmed_meta.rpt          06/29 02:02   722 kB        MdO  DONE: 4432 in 00:00:01

  capture
      cifsVol.cap             06/29 01:55   220 kB
      clientCap.cap           06/29 01:45   15 MB
      fsrvr.cap               06/29 02:05   52 kB
      nasTraffic.cap          06/29 01:55   212 kB
      ntaps_00001_20110629060427.cap 06/29 02:04   24 B
      proxyTraffic_01620_20110629055529.cap 06/29 01:55   150 kB
      proxyTraffic_01621_20110629055531.cap 06/29 01:55   94 kB

  diag-info

  auto-diag

  software
      A-VE_install.pdf        06/20 23:31   2.2 MB
      A-VE_quickstart.pdf     06/20 23:31   94 kB
      A1500_install.pdf       06/20 23:31   2.1 MB
      A1500_quickstart.pdf    06/20 23:31   491 kB
      A1k_install.pdf         06/20 23:31   1.3 MB
      A1k_quickstart.pdf      06/20 23:31   190 kB
      A2500_install.pdf       06/20 23:31   2.0 MB
      A2500_quickstart.pdf    06/20 23:30   301 kB
      A2k_install.pdf         06/20 23:31   7.8 MB
      A2k_quickstart.pdf      06/20 23:31   684 kB
      A4k_install.pdf         06/20 23:30   17 MB
      A4k_quickstart.pdf      06/20 23:31   1.0 MB
      A5c_install.pdf         06/20 23:31   2.3 MB
      A5c_quickstart.pdf      06/20 23:30   133 kB
      A6k_install.pdf         06/20 23:31   2.6 MB
      A6k_quickstart.pdf      06/20 23:31   808 kB
      HD_FRU.pdf              06/20 23:31   931 kB
      HW_Reference.pdf        06/20 23:31   6.1 MB
      Mibs.tgz                06/21 00:26   168 kB
      PwrSupply_FRU.pdf       06/20 23:30   953 kB
```

```
    acopiasmi.my            06/20 20:22   219 kB
    bridge.my               02/25/2003    45 kB
    cliMaintenance.pdf      06/20 23:31   3.8 MB
    cliNetwork.pdf          06/20 23:31   3.2 MB
    cliReference.pdf        06/20 23:31   13 MB
    cliStorage.pdf          06/20 23:31   5.7 MB
    compatibilityMatrix.pdf 06/20 23:31   206 kB
    dot3ad.my               02/25/2003    39 kB
    entity.my               06/28/2004    51 kB
    etherlike.my            10/12/2004    87 kB
    glossary.pdf            06/20 23:31   914 kB
    ifmib.my                10/12/2004    70 kB
    ifType.my               10/07/2004    4.4 kB
    logCatalog.pdf          06/20 23:31   2.9 MB
    masterIndex.pdf         06/20 23:30   2.5 MB
    mib-2.my                10/07/2004    100 kB
    openview.trapd.conf     06/20 20:00   218 kB
    pbridge.my              02/25/2003    31 kB
    SecureAgent.pdf         06/20 23:31   1.2 MB
    SlideRail_FRU.pdf       06/20 23:31   1.2 MB
    releaseNotes.html       06/20 23:31   295 kB
    rfc2668.my              10/12/2004    101 kB
    rmon.my                 02/25/2003    147 kB
    sitePlanning.pdf        06/20 23:31   2.0 MB
    snap-recon.pl           06/20 20:00   18 kB
    snmpReference.pdf       06/20 23:30   1.2 MB
    snmpv2-mib.my           10/07/2004    36 kB
    snmpv2-smi.my           10/07/2004    1.7 kB
    stamp-mibs-tgz          06/21 00:26    0 B
    vlan.my                 02/25/2003    69 kB

stats-logs
    cifs-service_20110629_041950.raw.stats.csv        06/29 01:28   2.9 kB
    cifs-service-auth_20110629_041950.raw.stats.csv   06/29 01:28   1.5 kB
    cifs-share_20110629_041950.raw.stats.csv          06/29 01:28   14 kB
    cifs-work-queue_20110629_041950.raw.stats.csv     06/29 01:28   4.0 kB
    domain-controller_20110629_041950.raw.stats.csv   06/29 01:28   13 kB
    metadata_20110629_041950.raw.stats.csv            06/29 01:28   791 B
    metalog_20110629_041950.raw.stats.csv             06/29 01:28   683 B
    migration_20110629_041950.raw.stats.csv           06/29 01:28   585 B
    nfs-service_20110629_041950.raw.stats.csv         06/29 01:28   1.4 kB
    nfs-share_20110629_041950.raw.stats.csv           06/29 01:28   8.1 kB
    nfs-share_20110629_060000.hourly.stats.csv        06/29 02:03   1.1 kB
```

show license

Purpose Use this command to display the contents of the license directory and/or the contents of a specified license file. A license file specifies the features and capacity that are licensed for use on this ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show license [file-name]`

file-name (optional, 1-255 characters) is the name of the file to display. If you omit this, the command displays the list of the directory's files.

Guidelines Issue the command to view the directory contents. Then issue the command with a file name to view the contents of the given license file. The documentation for show active-license describes the detailed output for a license file.

If the directory listing does not include an "active.license" file, no license is active for the current ARX. You can use the license activate command to automatically activate a license, but only if the ARX is connected to the Internet. For an ARX that is off of the Internet, you can manually activate the license. Start with the license create license-dossier command for manual activation.

Samples

`stkbrgA# show license`

shows all license files on the switch. For sample output, see Figure 7.12 on page 7-58.

`stkbrgA# show license active.license`

shows the currently-active license file. For sample output, see Figure 7.13 on page 7-58.

Related Commands

show active-license
license activate
license create license-dossier
copy ftp
copy scp
copy tftp
copy {nfs|cifs}
copy smtp
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
show directories

Figure 7.12 *Sample Output: show license*

```
stkbrgA# show license


  license
      active.license          06/29 01:08   9.3 kB
```

Figure 7.13 *Sample Output: show license active.license*

```
stkbrgA# show license active.license
License Information in active.license file.
-------------------------------------------

Auth Vers:                      5b
Usage:                          F5 Internal Product Development
Registration Key:               CRJGV-QPDYW-SATNK-RGBYY-DMTMOBL
Licensed version:               6.0.0
License Date:                   12/21/2010
License Start:                  12/20/2010
License End:                    07/29/2011
Service Check Date:             06/29/2011
Platform ID:                    Z100
Service Status:                 As of 2011-06-29 there is no active service contract.
 :                               This may inhibit your ability to upgrade your software.


Module List
-----------------
ARX VE:
Reg Key:                        Y837955-1236781


Feature List
-----------------
cifs_services_per_system:                8
cpu_cores:                               2
direct_attach_points_per_system:         16384
direct_shares_per_system:                2048
direct_shares_per_volume:                127
direct_shares_per_volume_group:          2048
disk_space_gb:                           40
files_per_system_4k:                     46875
files_per_volume_4k:                     15625
files_per_volume_group_4k:               46875
global_servers_per_system:               8
memory:                                  4096
namespaces_per_system:                   2
nfs_services_per_system:                 8
nic_interface_count:                     1
protocol_qty_allowed:                    2
shares_per_system:                       64
shares_per_volume:                       16
shares_per_volume_group:                 32
virtual_services_per_system:             8
volume_groups_per_system:                2
volumes_per_system:                      32
volumes_per_volume_group:                16
```

show logs

Purpose Use this command to display the contents of the logs directory and/or the contents of a specified log file.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show logs [`*`file-name`*`]`**

file-name (optional, 1-255 characters) is the name of the file to display. If you omit this, the command displays the list of the directory's files.

Guidelines Issue the command to view the directory contents. Then issue the command with a file name to view the file contents.

The syslog file is arguably the most-important file in the logs directory. It shows detailed logs from most of the switch software. You can use the logging level command to change the volume of messages from each software subsystem. Namespaces, volumes, and shares are denoted in the syslog by their internal, numeric IDs: use the show id-mappings command to translate these IDs into names used in the CLI.

Sample

```
bstnA# show logs


  logs
      error.log                06/29 02:03   185 kB
      fastpath                 06/29 01:55   33 kB
      ha-reboot-history.log    06/29 00:13   322 B
      syslog                   06/29 02:05   5.2 MB
      traplog                  06/29 02:04   594 kB
```

Related Commands

copy ftp
copy scp
copy tftp
copy {nfs|cifs}
copy smtp
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
show directories

show releases

Purpose A software release for the ARX is packaged in a *release file*. All release files are stored in the "releases" directory. Use this command to display the contents of the releases directory and/or the version of a specified release file.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show releases [file-name] [verbose]`

file-name (optional, 1-255 characters) is the name of the file to display. If you omit this, the command displays the list of the directory's files.

verbose (optional) runs a checksum test on the integrity of the downloaded release file.

Guidelines Issue the command without any options for a list of release files on the ARX. Then issue the command with a file name to view the specific release file.

The directory output shows one row for each release file. The table has five columns:

Flags appear in the first column to indicate how the release file is used:

- B is a flag for the Backup release, if any. This is the software release that was running prior to the current release. F5 personnel can roll the switch back to this release if needed.
- R is the flag for the currently-Running release.
- A flags the Armed release. This is the software to be loaded on the next reboot. Use boot system to arm the switch with a new release file. Use reload to reboot and put the new release file into service.

The release-file name is in the next column.

The date and time, in the third column, is the time the file was copied to the switch.

The size of the file is in the fourth column.

The next column, usually blank, says "incomplete" between the time that the file is downloaded and the switch has performed an integrity check. The switch does not use the release file until this check is complete. It should complete at the end of the download, automatically; a poorly-timed switch reload may prevent the check. In this case, replace the file.

The version of the release appears in the final column. The version appears in V*release*.*build* format, where *build* is a specific build in the given *release*. For example, "V5.02.000.12574" would appear for build 12574 of release 5.02.000.

If you issue the command with a file name, the output shows the version number, the date the release was built, and the username that ran the build. If you use the verbose option, the results of a checksum test appear. If this shows that the checksum failed, retry the copy of the release file.

Samples

```
provA# show releases
```

shows all releases on the switch. For sample output, see Figure 7.14, below.

```
provA# show releases test1.rel
Version 5.02.000.12568 (May 14 2010 20:19:21) [nbuilds]
```

shows one release file.

```
provA# show releases test1.rel verbose
Version 5.02.000.12568 (May 14 2010 20:19:21) [nbuilds]
Checksum: Passed
```

show that the release file's checksum passed.

Related Commands

copy ftp
copy scp
copy tftp
copy {nfs|cifs}
copy smtp
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
show directories

Figure 7.14 *Sample Output: show releases*

```
provA# show releases

  releases
  R A test1.rel              06/29 00:16   1.0 GB        V6.01.000.14032
  B   test2.rel              06/07 09:04   1.0 GB        V5.02.000.12577
      test3.rel              06/27 03:05   1.0 GB        V6.01.000.14030
```

show replicated-configs

Purpose Use this command to display the replicated-configs directory, which contains configuration files from a remote ARX. You can use these files to recreate the remote configuration in case of a disaster at the remote site.

You can use this command to list all of the replicated-configuration files on the current ARX, or to read one configuration file from the directory.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show replicated-configs [`*`file-name`*`]`**

file-name (optional, 1-255 characters) is the name of the file to display. If you omit this, the CLI shows a list of the directory's files.

Guidelines Use the show directories command to show the contents of all directories on the ARX. To maintain this directory, use the copy, move, grep, and delete commands.

This holds global-configuration files from remote ARX clusters. You can create a config-replication rule to regularly copy a global-config file from one cluster to another. The file's contents are an ordered list of CLI commands required to recreate the remote cluster's global configuration; they are the same as the output for show global-config.

Samples `newptA# show replicated-configs`

```
  configs-replication
      provSvcs.rcfg           04/02 03:08   12 kB
      testrun2.rcfg           04/02 03:08   12 kB
```

lists all replicated-config files on the "newptA" chassis.

`newptA# show replicated-configs provSvcs.rcfg`

shows the above file's contents. For sample output, see Figure 7.15 on page 7-63.

Related Commands

config-replication
show global-config
copy ftp
copy scp
copy tftp
copy smtp
copy {nfs|cifs}
copy ron
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
show directories

Figure 7.15 *Sample Output: show replicated-configs provSvcs.rcfg*

```
newptA# show replicated-configs provSvcs.rcfg
; ARX-500
; Version 6.01.000.14032 (Jun 20 2011 20:10:38) [nbuilds]
; Database version: 601000.103
; Generated global-config Wed Jun 29 03:26:19 2011
;
;================================ global ================================
global
  cluster-name providence member provA
  cluster-name newport member newptA

;================================ user ================================
  user adm1 encrypted-password hsLHMTN845U7MbyBvNyrJ8b4FaQzZtZp4MmunTMa+U8=
    exit

  user adm12 encrypted-password yfCVCDN845U7MbyBvNyrJ8b4FaQzZtZp0o1pIGWXwR/tMfOyEU1ALw==
    exit

  user admin encrypted-password E3uUIzN845U7MbyBvNyrJ8b4FaQzZtZpFkltcs+pEc0=
    exit

  user newadmin encrypted-password rzUtjzN845U7MbyBvNyrJ8b4FaQzZtZpA5GX80tLNo7ZpnkYzBhMOg==
    exit

;================================ group ================================
  group Administrators
    role backup-operator
    role crypto-officer
...
```

show reports

Purpose Use this command to display the contents of the reports directory and/or the contents of a specified file in the directory.

Mode (any)

Security Role(s) operator or backup operator (any)

Syntax

```
show reports
show reports report-name
show reports type report-type
show reports type report-type report-name
```

If you omit both options, the command displays the full list of the directory's files.

report-name (optional, 1-255 characters) is the name of the report to display.

report-type (optional, 1-12 characters) specifies the type of report to list. Enter **show reports type ?** for a full list of report types.

Guidelines Issue the command (perhaps with a type specification) to view the directory contents. Then issue the command with a file name to view the file contents.

When you use the command to list reports, the output is a table with one line per report. Each line shows a summary status for the job that generated the report. To see the one-line status for a particular report, use show reports status.

Some reports run over a long period of time. To wait for a report to finish (perhaps in a CLI script), you can use wait-for report.

Samples

```
bstnA# show reports
```

lists all reports. See Figure 7.16 on page 7-65 for sample output.

```
bstnA> show reports inconsistencies.12.rpt
```

shows one report. See Figure 7.17 on page 7-68.

```
bstnA> show reports type Imp
```

shows all import reports. See Figure 7.18 on page 7-69.

Related Commands

copy ftp
copy scp
copy tftp
copy {nfs|cifs}
copy smtp
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
show directories
show reports status
wait-for report

Figure 7.16 *Sample Output: show reports*

```
bstnA# show reports


  reports
    Codes: AbCh=Access-based Enum Changes,
           Act=DR activate configuration CLI output,
           AdUp=Active Directory Forest Update, ArC=Show Archive Contents,
           At=Command Scheduler, CLI=CLI Log, Diag=Collect Diag-Info,
           DS=Directory Structure, Dstg=Destage, ExMp=Export Mapping,
           FDR=File and Directory Tracking, Fs=Fileset, Imp=Import,
           Inc=Inconsistencies, iPl=Inline Place Rule,
           Load=DR load configuration CLI output, MdO=Metadata Only,
           MdU=Metadata Upgrade, Mem=Metalog Usage Statistics,
           Mem=Memory Usage Statistics, MgMd=Migrate Metadata,
           MgVg=Migrate Volume Group, MM=Manual Migrate Rule, NIS=NIS Update,
           Plc=Place Rule, Pmd=Policy Memory Debug,
           Proc=Processor Usage Statistics, PrSu=Promote Subshares,
           Rbld=Rebuild, RDbg=Rule Debug, Repl=DR Replicate configuration,
           Rm=Remove, RmNs=Remove Namespace, RmSh=Remove Share,
           RsD=Restore Data, RsSu=Remove Storage Subshares,
           RSSySu=Replica Snapshot Sync Subshares, SCp=Shadow Copy,
           Snapshot=Snapshot, SuCa=Subshare cache contents., Sum=Summary,
           SuNS=Subshare Sync New Storage, SuSv=Subshare Sync from Service,
           SuVo=Subshare Sync from Volume, SymL=Symlinks, Sync=Sync Files/Dirs,
           VPH=Virtual Path History
     aAdamsDatHist_ac1.MEDARCH.ORG_ARCHIVES_201201190637.rpt           01/19 01:37   4.0 kB
FDR  DONE: 4 in 00:00:00
      active-directory-MEDARCH.ORG.rpt            01/19 00:33   13 kB        AdUp DONE: 18 in
00:00:00
      active-directory-wells.me.org.rpt 01/19 00:34   3.4 kB        AdUp DONE: 0 in 00:00:05
      active-directory-vt.com.rpt 01/19 00:33   6.6 kB        AdUp DONE: 9 in 00:00:00
      adminSessions_201201190048.rpt 01/19 00:48   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201201190053.rpt 01/19 00:53   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201201190058.rpt 01/19 00:58   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201201190103.rpt 01/19 01:03   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201201190108.rpt 01/19 01:08   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201201190113.rpt 01/19 01:13   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201201190118.rpt 01/19 01:18   1.3 kB        At   DONE: 10 in 00:00:01
      adminSessions_201201190123.rpt 01/19 01:23   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201201190128.rpt 01/19 01:28   1.3 kB        At   DONE: 10 in 00:00:00
      adminSessions_201201190133.rpt 01/19 01:33   1.3 kB        At   DONE: 10 in 00:00:00
```

```
      adminSessions_201201190138.rpt 01/19 01:38   1.3 kB        At   DONE: 10 in 00:00:00
      bePaths_all_201201190637.rpt 01/19 01:37   4.5 kB        VPH  DONE: 11 in 00:00:00
      bePathsVerbose_all_201201190637.rpt 01/19 01:37   4.7 kB        VPH  DONE: 11 in 00:00:00
      bePathsVerbose_fileRecordsMed_201201190637.rpt 01/19 01:37   4.8 kB        VPH  DONE: 11 in
00:00:00
      FA_lab_0_create_20120119005203858.rpt 01/19 00:52   2.4 kB        Snapshot DONE: 0 in
00:00:00
      FA_rcrds_0_create_20120119005703542.rpt 01/19 00:57   2.4 kB        Snapshot DONE: 0 in
00:00:00
      bePaths_fileRecordsMed_201201190637.rpt 01/19 01:37   4.5 kB        VPH  DONE: 11 in 00:00:00
      cifs_only.rpt           01/19 01:30   5.0 kB        ExMp DONE: 10 in 00:00:00
      cifsPromoteSubshares_201201190040.rpt 01/19 00:40   787 B        PrSu DONE: 1 in 00:00:00
      cifsUsers_201201190050.rpt 01/19 00:50   1.0 kB        At   DONE: 10 in 00:00:00
      collect_diag_201201190536.rpt 01/19 00:36   1.5 kB        Diag DONE: 0 in 00:00:13
      daily_archive_201201190038.rpt 01/19 00:39   2.3 kB        Plc  DONE: 0 in 00:00:10
      daily_archive_201201190136.rpt 01/19 01:36   2.3 kB        Plc  DONE: 0 in 00:00:11
      daily_archive_201201190140.rpt 01/19 01:40   2.3 kB        Plc  DONE: 203 in 00:00:41
      dir_structure.26.rpt    01/19 01:29   8.0 kB        DS   DONE: 223 in 00:00:00
      dir_structure.27.rpt    01/19 01:29   2.8 kB        DS   DONE: 63 in 00:00:00
      dir_structure.28.rpt    01/19 01:29   2.5 kB        DS   DONE: 3 in 00:00:00
      dir_structure.29.rpt    01/19 01:29   1.8 kB        DS   DONE: 32 in 00:00:00
      dir_structure.30.rpt    01/19 01:29   2.1 kB        DS   DONE: 20 in 00:00:00
      dir_structure.31.rpt    01/19 01:29   2.0 kB        DS   DONE: 3 in 00:00:00
      dir_structure.33.rpt    01/19 01:29   8.0 kB        DS   DONE: 223 in 00:00:00
      dir_structure.34.rpt    01/19 01:29   1.7 kB        DS   DONE: 223 in 00:00:00
      dir_structure.37.rpt    01/19 01:30   1.9 kB        DS   FAILED: 0 in 00:00:00
      dir_structure.7.rpt     01/19 01:28   93 kB        DS   DONE: 4432 in 00:00:00
      dir_structure.8.rpt     01/19 01:28   2.0 kB        DS   DONE: 109 in 00:00:00
      dir_structure.9.rpt     01/19 01:28   2.0 kB        DS   DONE: 19 in 00:00:00
      docsPlc_201201190037.rpt 01/19 00:37   1.4 kB        iPl  RUNNING: 0 in 01:04:26
      docsPlc_20120119003752.rpt 01/19 00:37   2.1 kB        Plc  FAILED: 0 in 00:00:01
      export-mapping.rpt      01/19 01:38   6.1 kB        ExMp DONE: 13 in 00:00:01
      fe2beMap.rpt            01/19 01:38   6.1 kB        ExMp DONE: 13 in 00:00:00
      fiveDayFiles_fileRecordsMed_201201190637.rpt 01/19 01:37   1.2 kB        ArC  DONE: 2 in
00:00:00
      flRcrdsSinceJan_fileRecordsMed_201201190637.rpt 01/19 01:37   1.2 kB        ArC  DONE: 2 in
00:00:00
      fs2_sbshrs_201201190636.rpt 01/19 01:36   28 kB        SuCa DONE: 79 in 00:00:00
      ft-config_201201190047.rpt 01/19 00:47   3.0 kB        CLI  DONE: 40 in 00:00:02
      idxPaths_ac1.MEDARCH.ORG_labs_201201190637.rpt            01/19 01:37   3.1 kB        FDR
DONE: 3 in 00:00:00
      idxPathsVerbose_ac1.MEDARCH.ORG_labs_201201190637.rpt            01/19 01:37   3.2 kB
FDR  DONE: 3 in 00:00:00
      import.1.budget.5.rpt   01/19 00:37   2.0 kB        Imp  DONE: 713 in 00:00:10
      import.10.backlots.17.rpt 01/19 00:38   2.3 kB        Imp  DONE: 1 in 00:00:04
      import.11.scanners.18.rpt 01/19 00:38   2.4 kB        Imp  DONE: 5 in 00:00:03
      import.2.bills.6.rpt    01/19 00:37   2.2 kB        Imp  DONE: 122 in 00:00:08
      import.3.bills2.7.rpt   01/19 00:37   2.6 kB        Imp  DONE: 415 in 00:00:08
      import.4.it5.8.rpt      01/19 00:37   2.1 kB        Imp  DONE: 131 in 00:00:08
      import.5.rx.11.rpt      01/19 00:38   2.2 kB        Imp  DONE: 34 in 00:00:03
      import.6.charts.12.rpt  01/19 00:38   2.5 kB        Imp  DONE: 165 in 00:00:04
      import.7.bulk.13.rpt    01/19 00:38   2.2 kB        Imp  DONE: 5 in 00:00:04
      import.8.equip.15.rpt   01/19 00:38   2.1 kB        Imp  DONE: 20 in 00:00:04
      import.9.leased.16.rpt  01/19 00:38   2.1 kB        Imp  DONE: 43 in 00:00:03
      inconsistencies.1.rpt   01/19 01:27   1.8 kB        Inc  DONE: 4432 in 00:00:04
      inconsistencies.14.rpt  01/19 01:29   2.0 kB        Inc  DONE: 223 in 00:00:00
      inconsistencies.15.rpt  01/19 01:29   2.1 kB        Inc  DONE: 63 in 00:00:00
      inconsistencies.16.rpt  01/19 01:29   2.1 kB        Inc  DONE: 3 in 00:00:00
      inconsistencies.17.rpt  01/19 01:29   1.9 kB        Inc  DONE: 32 in 00:00:00
      inconsistencies.18.rpt  01/19 01:29   2.0 kB        Inc  DONE: 20 in 00:00:00
      inconsistencies.19.rpt  01/19 01:29   2.1 kB        Inc  DONE: 3 in 00:00:00
      inconsistencies.2.rpt   01/19 01:27   1.8 kB        Inc  DONE: 109 in 00:00:00
```

```
      inconsistencies.3.rpt   01/19 01:27   1.8 kB        Inc  DONE: 19 in 00:00:00
      indecesSinceNov_ac1.MEDARCH.ORG_labs_201201190637.rpt           01/19 01:37   3.2 kB
FDR  DONE: 3 in 00:00:00
     indecesSinceNov_ac1.MEDARCH.ORG_labs_20120119063755.rpt          01/19 01:37   3.2 kB
FDR  DONE: 3 in 00:00:01
      indexPaths_ac1.MEDARCH.ORG_labs_201201190637.rpt           01/19 01:37   3.1 kB        FDR
DONE: 3 in 00:00:00
     indexPaths_ac1.MEDARCH.ORG_labs_20120119063748.rpt          01/19 01:37   3.2 kB        FDR
DONE: 3 in 00:00:00
      labArchive_0_remove_20120119055205741.rpt 01/19 00:52   2.3 kB        Snapshot DONE: 0 in
00:00:01
      leTier1_201201190039.rpt 01/19 01:39   2.1 kB        iPl  DONE: 0 in 01:00:00
      leTier1_201201190139.rpt 01/19 01:39   1.4 kB        iPl  RUNNING: 0 in 00:03:07
      medarcv_meta.rpt        01/19 01:29   2.1 kB        Inc  DONE: 63 in 00:00:00
      metadata_only.12.rpt    01/19 01:28   9.8 kB        MdO  DONE: 65 in 00:00:00
      metadata_only.20.rpt    01/19 01:29   31 kB        MdO  DONE: 223 in 00:00:00
      metadata_only.21.rpt    01/19 01:29   9.8 kB        MdO  DONE: 63 in 00:00:00
      metadata_only.22.rpt    01/19 01:29   2.0 kB        MdO  DONE: 3 in 00:00:00
      metadata_only.23.rpt    01/19 01:29   5.4 kB        MdO  DONE: 32 in 00:00:00
      metadata_only.24.rpt    01/19 01:29   4.2 kB        MdO  DONE: 20 in 00:00:00
      metadata_only.25.rpt    01/19 01:29   2.0 kB        MdO  DONE: 3 in 00:00:00
      metadata_only.36.rpt    01/19 01:30   1.7 kB        MdO  FAILED: 0 in 00:00:00
      metadata_only.4.rpt     01/19 01:28   722 kB        MdO  DONE: 4432 in 00:00:00
      metadata_only.5.rpt     01/19 01:28   15 kB        MdO  DONE: 109 in 00:00:00
      metadata_only.6.rpt     01/19 01:28   4.1 kB        MdO  DONE: 19 in 00:00:00
      nfs_only.rpt            01/19 01:38   3.1 kB        ExMp DONE: 3 in 00:00:01
      nis-update.wwmed.com.rpt           01/19 01:42   2.0 MB        NIS  RESOLVING: 19063 in
01:05:59
      pathsSinceJan_all_201201190637.rpt 01/19 01:37   4.5 kB        VPH  DONE: 11 in 00:00:00
      pathsSinceJan_all_20120119063733.rpt 01/19 01:37   4.7 kB        VPH  DONE: 11 in 00:00:00
      pathsSinceJan_fileRecordsMed_201201190637.rpt 01/19 01:37   4.5 kB        VPH  DONE: 11 in
00:00:00
      pathsSinceJan_fileRecordsMed_20120119063737.rpt 01/19 01:37   4.8 kB        VPH  DONE: 11 in
00:00:00
      pathsTilJune_all_201201190637.rpt 01/19 01:37   589 B        VPH  DONE: 0 in 00:00:00
      pathsTilJune_fileRecordsMed_201201190637.rpt 01/19 01:37   600 B        VPH  DONE: 0 in
00:00:00
      processor_usage_201201190131.rpt 01/19 01:31   915 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.0.rpt 01/19 01:31   918 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.1.rpt 01/19 01:31   919 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.10.rpt 01/19 01:31   919 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.2.rpt 01/19 01:31   918 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.3.rpt 01/19 01:31   918 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.4.rpt 01/19 01:31   918 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.5.rpt 01/19 01:31   918 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.6.rpt 01/19 01:31   918 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.7.rpt 01/19 01:31   918 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.8.rpt 01/19 01:31   919 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.9.rpt 01/19 01:31   919 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013144.rpt 01/19 01:31   918 B        Proc DONE: 1 in 00:00:00
      processor_usage_20120119013145.rpt 01/19 01:31   6.2 kB        Proc DONE: 13 in 00:00:00
      rcrds-dirs.rpt          01/19 01:30   8.0 kB        DS   DONE: 223 in 00:00:00
      rcrdsArchive_0_remove_20120119055705306.rpt 01/19 00:57   2.3 kB        Snapshot DONE: 0 in
00:00:00
      rcrdsIssues.rpt         01/19 01:29   2.0 kB        Inc  DONE: 223 in 00:00:01
     reagentListsSinceNov_ac1.MEDARCH.ORG_labs_201201190637.rpt          01/19 01:37   762 B
FDR  DONE: 0 in 00:00:00
     reagentListsSinceNov_ac1.MEDARCH.ORG_labs_20120119063756.rpt         01/19 01:37   773 B
FDR  DONE: 0 in 00:00:00
      sbshrs_201201190636.rpt 01/19 01:36   138 kB        SuCa DONE: 397 in 00:00:00
      sendback_201201190139.rpt 01/19 01:39   2.3 kB        Plc  DONE: 163 in 00:00:12
```

```
      snap_daily_0_create_20120119004024903.rpt 01/19 00:40   3.1 kB        Snapshot DONE: 0 in
00:00:00
      snap_daily_0_create_20120119004100329.rpt 01/19 00:41   3.1 kB        Snapshot DONE: 0 in
00:00:00
      snap_daily_0_create_20120119004130708.rpt 01/19 00:41   3.1 kB        Snapshot DONE: 0 in
00:00:01
      snap_daily_2_remove_20120119004218241.rpt 01/19 00:42   3.1 kB        Snapshot DONE: 0 in
00:00:01
      snap_hourly_0_create_20120119004258666.rpt 01/19 00:43   3.1 kB        Snapshot DONE: 0 in
00:00:00
      snap_hourly_0_create_20120119004329008.rpt 01/19 00:43   3.1 kB        Snapshot DONE: 0 in
00:00:01
      snap_hourly_0_create_20120119013600892.rpt 01/19 01:36   3.4 kB        Snapshot DONE: 0 in
00:00:00
     snap_offsite_weekly_0_create_20120119004631404.rpt 01/19 00:46   3.3 kB       Snapshot DONE:
0 in 00:00:01
     snap_offsite_weekly_0_create_20120119004718925.rpt 01/19 00:47   3.3 kB       Snapshot DONE:
0 in 00:00:00
      sync.1._acct.rpt        01/19 01:28   1.4 kB        Sync DONE: 32 in 00:00:01
      sync.2._rcrds.rpt       01/19 01:30   1.8 kB        Sync DONE: 11 in 00:00:00
      sync.3._rcrds.rpt       01/19 01:30   2.0 kB        Sync DONE: 293 in 00:00:00
      sync.4._lab_equipment.rpt 01/19 01:30   1.5 kB        Sync DONE: 0 in 00:00:00
      sync.5._lab_equipment.rpt 01/19 01:30   2.2 kB        Sync DONE: 0 in 00:00:00
      syncSshrNewStorageReport_201201190538.rpt 01/19 00:38   3.1 kB        SuNS DONE: 10 in
00:00:00
      syncSshrVolToService_201201190040.rpt 01/19 00:40   1.4 kB        SuVo DONE: 5 in 00:00:00
      syncSshrVolToService_20120119004007.rpt 01/19 00:40   1.4 kB        SuVo DONE: 5 in 00:00:00
      syncSshrVolToService_20120119004010.rpt 01/19 00:40   1.4 kB        SuVo DONE: 5 in 00:00:00
      todayFiles_fileRecordsMed_201201190637.rpt 01/19 01:37   1.2 kB        ArC  DONE: 2 in
00:00:00
      tools_dir._claims.rpt   01/19 01:30   1.9 kB        DS   FAILED: 0 in 00:00:00
      wwmed_bills2.rpt        01/19 01:28   1.8 kB        Inc  DONE: 4432 in 00:00:04
      wwmed_meta.rpt          01/19 01:28   722 kB        MdO  DONE: 4432 in 00:00:00
```

Figure 7.17 *Sample Output: show reports inconsistencies.12.rpt*

```
bstnA> show reports inconsistencies.14.rpt
**** Inconsistencies Report: Started at 01/19/2012 01:29:11 -0500 ****
**** Software Version: 6.02.000.14314 (Jan 16 2012 20:04:23) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:
**** Namespace: medarcv
**** Volume: /rcrds
**** Path: /rcrds

Share                Physical Filer
-------------------  ------------------------------------------------------------
[rx               ]  192.168.25.29:prescriptions
[charts           ]  192.168.25.20:histories
[bulk             ]  192.168.25.27:bulkstorage

**** Legend:
****   LF = File exists in the metadata, but is missing from the physical filer.
****   LD = Directory exists in the metadata, but is missing from the physical filer.
****   FF = File exists on the physical filer, but is missing from the metadata.
****   FD = Directory exists on the physical filer, but is missing from the metadata.
****   LL = File is a symlink in the metadata, but is a regular file on the filer.
****   FL = File is a symlink on the filer, but is a regular file in the metadata.
****   IF = Filehandles in the metadata do not match the filehandles on the physical filer.
****   MF = The file is currently being migrated.
****   NL = Unable to lock parent directory during report.
****   FE = Error contacting filer during report.
```

```
****   FO = Filer Offline: The filer is offline or disabled.
****   F8 = A file name matches a CIFS alternate "8.3" name on another share.
****   D8 = A directory name matches a CIFS alternate "8.3" name; its contents will be skipped.
****   DC = A client has the file or directory open for delete-on-close, but the filer has already
deleted it.
****   SD = Striped leaf directory found on filer, expected on other shares.
****   SL = File is a symlink.
****   UT = Name contains characters that are invalid UTF-8; must solve issue directly on the filer
****   IS = Inconsistent attributes on one of this directory's stripes (discovered)
****   MI = Attributes are consistent, metadata marked as inconsistent
****   SI = Attributes are inconsistent, metadata not marked as inconsistent


Type                        Share                 Path
--------------------------  --------------------  -------------------------------
[LF                      ]  [charts            ]  /copyRandomx.exe
[   F8                   ]  [charts            ]  /KMO_ME~1.DAT -> kmo_medical_record.dat
[   D8                   ]  [charts            ]  /RECORD~1/ -> records_predating_y2k

**** Total Found Items:                    0
**** Total Lost Items:                     1
**** Total Invalid Filehandles:            0
**** Total Migrating Files:                0
**** Total Deleted Before Close:           0
**** Total Locking Errors:                 0
**** Total Filer Errors:                   0
**** Total 8.3 Errors:                     2
**** Total Found Stripes:                  0
**** Total Inconsistent Attrs:             0

**** Total processed:             223
**** Elapsed time:           00:00:00
**** Inconsistencies Report: DONE at 01/19/2012 01:29:11 -0500 ****
```

Figure 7.18 *Sample Output: show reports type Imp*

```
bstnA> show reports type Imp


  reports
    Codes: Imp=Import
      import.1.budget.5.rpt   06/17 00:47   2.8 kB        Imp  DONE: 928 in 00:00:03
      import.10.backlots.17.rpt 06/17 00:48   2.0 kB        Imp  DONE: 1 in 00:00:03
      import.11.scanners.18.rpt 06/17 00:48   2.2 kB        Imp  DONE: 5 in 00:00:04
      import.12.shr1-old.28.rpt 06/17 01:07   6.1 kB        Imp  DONE: 161 in 00:00:10
      import.13.shr1-next.29.rpt 06/17 01:07   2.4 kB        Imp  DONE: 14 in 00:00:04
      import.2.bills.6.rpt    06/17 00:47   1.9 kB        Imp  DONE: 132 in 00:00:01
      import.3.bills2.7.rpt   06/17 00:47   2.5 kB        Imp  DONE: 415 in 00:00:01
      import.4.it5.8.rpt      06/17 00:47   1.9 kB        Imp  DONE: 131 in 00:00:01
      import.5.rx.11.rpt      06/17 00:48   2.0 kB        Imp  DONE: 34 in 00:00:03
      import.6.charts.12.rpt  06/17 00:48   2.3 kB        Imp  DONE: 165 in 00:00:04
      import.7.bulk.13.rpt    06/17 00:48   1.9 kB        Imp  DONE: 5 in 00:00:05
      import.8.equip.15.rpt   06/17 00:48   2.0 kB        Imp  DONE: 19 in 00:00:04
      import.9.leased.16.rpt  06/17 00:48   1.9 kB        Imp  DONE: 42 in 00:00:03
```

show reports status

Purpose Use this command to show a one-line summary of a given report.

Mode (any)

Security Role(s) operator (any except backup operator)

Syntax **`show reports status`** ***`report-name`***

report-name (optional, 1-255 characters) is the name of the report to summarize.

Guidelines Use this command to get a one-line status for the job that created the report. The line is the same as the one shown by show reports. The line shows a summary status for the job that generated the report, with the following fields:

report-name`.rpt` *last-mod-time size type status*

report-name identifies the report. For time-sensitive reports, this name contains a date stamp for when the report was created, either manually or on a schedule. The report is created when an administrator invokes a report-generating command or when a rule's schedule fires.

last-mod-time is the time of the report's most-recent change.

size is the amount of internal-disk space that the report uses. You can use copy ftp and similar commands to copy reports off the ARX, and delete reports report-name to clear a report from the disk.

type indicates the process that created the report. This column uses the same abbreviations used for the type report-type option, above.

status shows the current or final state of the operation (typically DONE or FAILED) followed by a summary of the results. The summary shows a number of objects processed (such as files and directories) and the elapsed time for the report run. The elapsed time is the time since the report was created, which may be longer than the actual time the process ran; view the report contents to find actual processing times.

Sample `bstnA# show reports status status import.2.bills.6.rpt`
shows the status for an import report. See Figure 7.19 for sample output.

Related Commands show reports

Figure 7.19 *Sample Output: show reports status*

```
bstnA# show reports status import.2.bills.6.rpt


  reports
    Codes: Imp=Import
      import.2.bills.6.rpt    06/17 00:47   1.9 kB        Imp  DONE: 132 in 00:00:01
```

show scripts

Purpose Use this command to display the contents of the scripts directory and/or the contents of a specified script file.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show scripts [`*`file-name`*`]`**

file-name (optional, 1-255 characters) is the name of the file to display. If you omit this, the command displays the list of the directory's files.

Guidelines Issue the show scripts command to view the directory contents. Then issue the command with a file name to view the file contents.

Samples

```
bstnA# show scripts


  scripts
      analyze-failover        06/20 20:00   95 B
      change-mfg-date.scr     06/20 20:00   631 B
      check-global.scr        06/20 20:00   1.9 kB
      check-hw.scr            06/20 20:00   1.8 kB
      check-run.scr           06/20 20:00   3.0 kB
      clean_reports           06/20 20:00   1.3 kB
      cli_script_testdir.scr  11/16/2010     0 B
      cli_script.scr          11/16/2010     0 B
      global                  06/29 02:05   24 kB
      import_rate             06/20 20:00   16 kB
      jiltdump                06/20 20:00   2.1 kB
      logging.scr             07/10/2010    218 B
      monitor                 06/20 20:00   1.3 kB
      nslookup                06/20 20:00   773 B
      power                   06/20 20:00   1.6 kB
      run_cfg.scr             07/08/2010    10 kB
      running                 06/29 02:05   12 kB
      schemadump.sql          06/08 14:34   696 B
      share_status            06/20 20:00   4.1 kB
      show                    06/20 20:00   660 B
      show_chassis.scr        05/18 10:37   70 B
      start_conf              06/29 02:05   36 kB
      test-global-config_orig.scr 06/15 17:27   24 kB
      test-global-config.scr  06/15 17:38   24 kB
      test-running-config_orig.scr 06/15 17:27   10 kB
      test-running-config.scr 06/15 17:38   10 kB
      traceroute              06/20 20:00   922 B
      ttcp                    06/20 20:00   2.1 kB
```

Related Commands

copy ftp
copy scp
copy tftp
copy smtp
copy {nfs|cifs}
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
show directories

show software

Purpose Use this command to display the contents of the ARX software directory and/or the contents of a specified MIB file in that directory.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show software [file-name]`

file-name (optional, 1-1024 characters) is the name of the file to display. If you omit this, the command displays the list of the directory's files.

Guidelines The software directory contains MIBs, user documentation, and some useful scripts. Issue the show software command to view the directory contents. Then issue the command with a file name to view the file contents (this only works with the text-based files, such as MIBs and scripts).

To view the documentation, use one of the copy commands (copy ftp, copy scp, copy {nfslcifs}, copy tftp, or copy smtp) to upload it to an external machine. Then use a standard PDF reader or plugin. PDF readers are freely available on the Internet.

Sample `bstnA# show software`

shows the full contents of the software directory. See Figure 7.20 for sample output.

Related Commands
- copy ftp
- copy scp
- copy tftp
- copy {nfslcifs}
- copy smtp
- delete
- rename
- show directories

Figure 7.20 *Sample Output: show software*

```
bstnA# show software


  software
      A-VE_install.pdf        06/20 23:31   2.2 MB
      A-VE_quickstart.pdf     06/20 23:31   94 kB
      A1500_install.pdf       06/20 23:31   2.1 MB
      A1500_quickstart.pdf    06/20 23:31   491 kB
      A1k_install.pdf         06/20 23:31   1.3 MB
      A1k_quickstart.pdf      06/20 23:31   190 kB
      A2500_install.pdf       06/20 23:31   2.0 MB
      A2500_quickstart.pdf    06/20 23:30   301 kB
      A2k_install.pdf         06/20 23:31   7.8 MB
      A2k_quickstart.pdf      06/20 23:31   684 kB
      A4k_install.pdf         06/20 23:30   17 MB
      A4k_quickstart.pdf      06/20 23:31   1.0 MB
```

```
A5c_install.pdf          06/20 23:31   2.3 MB
A5c_quickstart.pdf       06/20 23:30   133 kB
A6k_install.pdf          06/20 23:31   2.6 MB
A6k_quickstart.pdf       06/20 23:31   808 kB
HD_FRU.pdf               06/20 23:31   931 kB
HW_Reference.pdf         06/20 23:31   6.1 MB
Mibs.tgz                 06/21 00:26   168 kB
PwrSupply_FRU.pdf        06/20 23:30   953 kB
acopiasmi.my             06/20 20:22   219 kB
bridge.my                02/25/2003    45 kB
cliMaintenance.pdf       06/20 23:31   3.8 MB
cliNetwork.pdf           06/20 23:31   3.2 MB
cliReference.pdf         06/20 23:31   13 MB
cliStorage.pdf           06/20 23:31   5.7 MB
compatibilityMatrix.pdf 06/20 23:31   206 kB
dot3ad.my                02/25/2003    39 kB
entity.my                06/28/2004    51 kB
etherlike.my             10/12/2004    87 kB
glossary.pdf             06/20 23:31   914 kB
ifmib.my                 10/12/2004    70 kB
ifType.my                10/07/2004    4.4 kB
logCatalog.pdf           06/20 23:31   2.9 MB
masterIndex.pdf          06/20 23:30   2.5 MB
mib-2.my                 10/07/2004    100 kB
openview.trapd.conf      06/20 20:00   218 kB
pbridge.my               02/25/2003    31 kB
SecureAgent.pdf          06/20 23:31   1.2 MB
SlideRail_FRU.pdf        06/20 23:31   1.2 MB
releaseNotes.html        06/20 23:31   295 kB
rfc2668.my               10/12/2004    101 kB
rmon.my                  02/25/2003    147 kB
sitePlanning.pdf         06/20 23:31   2.0 MB
snap-recon.pl            06/20 20:00   18 kB
snmpReference.pdf        06/20 23:30   1.2 MB
snmpv2-mib.my            10/07/2004    36 kB
snmpv2-smi.my            10/07/2004    1.7 kB
stamp-mibs-tgz           06/21 00:26    0 B
vlan.my                  02/25/2003    69 kB
```

show stats-logs

Purpose

The ARX has an internal stats-monitor process that monitors communication with external devices (such as filers and clients) and places statistical data about this communication in the "stats-logs" directory. Use this command to display the contents of the stats-logs directory or the contents of a specific stats-log file.

Mode

(any)

Security Role(s)

crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

`show stats-logs [`***`file-name`***`]`

file-name (optional, 1-255 characters) is the name of the file to display. If you omit this, the command displays the list of the directory's files.

Guidelines

Issue the command to view the directory contents. Then issue the command with a file name to view the file contents.

The stats-logs are all in CSV format, so that you can copy them onto an external device and use Microsoft Excel (or a similar spreadsheet application) to view them as spreadsheets. You can use copy ftp, copy {nfs|cifs}, copy smtp, or similar copy commands to copy these files off of the ARX.

If you select a stats-monitor sampling interval that is shorter than the default, the stats-monitor process writes to these .csv files more frequently. These .csv files are kept (together with all other log files) on a separate disk partition from configuration files and other important data. Each hour, the stats-monitor creates an hourly file with the first line of the new hour; this is a summary of the data in the raw file. The stats-monitor files are not permitted to consume more than 80% of the log partition. If the .csv files approach their size limit, the oldest raw-data files are removed and the hourly-summary files remain.

Guidelines: File Names

The stats-log files use the following naming convention:

source_YYYYMMDD_hhmmss.{raw|hourly}.stats.csv

where

- ***source*** is the source process or group of devices that generated the statistics. These are
  - "cifs-service" (related to a cifs front-end service and its clients; these statistics are similar to the CIFS output from show statistics namespace ... summary),
  - "cifs-service-auth" (CIFS-authentication statistics, also visible with show statistics cifs authentication),
  - "nfs-service" (for an nfs service and its clients, similar to the NFS output from show statistics namespace ... summary),
  - "cifs-share" or "nfs-share" (back-end shares behind an ARX volume, viewable with show statistics namespace ... summary on an individual share),
  - "metadata" (used to locate files behind a managed volume, and stored on an external metadata share; you can also see these statistics with show statistics metadata),

Guidelines: File Names (Cont.)

  - "metalog" (volume-state information required for redundancy failovers; these statistics are also viewable with show statistics metalog),
  - "cifs-work-queue" (internal data structures for managing CIFS tasks, also viewable with show statistics cifs work-queues),
  - "domain-controller" (related to traffic between internal CIFS processes and external DCs; these statistics are similar to the output from show statistics domain-controller), and
  - "migration" (related to place-rules and other policies that migrate files between back-end shares; these statistics are related to the output from show statistics migration).
- ***YYYYMMDD_hhmmss*** is the date and time the file was created.
- **raw|hourly** differentiates between a file with raw data and an hourly file with roll-up data. The stats-monitor process writes to each raw-data file as often as dictated by the sampling interval setting. It creates an hourly file for each source once per hour, with the first set of data for the next hour. If the sampling interval is less than one hour, the hourly files act as summaries of the raw-data files: they show the data samples at the beginning of the hour, but not the changes in those values during the hour.

Guidelines: File Format

The top line of any stats-log file has the following format:

```
BEGIN source version=v timestamp=nnnnnnnnnn
```

where

- *source* is the same as the source in the file name, defined above,
- *v* is a version number for the stats-monitor, and
- *nnnnnnnnnn* is the number of seconds since the first second of 1970, UTC.

The second line shows the data types of each field in the file's rows. The possible data types are

- TIMESTAMP identifies a time as the number of seconds since epoch, as above.
- KEY is a unique value that identifies the object being monitored (such as a volume-group ID).
- COUNTER is an integer to show a count.
- INTERVAL and
- TOTALTIME are time intervals, in microseconds (millionths of a second).
- TEXT

The third row is the name of each field.

The remaining rows contain data, arranged in columns below the names in the above row.

STATS_RESET marks when a reboot, a clear statistics cifs work-queues, or a clear statistics filer occurred to clear some statistics. This has a "ts" (time-stamp) field, followed by some KEY fields to identify the object statistics that were cleared.

Samples

```
bstnA# show stats-logs
```

shows a list of all stats-log CSV files. See Figure 7.21 for sample output.

```
provA# show stats-logs metadata_20110705_041347.raw.stats.csv
```

shows a particular stats-log CSV file. See Figure 7.22 on page 7-78 for sample output.

Related Commands

stats-monitor -> sampling interval
copy ftp
copy scp
copy tftp
copy {nfs|cifs}
copy smtp
move
move ... ftp
move ... scp
move ... tftp
move ... {nfs|cifs}
delete
show directories

Figure 7.21 *Sample Output: show stats-logs*

```
bstnA# show stats-logs


  stats-logs
      cifs-service_20121109_051448.raw.stats.csv             11/09 03:04   14 kB
      cifs-service_20121109_060000.hourly.stats.csv          11/09 03:04   5.1 kB
      cifs-service-auth_20121109_051448.raw.stats.csv        11/09 03:04   5.0 kB
      cifs-service-auth_20121109_060000.hourly.stats.csv     11/09 03:04   1.8 kB
      cifs-share_20121109_051448.raw.stats.csv               11/09 03:04   384 kB
      cifs-share_20121109_060000.hourly.stats.csv            11/09 03:04   45 kB
      cifs-work-queue_20121109_051448.raw.stats.csv          11/09 03:04   12 kB
      cifs-work-queue_20121109_060000.hourly.stats.csv       11/09 03:04   6.0 kB
      domain-controller_20121109_051448.raw.stats.csv        11/09 03:04   99 kB
      domain-controller_20121109_060000.hourly.stats.csv     11/09 03:04   11 kB
      metadata_20121109_051448.raw.stats.csv                 11/09 03:04   7.8 kB
      metadata_20121109_060000.hourly.stats.csv              11/09 03:04   1.2 kB
      metalog_20121109_051448.raw.stats.csv                  11/09 03:04   3.6 kB
      metalog_20121109_060000.hourly.stats.csv               11/09 03:04   868 B
      migration_20121109_051448.raw.stats.csv                11/09 03:04   5.1 kB
      migration_20121109_060000.hourly.stats.csv             11/09 03:04   741 B
      nfs-service_20121109_051448.raw.stats.csv              11/09 03:04   8.2 kB
      nfs-service_20121109_060000.hourly.stats.csv           11/09 03:04   1.9 kB
      nfs-share_20121109_051448.raw.stats.csv                11/09 03:04   98 kB
      nfs-share_20121109_060000.hourly.stats.csv             11/09 03:04   11 kB
```

Figure 7.22 *Sample Output: show stats-logs metadata_...*

```
provA# show stats-logs metadata_20121109_051605.raw.stats.csv
BEGIN metadata version=1 timestamp=1352438165
TIMESTAMP,KEY,KEY,COUNTER,COUNTER,TOTALTIME,COUNTER,COUNTER,TIMESTAMP,COUNTER,COUNTER,TOTALTIME,CO
UNTER,COUNTER,TIMESTAMP,COUNTER,TOTALTIME,COUNTER,COUNTER,TIMESTAMP
Timestamp,namespace,volume,reads,read_bytes,read_time,read_errors,read_current,last_read,writes,wr
ite_bytes,write_time,write_errors,write_current,last_write,syncs,sync_time,sync_errors,sync_curren
t,last_sync
STATS_RESET ts=1352451714,provMed,/mds
1352451811,provMed,/mds,29,61440,93,0,0,0,179,356864,585,0,0,0,67,44566,0,0,0
STATS_RESET ts=1352451719,provMed,/rns
1352451811,provMed,/rns,18,52224,58,0,0,0,66,191488,357,0,0,0,55,21676,0,0,0
1352452111,provMed,/mds,29,61440,93,0,0,0,293,540672,988,0,0,0,88,87580,0,0,0
1352452111,provMed,/rns,18,52224,58,0,0,0,96,249856,574,0,0,0,76,49023,0,0,0
1352452411,provMed,/mds,29,61440,93,0,0,0,293,540672,988,0,0,0,88,87580,0,0,0
1352452411,provMed,/rns,18,52224,58,0,0,0,96,249856,574,0,0,0,76,49023,0,0,0
1352452711,provMed,/mds,29,61440,93,0,0,0,293,540672,988,0,0,0,88,87580,0,0,0
1352452711,provMed,/rns,18,52224,58,0,0,0,96,249856,574,0,0,0,76,49023,0,0,0
1352453011,provMed,/mds,39,66560,115,0,0,0,305,553984,1033,0,0,0,120,90290,0,0,0
1352453011,provMed,/rns,18,52224,58,0,0,0,106,254976,593,0,0,0,100,50141,0,0,0
1352453311,provMed,/mds,39,66560,115,0,0,0,308,566272,1065,0,0,0,123,92382,0,0,0
1352453311,provMed,/rns,18,52224,58,0,0,0,106,254976,593,0,0,0,103,50152,0,0,0
```

tail

Purpose Use the tail command to display the end of a file selected from a specified directory.

Mode (any)

Security Role(s) network-technician

Syntax

```
tail {releases | logs | stats-logs | cores | configs |
  replicated-configs | reports | software | scripts | capture |
  license}
  file-name [lines | follow]
tail capture file-name [lines]
```

releases | ... | license identifies the directory. This is a required choice.

file-name (1-255 characters) identifies the file to display.

This command shows the last 24 lines of the file. You can use the optional ***lines*** argument or **follow** flag to change the number of lines:

lines (optional; 1-4096) is the number of lines at the end of the file to display.

follow (optional) makes the CLI follow the file as it grows.

Guidelines Use show directories to view a list of all files in all directories. Use show *directory-name* (for example, show logs or show capture) to list the files in one directory.

Samples

```
bstnA# tail logs syslog follow
```

tails the syslog file as it grows. This displays all syslog messages in real time.

```
bstnA# tail logs syslog 100
```

displays the last 100 lines of the syslog file.

```
bstnA# tail scripts myscript.scr
```

displays the last 24 lines of the "myscript.scr" script file.

```
bstnA# tail capture cifsVol.cap follow
```

tails a packet-capture file, "cifsVol.cap," as it grows. You can use the capture session command to start capturing packets and sending them to a file.

Related Commands

grep
pause
show directories
show releases, show logs, show stats-logs, show cores, show configs, show replicated-configs, show reports, show software, show scripts, show capture, show license

truncate-report

Purpose To stop a process from generating a report, use the truncate-report command. This truncates the report without otherwise affecting the process.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`truncate-report`** ***`file-name`***

file-name (1-255 characters) identifies the report file to truncate. This must be a file in the "reports" directory; use show reports for a file listing.

Guidelines This command prompts for confirmation before truncating the report; enter **yes** to continue.

The nsck utility can generate several report types, one per nsck job. To stop the nsck utility from writing to the report, use this command. The utility continues processing after the report is truncated.

This command is used the same way for any report, from any process.

Samples

```
bstnA# truncate-report last-modified.22.rpt
Truncate report ''last-modified.22.rpt''? [yes/no] yes
bstnA#
```

truncates a report.

Related Commands show reports

wait-for report

Purpose Use the wait-for report command to wait until the ARX finishes writing a report.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`wait-for report`** ***`report-name`*** **`[timeout`** ***`timeout`*****`]`**

report-name (1-255 characters) identifies the report. You can use the show reports command to list all of them.

timeout (optional, 1-2096) is a timeout value, in seconds. This sets a time limit; if the report takes longer than this to complete, the wait-for operation quits.

Default(s) ***timeout*** - 0 (zero, meaning that the wait-for command should wait indefinitely)

Guidelines This command is useful in CLI scripts.

Many commands create reports to show their progress (such as nsck ... report metadata-only), and some internal operations generate reports (such as the various reports produced as a by-product of a share enable (gbl-ns-vol-shr)). You can use this command to wait for any given process to finish writing a report.

If you set a ***timeout*** and it expires before the report is finished, the command exits with a warning.

Sample

```
bstnA# wait-for report remService_medco_201005260714.rpt
```

waits for the ARX to finish writing a remove-service report.

Related Commands show reports

8

Master Key

The *master key* encrypts and decrypts all of the Critical-Security Parameters (CSPs), such as passwords, on the ARX. You can use the master key in conjunction with the show running-config and show global-config commands to backup and restore the full switch configuration, including passwords.

show master-key

Purpose The ARX supports a single *master key* that encrypts and decrypts all of its CSPs (such as passwords). You generate the master key as part of the switch's initial boot process; use the show master-key command to get an encrypted copy of the master key.

Mode (any)

Security Role(s) crypto-officer

Syntax

```
show master-key
```

Guidelines This command prompts you for two passwords:

- System Password is a password entered at initial-boot time. It is 12-32 characters long. This validates that you have permission to access the master key.
- Wrapping Password is set with this command. The security software uses this to encrypt (and later decrypt) the master-key string.

  Enter 12-32 characters. At least one character in this password must be a number (0-9) or a symbol (!, @, #, $, and so on).

  Save this password: you will need it to decrypt the master key later, on the secondary switch.

This command outputs a base64-encoded string that is the encrypted master key. Save this string *and* the wrapping password that you set in the command.

You can use these pieces of information to duplicate the master key later on a redundant switch; both switches in a redundant pair must share the same master key. If you set up two redundant pairs in a disaster-recovery configuration, where one pair is an active cluster (see cluster-name) and the other is a backup cluster, all four switches must share the same master key.

For maximum security, the encrypted master key and its wrapping password *must be saved separately*.

Guidelines: Resetting the Master Key

There are occasions where you may need to reset your master key. For example, your chassis may be designated for use in a backup cluster for a disaster-recovery setup, and may need the same master key as the switches in the active cluster. If the master key was not copied during installation, you must reset the switch to its factory defaults to change it.

You must clear your entire configuration to reset your master key. You can restore the running-config (network parameters), but the global-config (storage parameters) should remain clear for a backup switch. Follow these steps to reset a switch back to its factory defaults and reset its master key:

- Use the copy running-config command to copy the entire running configuration (network-level parameters, not storage parameters) into a file on the chassis.
- The ARX should not be running any storage services if it is designated as a backup; a backup is designed to take over services from the active switch or cluster after a failover. Use the remove service command to clear each service that is currently running, if any. This cleanly de-couples the ARX software from all back-end filers.
- Connect a serial cable to the Console port. Only the Console port is available after you take the next step.
- Use delete startup-config and delete configs boot-config to delete the entire configuration, and then run the reload command. This resets the machine to its factory defaults, disables all management-IP interfaces, and reruns the initial-boot script at the Console port.
- Use the initial-boot script to reset the master key. For detailed instructions on the initial-boot process, refer to the hardware-installation guide for your chassis.
- Invoke the run command on the running-config file that you saved onto the chassis earlier. This re-establishes all of your network parameters.

Sample

```
bstnA(cfg)# show master-key
System Password: Sup3r$ecretpw
Wrapping Password: An0ther$ecretpw
Validate Wrapping Password: An0ther$ecretpw

Encrypted master key:

2oftVCwAAAAgAAAApwazSRFd2ww/H1pi7R7JMDZ9SoIg4WGA/XsZP+HcXjsIAAAADDRbMC
xE/bc=
bstnA(cfg)#
```

Related Commands

9

Layer 2

blocked-vlan (cfg-vlan)

Purpose Use this command to block an interface (a VLAN member) or a range of interfaces from receiving ingress traffic from the current VLAN. (Contrast this with no members (cfg-vlan), which removes the interface from the VLAN altogether.)

Mode cfg-vlan

Security Role(s) network-engineer or crypto-officer

Syntax

blocked-vlan *slot/port* [to *slot/port*]
no blocked-vlan *slot/port* [to *slot/port*]

slot/port (2/1-14 on ARX-4000; 1/1-4 or 2/1-2 on ARX-2500; 1/1-12 on ARX-2000; 1/1-8 on ARX-1500) is the first (or only) Ethernet port.

to *slot/port* (optional) is the last port in a range of ports.

Default(s) None

Valid Platforms ARX-2000 and ARX-4000

Guidelines This blocks the interface from accepting incoming packets, but keeps using the interface for outbound packets. This only applies to packets from the current VLAN.

Use the show interface summary command to locate all Ethernet ports on the chassis.

Samples

```
bstnA(cfg-vlan[1])# blocked-vlan 2/5 to 2/6
```

blocks ports 2/5 - 2/6 from VLAN 1.

```
bstnA(cfg-vlan[7])# no blocked-vlan 2/3
```

removes port 2/3 from being blocked on VLAN 7.

Related Commands

show interface
show vlan

channel

Purpose

A *channel* is an aggregated group of Ethernet ports that function as one link, as defined in IEEE 802.3ad. You can configure up to eight channels on the ARX. From cfg mode, use the channel command to begin configuring a channel.

Use no channel to remove the channel number configuration and return all member ports to their independent roles.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
channel number
no channel number
```

number (1-8) is an ID you choose for the channel. If the channel is already configured, this command edits its configuration. The no form of the command removes the channel configuration.

Default(s)

None

Valid Platforms

ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines

We recommend disconnecting the ARX from the peer station before configuring the channel on either side. Once the channel is configured on both switches, connect the cables between them.

This command puts you into cfg-channel mode, where you use members (cfg-channel) to configure the channel's member ports and no shutdown (cfg-channel) to start traffic on the channel. For a channel that connects a redundant pair, configure the channel's member ports with redundancy protocol (cfg-channel). You can optionally use the description (cfg-channel) command to set an optional name for the channel. The vlan (cfg-channel) command assigns the channel to a VLAN, untagged. The vlan-tag command assigns the channel to a VLAN in tagged mode; you can assign the channel to multiple VLANs by invoking this command once for each VLAN.

By default, outbound traffic is load-balanced by hashing both the source and destination IP of an outbound packet. On some platforms you have the option to change the IP addresses used in this hash with the load-balance command.

Use the show channel command to view the channel's configuration.

All ports in a channel must be in the same VLAN(s). The station on the other side of the channel must have the same member ports and VLANs as the ARX; traffic cannot flow through the channel unless the channel configurations match. Additionally, the ports cannot have an "auto" speed.

If spanning tree is running, the channel's VLAN determines its spanning-tree membership.

Guidelines: Shutting Down the Redundant-Pair Link

If you use the no channel command on the channel that carries the redundant-pair link (see redundancy protocol (cfg-channel)) while redundancy is enabled (enable (cfg-redundancy)), the command causes the standby peer to reboot. The reboot does not disrupt any storage services, but the ARX peers cannot function as a redundant pair while the link is shut down. Additionally, a quorum-disk failure or disconnection would cause the active peer to reboot, too. If you proceed with shutting down the link, you should establish a new one as soon as possible: use the redundancy protocol or redundancy protocol (cfg-channel) command on another port or channel to establish a new redundant-pair link.

The CLI prompts for confirmation before removing the channel and shutting down the redundant-pair link; enter **yes** to proceed with the channel removal and the reboot.

Sample

```
bstnA(cfg)# channel 1
bstnA(cfg-channel[1])#
```

creates channel 1.

Related Commands

members (cfg-channel)
redundancy protocol (cfg-channel)
description (cfg-channel)
vlan (cfg-channel)
vlan-tag
load-balance
shutdown (cfg-channel)
show channel

clear counters channel

Purpose Use this command to clear the current Ethernet statistics for all channels or for a specified channel. This command restarts the count for all channel statistics. It also restarts statistics for every member port within the channel(s).

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `clear counters channel [channel-id]`

channel-id (1-8) identifies a single channel to clear. If you omit this option, the command clears the statistics for all channels. Use the show channel summary command to enumerate all configured channels.

Default(s) None

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines The CLI prompts for confirmation before it clears any counters; type **yes** to proceed with the operation.

Use this command to clear and restart the statistics counter for troubleshooting and monitoring channels. Use show channel ... stats to view these statistics. To view the individual statistics for the channel's member ports, which are also cleared with this command, use show interface gigabit stats or show interface ten-gigabit stats.

Sample

```
bstnA# clear counters channel 2

Clear the counters for all of the interfaces associated with channel 2?
[yes/no] yes
bstnA#
```

clears the current statistics count for channel 2.

Related Commands

show channel
show channel ... stats
show interface gigabit stats
show interface ten-gigabit stats

clear counters gigabit

Purpose Use this command to clear the current Gigabit Ethernet statistics for all Gigabit ports or for a specified slot/port and restart the count.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear counters gigabit [`*`slot/port`*`]`**

slot/port (2/1-14 on ARX-4000; 1/1-4 on ARX-2500; 1/1-12 on ARX-2000; 1/1-8 on ARX-1500) is a Gigabit Ethernet port. Use the show interface summary command to locate all Gigabit-Ethernet ports and their slot(s).

Default(s) None

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines Use this command to clear and restart the statistics counter for troubleshooting and monitoring Gigabit Ethernet ports. Use show interface gigabit stats to view the interface's traffic statistics.

Sample

```
bstnA# clear counters gigabit 2/7
bstnA#
```

clears the current statistics count at port 2/7.

Related Commands show interface gigabit **`stats`**
show interface **`summary`**

clear counters lacp

Purpose

Link-Aggregation Control Protocol (LACP) is a control protocol for dynamically managing the member ports in a channel. The peers on both ends of the channel use LACP to exchange information about member ports; the peers can automatically remove or replace member links if configuration changes disqualify/re-qualify them for channel membership.

For LACP statistics, use the show channel ... lacp stats command (see the documentation for show channel ... stats). Use this command to clear the current LACP statistics for all channels (or for a specified channel) and restart the count.

Mode

priv-exec

Security Role(s)

network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

`clear counters lacp [channel` ***`id`*****`]`**

id (1-8) identifies a single channel to clear. If you omit this option, the command clears the statistics for all channels. Use the show channel summary command to enumerate all configured channels.

Default(s)

None

Valid Platforms

ARX-1500 and ARX-2500 only

Guidelines

The channel command creates a channel, and the lacp active or lacp passive command enables LACP on the channel. Use this command to clear and restart the statistics counters for LACP. Use show channel ... stats to view a channel's statistics (including LACP statistics).

The CLI prompts for confirmation before clearing the LACP counters; enter **yes** to proceed.

Sample

```
stoweA# clear counters lacp channel 2
Clear the LACP counters for channel 2? [yes/no] yes
```

clears the current statistics count for channel 2.

Related Commands

show channel ... stats
lacp active
lacp passive
channel

clear counters ten-gigabit

Purpose Use this command to clear the current Ethernet statistics for all Ten-Gigabit ports or for a specified slot/port and restart the count.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear counters ten-gigabit [`*`slot/port`*`]`**

slot/port (2/1-2) is a Gigabit Ethernet port. Use the show interface summary command to locate all Gigabit-Ethernet ports and their slot(s).

Default(s) None

Valid Platforms ARX-2500 and ARX-4000 only

Guidelines Use this command to clear and restart the statistics counter for troubleshooting and monitoring Gigabit Ethernet ports. Use show interface ten-gigabit stats to view the interface's traffic statistics.

Sample

```
bstnA# clear counters ten-gigabit 2/2
```

clears the current statistics count at port 2/2.

Related Commands show interface gigabit **`stats`**
show interface **`summary`**

clear counters redundancy network

Purpose Use this command to clear the link-transition counters in the show redundancy network output.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `clear counters redundancy network`

Default(s) None

Valid Platforms any *except* ARX-VE

Guidelines The link-transition counter increments every time the link status changes for a redundancy-protocol port (the link goes up or down). A redundancy-protocol port is used as one end of a redundant-pair link. Use the show redundancy network command to see the link-transition counter.

Sample

```
bstnA# clear counters redundancy network
bstnA#
```

clears the link-transition counters.

Related Commands show redundancy network
clear counters redundancy

description (cfg-channel)

Purpose

A link-aggregation *channel* (IEEE 802.3ad) can optionally have a description to display in its show commands. From cfg-channel mode, use the description command to create a description for the current channel.

Use no description to remove the description.

Mode

cfg-channel

Security Role(s)

network-engineer or crypto-officer

Syntax

`description` ***`description`***
`no description`

description (1-15 characters) is a text string description for the current channel. Insert quotation marks around the description if it contains spaces.

Default(s)

None

Valid Platforms

ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines

The description appears in the output from the show channel command.

Samples

```
bstnA(cfg-channel[1])# description “trunk 2”
```

sets a description for channel 1.

```
bstnA(cfg-channel[6])# no description
```

erases the description for channel 6.

Related Commands

show channel

description (cfg-if-gig)

Purpose

A port can optionally have a description for its show commands. Use the description command to create a description for the current port.

Use no description to remove the description.

Mode

cfg-if-gig

Security Role(s)

network-engineer or crypto-officer

Syntax

`description` ***`description`***
`no description`

description (1-60 characters) is the description you choose for the current port. Quote the description if it contains spaces.

Default(s)

No description

Guidelines

The description appears in the output from the show interface gigabit command.

Samples

```
bstnA(cfg-if-gig[2/4])# description “link to back-end filers”
```

sets a description for the current port, 2/4.

```
bstnA(cfg-if-gig[2/8])# no description
```

erases the description for port 2/8.

Related Commands

interface gigabit
show interface gigabit

description (cfg-if-ten-gig)

Purpose

A ten-gigabit port can optionally have a description for its show commands. Use the description command to create a description for the current ten-gigabit port.

Use no description to remove the description.

Mode

cfg-if-ten-gig

Security Role(s)

network-engineer or crypto-officer

Syntax

`description` ***`description`***
`no description`

description (1-60 characters) is the description you choose for the current ten-gigabit port. Quote the description if it contains spaces.

Default(s)

No description

Valid Platforms

ARX-2500 and ARX-4000 only

Guidelines

The description appears in the output from the show interface ten-gigabit command.

Samples

```
bstnA(cfg-if-ten-gig[2/2])# description “link to big-ip”
```

sets a description for the current ten-gigabit port, 2/2.

```
bstnA(cfg-if-ten-gig[2/1])# no description
```

erases the description for port 2/1.

Related Commands

interface ten-gigabit
show interface ten-gigabit

description (cfg-vlan)

Purpose

A VLAN can optionally have a description for its show commands. Use the description command to create a description for the current VLAN.

Use no description to remove the description.

Mode

cfg-vlan

Security Role(s)

network-engineer or crypto-officer

Syntax

`description` ***`description`***
`no description`

description (up to 80 characters) is the description you choose for the current VLAN. Quote the description if it contains spaces.

Default(s)

"default"

Valid Platforms

any *except* ARX-VE

Guidelines

The description appears in the output from the show vlan and show vlan summary commands.

Samples

```
bstnA(cfg-vlan[1])# description "ARX-defined VLAN"
```

sets a description for the current VLAN.

```
bstnA(cfg-vlan[7])# no description
```

erases the description for VLAN 7.

Related Commands

show vlan
show vlan summary

forward-delay

Purpose

The *Forward Delay* is the time for spanning-tree ports to stay in the listen and learn states, waiting for the best BPDU frame to reach the ARX. Use the forward-delay command to set the Forward-Delay time.

Use the no form to revert to the default Forward Delay.

Mode

cfg-stp

Security Role(s)

network-engineer or crypto-officer

Syntax

```
forward-delay seconds
no forward-delay
```

seconds is a number from 4 to 30.

Default(s)

15 (seconds)

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

The Spanning-Tree protocol is defined in IEEE 802.1D.

The Forward Delay time should be at least twice the maximum transit time for a BPDU to traverse the entire network. This allows the bridges enough time to establish a new spanning-tree topology in case of a bridge or link failure.

Sample

```
bstnA(cfg-stp)# forward-delay 10
```

sets the Forward Delay to 10 (seconds).

Related Commands

spanning-tree

flowcontrol

Purpose

A high-speed (1000-tx-full or ten-gigabit) interface can transmit a flow-control request and/or accept and enforce a flow-control request for an over-burdened peer. Use the flowcontrol command to enable flow control.

Use the no form to disable flow control.

Mode

cfg-if-gig

cfg-if-ten-gig

Security Role(s)

network-engineer or crypto-officer

Syntax

```
flowcontrol send {on | off}
flowcontrol receive {on | off}
no flowcontrol
```

on | **off** is a required choice.

Default(s)

Disabled

Valid Platforms

In cfg-if-gig mode: ARX-1500, ARX-2000, ARX-2500, and ARX-4000

In cfg-if-ten-gig mode: ARX-2500 and ARX-4000 only

Guidelines

On a one-gigabit interface, the speed setting must be 1000-tx-full to enable flow control. Use the speed (cfg-if-gig) command to set the speed on a single-gigabit interface. The show interface gigabit and show interface ten-gigabit commands show the current flow-control setting on a particular one-gigabit or ten-gigabit interface, respectively.

This command is not supported on the ARX-500 or ARX-VE platforms.

Sample

```
bstnA(cfg-if-gig[2/4])# flowcontrol send on
```

sets the interface to transmit flow-control requests. This does not necessarily accept any flow-control requests from peers.

Related Commands

interface gigabit
speed (cfg-if-gig)
show interface gigabit
show interface ten-gigabit

hello-time

Purpose

The *Hello Time* is the interval (in seconds) between broadcasts of Bridge Protocol Data Units (BPDUs) to neighboring bridges. The BPDUs have spanning-tree topology information that a bridge uses to determine its role in the spanning tree. Use the hello-time command to set the Hello Time.

Use the no form to revert to the default Hello Time.

Mode

cfg-stp

Security Role(s)

network-engineer or crypto-officer

Syntax

`hello-time` ***`seconds`***
`no hello-time`

seconds is a number from 1 to 10.

Default(s)

2 (seconds)

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

The Spanning-Tree protocol is defined in IEEE 802.1D.

Samples

```
bstnA(cfg-stp)# hello-time 5
```

sets the Hello Time to 5 (seconds).

```
bstnA(cfg-stp)# no hello-time
```

sets the Hello Time to its default, 2 (seconds).

Related Commands

spanning-tree

interface gigabit

Purpose

A layer-2 port is called an *interface* in the CLI. Use the interface gigabit command to begin configuring an interface.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`interface gigabit` ***`slot/port`***

slot/port (2/1-14 on ARX-4000; 1/1-4 on ARX-2500; 1/1-12 on ARX-2000; 1/1-8 on ARX-1500; 1/2 on ARX-500; 1/1 on ARX-VE) is a Gigabit Ethernet port. Use the show interface summary command to show all Gigabit Ethernet ports and their slot(s).

Default(s)

None

Guidelines

This command puts you into cfg-if-gig mode, where you can set several configuration parameters for the port. Use the speed (cfg-if-gig) command to manually set the port speed and duplex configuration. If the device at the other end of the connection supports flow control, you can use the flowcontrol command to configure it. Use the description (cfg-if-gig) command to set an optional description for the port, for show commands. Use the no shutdown (cfg-if-gig) command to start the port.

For every platform except the ARX-1500, ARX-2500, and ARX-VE, traffic-storm control is configured automatically at every gigabit interface. The ARX forwards a maximum of 1000 packets per second for each of the following frame types: broadcast, multicast, or unicast frames with unknown destination addresses. If the ingress-side of the port reaches the maximum in a given second, the port drops packets of the chosen type until the second is over. Traffic-storm control does not suppress spanning-tree BPDU packets. The ARX-1500, ARX-2500, and ARX-VE do not support traffic-storm control.

Guidelines: Port 1/1 on the ARX-1500 and ARX-2500

By default, port 1/1 on an ARX-1500 or ARX-2500 is configured as an out-of-band (OOB) management interface. You can use the interface mgmt command and its sub commands to manage this interface. The CLI returns an error message if you enter interface gigabit 1/1 while port 1/1 is being used for out-of-band management.

To re-assign this port to client/server traffic, use no interface mgmt to delete the out-of-band management interface. Then use this command on port 1/1 (interface gigabit 1/1) to enter cfg-if-gig mode and edit the port for client/server traffic.

Sample

```
bstnA(cfg)# interface gigabit 2/6
bstnA(cfg-if-gig[2/6])#
```

edits the interface at slot 2, port 6.

Related Commands

show interface **`summary`**
speed (cfg-if-gig)
flowcontrol
description (cfg-if-gig)
shutdown (cfg-if-gig)
show interface gigabit

interface ten-gigabit

Purpose The ARX-4000 supports two ten-gigabit ports, which are called *interfaces* in the CLI. Use the interface ten-gigabit command to begin configuring a ten-gigabit interface.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax **`interface ten-gigabit`** ***`slot/port`***

slot/port (2/1-2) identifies a ten-Gigabit Ethernet port. Use the show interface summary command to show all Ethernet ports and their slot(s).

Default(s) None

Valid Platforms ARX-2500 and ARX-4000 only

Guidelines This command puts you into cfg-if-ten-gig mode, where you can set several configuration parameters for the port. The speed is fixed at 10 gigabits/second, full duplex. Use the description (cfg-if-ten-gig) command to set an optional description for the port, for show commands. If the device at the other end of the connection supports flow control, you can use the flowcontrol command to configure it. Use the no shutdown (cfg-if-ten-gig) command to start the port.

For every platform except the ARX-1500, ARX-2500, and ARX-VE, traffic-storm control is configured automatically at every ten-gigabit interface. The ARX forwards a maximum of 1000 packets per second for each of the following frame types: broadcast, multicast, or unicast frames with unknown destination addresses. If the ingress-side of the port reaches the maximum in a given second, the port drops packets of the chosen type until the second is over. Traffic-storm control does not suppress spanning-tree BPDU packets. The ARX-1500, ARX-2500, and ARX-VE do not support traffic-storm control.

Sample

```
bstnA(cfg)# interface ten-gigabit 2/1
bstnA(cfg-if-ten-gig[2/1])#
```

edits the ten-gigabit interface at slot 2, port 1.

Related Commands

show interface **`summary`**
description (cfg-if-ten-gig)
flowcontrol
shutdown (cfg-if-ten-gig)
show interface ten-gigabit

ip private vlan

Purpose Two private VLANs are used for inter-process communication in the ARX: the private VLAN and the metalog VLAN. The private VLAN carries a private IP network, used for inter-process communication. The metalog VLAN carries important namespace-software records to a battery-backed NVRAM device. These VLANs are configured during the initial boot process; it is rarely necessary to change them. Use the ip private vlan command to change one of these VLANs, or the IP subnet carried by the private VLAN.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax **`ip private vlan internal`** ***`vlan-id`*** **`[metalog`** ***`meta-vlan-id`*****`] [subnet`** ***`ip-subnet mask`*****`]`**

vlan-id (1-4095) is the number for the private VLAN.

metalog ***meta-vlan-id*** (optional; 1-4095) sets the number for the metalog VLAN. This *must* be different from the private VLAN, above.

subnet ***ip-subnet mask*** (optional) is the IP address and mask for the private subnet (for example, 169.254.14.0). The mask must be 24 bits (255.255.255.0) or less for an ARX-2000 or ARX-4000; it must be 26 bits (255.255.255.192) or less for an ARX-500. 255.255.255.0 defines a large enough subnet for any platform.

Valid Platforms ARX-500, ARX-2000, and ARX-4000

Guidelines The private and metalog VLANs must not be supported on any external LAN segments. If either VLAN is in use, use this command to change it/them.

You cannot change the private VLAN or subnet in a redundant pair (see redundancy).

You may need to change the ***ip-subnet*** in a large RON where two switches have the same private subnet. This is a rare situation, but possible. If it occurs, the current switch can only reach one of the conflicting switches over the RON: the switch that was connected to the current switch first. To reach the other switch, you must make their private subnets unique within the RON. The show ron conflicts command indicates which switches have the conflict, and shows all of the private subnets that are currently in the RON (and should therefore be avoided). The ip private subnet reassign command is designed to fix this by automatically choosing a unique private subnet; alternatively, you can set the subnet manually with this command. Go to the CLI for one of the conflicting switches and use either command to change its private subnet.

Samples

```
bstnA(cfg)# ip private vlan internal 2222
Change the private VLAN and reboot the chassis? [yes/no] yes
...
```

changes the private VLAN to 2222.

```
bstnA(cfg)# ip private vlan internal 2222 subnet 169.254.166.0
255.255.255.0
```

changes the private VLAN to 2222, leaves the metalog VLAN as is, and changes the private subnet to 169.254.166.0/24.

```
bstnA(cfg)# ip private vlan internal 2222 metalog 2223
```

changes the private VLAN to 2222 and the metalog VLAN to 2223.

```
prtlndA(cfg)# ip private vlan internal 1002 subnet 169.254.200.0
255.255.255.192
```

changes the private subnet only, perhaps because of a conflict with another switch in the RON.

Related Commands

show ron conflicts
ip private subnet reassign

jumbo mtu

Purpose The ARX supports both standard Ethernet and jumbo-frame packets. Use this command to enable or disable jumbo-frame transmission on the current VLAN.

Mode cfg-vlan

Security Role(s) network-engineer or crypto-officer

Syntax

```
jumbo mtu bytes
no jumbo mtu
```

bytes (1530-9198) establishes the size of frames on this VLAN.

Default(s) 1500 bytes (standard-frame size)

Valid Platforms any *except* ARX-VE

Guidelines Use this command only if your client-server network supports jumbo frames.

Use the no form of the command to disable jumbo-frame transmission on the switch.

Samples

```
bstnA(cfg-vlan[2])# jumbo mtu 9000
```

enables 9000-byte jumbo frames on vlan 2.

```
bstnA(cfg-vlan[2])# no jumbo mtu
```

disables jumbo frames on vlan 2.

Related Commands show vlan

lacp active

Purpose Link-Aggregation Control Protocol (LACP) is a control protocol for dynamically managing the member ports in a channel. The peers on both ends of the channel use LACP to exchange information about member ports; the peers can automatically remove or replace member links if configuration changes disqualify/re-qualify them for channel membership. Use the lacp active command to enable LACP on the current channel, with the ARX as the active peer.

Use no lacp to stop sending LACPDUs to the device at the other end of the channel.

Mode cfg-channel

Security Role(s) network-engineer or crypto-officer

Syntax

```
lacp active
no lacp active
```

Default(s) **no lacp active** - static LACP (that is, no LACP) runs by default.

Valid Platforms ARX-1500 and ARX-2500

Guidelines LACP is defined in IEEE 802.1AX-2008. If you configure LACP for the channel, the peers can exchange Link Aggregation Control Protocol Data Units (LACPDUs) over each member port. The peers can use the state information in these LACPDUs to dynamically react to configuration or topology changes. For example, if a configuration change disqualifies one of the channel's member ports for aggregation, the ARX and its peer can use LACP to automatically remove the port's link from the channel. If another configuration change makes the link suitable for aggregation again, the LACP processes can automatically restore it to the channel.

To establish LACP on the channel, enable passive LACP at the peer and use this command on the ARX. If you connect two ARX peers over a channel (see redundancy protocol (cfg-channel)), you can use this command on both ARX peers to establish LACP; one of them assumes the passive LACP role automatically.

To support LACP on chassis types other than the ARX-1500 or ARX-2500, you can use the lacp passive command.

If LACP is disabled, all member ports remain in the channel no matter what configuration or topology changes occur later.

Important

The no lacp active *command restarts all of the channel's member ports. This stops all traffic on the channel for a brief time. This is not recommended for a busy channel; perform this operation only during off hours, or on an inactive channel.*

For a channel used in a redundant-pair link (see the documentation for redundancy protocol (cfg-channel)*), this causes the backup ARX to reboot. In most cases, the reboot has no effect on client traffic.*

Guidelines (Cont.)

Use the channel command to create a channel, and use the members (cfg-channel) command to add a port to the channel. Each end of the channel should have the same LACP timeout settings; you can use the lacp rate command to change this end of the channel to a long timeout. The show channel ... lacp command shows the current configuration and status of LACP on a given channel. For LACP statistics, use the show channel ... lacp stats command (see the documentation for show channel ... stats).

Samples

```
stoweA(cfg)# channel 6
stoweA(cfg-channel[6])# no lacp active

Warning: Disabling LACP will result in temporarily loss of network
connectivity for all members of this channel.

Are you sure? [yes/no] yes
```

stops LACP processing on channel 6.

```
stoweA(cfg)# channel 1
stoweA(cfg-channel[1])# lacp active
```

starts LACP processing on channel 1.

Related Commands

channel
members (cfg-channel)
lacp passive
lacp rate
show channel ... lacp
show channel ... stats

lacp passive

Purpose

Link-Aggregation Control Protocol (LACP) is a control protocol for dynamically managing the member ports in a channel. The peers on both ends of the channel use LACP to exchange information about member ports; the peers can automatically remove or replace member links if configuration changes disqualify/re-qualify them for channel membership. Use the lacp passive command to enable LACP on the current channel.

Use no lacp passive to ignore all LACPDUs from the device at the other end of the channel.

Mode

cfg-channel

Security Role(s)

network-engineer or crypto-officer

Syntax

```
lacp passive
no lacp passive
```

Default(s)

no lacp passive - static LACP (that is, no LACP) runs by default.

Valid Platforms

ARX-2000 or ARX-4000

Guidelines

LACP is defined in IEEE 802.1AX-2008. If you configure LACP for the channel, the peers can exchange Link Aggregation Control Protocol Data Units (LACPDUs) over each member port. The peers can use the state information in these LACPDUs to dynamically react to configuration or topology changes. For example, if a configuration change disqualifies one of the channel's member ports for aggregation, the ARX and its peer can use LACP to automatically remove the port's link from the channel. If another configuration change makes the link suitable for aggregation again, the LACP processes can automatically restore it to the channel.

To establish LACP on the channel, enable active LACP at the peer and use this command on the ARX. If you connect two ARX peers over a channel (see redundancy protocol (cfg-channel)), you can use this command on both ARX peers to establish LACP; one of them assumes the active LACP role automatically.

The ARX-1500 and ARX-2500 only support active LACP; you can use the lacp active command to enable active LACP on those devices.

If LACP is disabled, all member ports remain in the channel no matter what configuration or topology changes occur later.

Important

The no lacp passive *command restarts all of the channel's member ports. This stops all traffic on the channel for a brief time. This is not recommended for a busy channel; perform this operation only during off hours, or on an inactive channel.*

For a channel used in a redundant-pair link (see the documentation for redundancy protocol (cfg-channel)*), this causes the backup ARX to reboot. In most cases, the reboot has no effect on client traffic.*

Guidelines (Cont.)

Use the channel command to create a channel, and use the members (cfg-channel) command to add a port to the channel. Each end of the channel should have the same LACP timeout settings; you can use the lacp rate command to change this end of the channel to a long timeout. The show channel ... lacp command shows the current configuration and status of LACP on a given channel. For LACP statistics, use the show channel ... lacp stats command (see the documentation for show channel ... stats).

Samples

```
prtlndA(cfg)# channel 2
prtlndA(cfg-channel[2])# no lacp passive

Warning: Disabling LACP will result in temporarily loss of network
connectivity for all members of this channel.

Are you sure? [yes/no] yes
```

stops LACP processing on channel 2.

```
bstnA(cfg)# channel 1
bstnA(cfg-channel[1])# lacp passive
```

starts LACP processing on channel 1.

Related Commands

channel
members (cfg-channel)
lacp active
lacp rate
show channel ... lacp
show channel ... stats

lacp rate

Purpose

Link-Aggregation Control Protocol (LACP) is a control protocol for dynamically managing the member ports in a channel. The peers on both ends of the channel use LACP to exchange information about member ports; the peers can automatically remove or replace member links if configuration changes disqualify/re-qualify them for channel membership.

By default, the ARX uses a Fast_Periodic_Time for its transmissions (1 LACPDU each second), and a Short_Timeout_Time (time out if the peer takes longer than 3 seconds to respond). If the peer has different LACP timer settings, the peers may periodically disconnect from each other. Use the lacp rate command to lengthen the rate to one LACPDU every 30 seconds, and to lengthen the timeout to 90 seconds.

Use no lacp rate to return the timeout to its faster defaults.

Mode

cfg-channel

Security Role(s)

network-engineer or crypto-officer

Syntax

```
lacp rate long-timeout
no lacp rate
```

Default(s)

no lacp rate -

- Fast_Periodic_Time - transmit an LACPDU once per second.
- Short_Timeout_Time - time out after 3 seconds.

Valid Platforms

ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines

The LACP timer settings should be the same at both ends of the channel.

This command is irrelevant unless you also use lacp active or lacp passive to enable LACP on the current channel.

This command disables and re-enables the channel to activate the change. The CLI prompts you with a warning about the temporary connection loss; enter **yes** to proceed.

LACP is defined in IEEE 802.3ad, and later in IEEE 802.1AX-2008. Channel peers configured with LACP can exchange Link Aggregation Control Protocol Data Units (LACPDUs) over each member port. The peers can use the state information in these LACPDUs to dynamically react to configuration or topology changes. For example, if a configuration change disqualifies one of the channel's member ports for aggregation, the ARX and its peer can use LACP to automatically remove the port's link from the channel. If another configuration change makes the link suitable for aggregation again, the LACP processes can automatically restore it to the channel.

These LACP timer settings are also defined in IEEE 802.1AX-2008. This command toggles between Fast_Periodic_Time/Short_Timeout_Time and Slow_Periodic_Time/Long_Timeout_Time.

Use the channel command to create a channel, and use the members (cfg-channel) command to add a port to the channel. The show channel ... lacp command shows the current configuration and status of LACP on a given channel. For LACP statistics, use the show channel ... lacp stats command (see the documentation for show channel ... stats).

Sample

```
bstnA(cfg-channel[1])# lacp rate long-timeout

Warning: Changing the LACP Rate will result in temporarily loss of
network connectivity for all members of this channel.

Are you sure? [yes/no] yes
```

sets a 30-second interval (and a 90-second timeout) for LACP transmissions on channel 1.

Related Commands

lacp active
lacp passive
channel
members (cfg-channel)
show channel **... lacp**
show channel ... stats

load-balance

Purpose A channel balances out-bound load amongst its ports by hashing packet IPs together and choosing a port based on the hash. Use the load-balance command to change the IP(s) used in the hash.

Mode cfg-channel

Security Role(s) network-engineer or crypto-officer

Syntax

```
load-balance {src-ip | dst-ip | src-dst-ip}
no load-balance
```

src-ip uses only the packet's source-IP address in the hash. This may not produce the best hash: the source IPs are limited to a small set of VIPs (see virtual server), proxy IPs (see ip proxy-address), and management IPs (see ip address (cfg-if-vlan) and ip address (cfg-mgmt)).

dst-ip uses the packet's destination-IP address, ignoring the source address. Destination addresses are those of clients, filers, and management stations, so this is typically a better hash than one that uses the limited set of source IPs.

src-dst-ip combines the source and destination IPs with a bit-wise XOR operation. This hash typically produces the best traffic distribution within the channel.

Default(s) **src-dst-ip**

Valid Platforms ARX-2000 or ARX-4000

Guidelines This stops and restarts the channel to reset the traffic-distribution hash. The CLI therefore prompts for confirmation before stopping traffic; type **yes** to confirm.

This is relevant to outbound traffic only.

Use show channel [load-balance] to view the current load-balancing configuration for all channels. To find the results of the hash for a packet with a particular source and destination IP, use show load-balancing.

Samples

```
bstnA(cfg-channel[9])# load-balance src-ip
Changing the load-balancing algorithm will cause the channel to disrupt
traffic.

Are you sure? [yes/no] yes
```

uses source-IP addresses to choose outbound ports in channel 9. This ignores destination-IP addresses for the hash.

```
bstnA(cfg-channel[2])# no load-balance
Changing the load-balancing algorithm will cause the channel to disrupt
traffic.

Are you sure? [yes/no] yes
```

returns channel 2 to the default; choose an outbound port based on an XOR of the packet's source and destination IPs.

Related Commands channel
show load-balancing
show channel

mac-address aging-time

Purpose

The ARX learns the MAC addresses of neighboring bridges by examining the source MACs in the frames that it receives. Each MAC address is added to (or updated in) an internal table of MAC addresses. If no new updates are received for a MAC address over the *aging time*, the switch erases a MAC address from the table. Use the mac-address aging time command to set the aging time.

Use the no form of the command to revert to the default aging time.

Mode

cfg-stp

Security Role(s)

network-engineer or crypto-officer

Syntax

```
mac-address aging-time seconds
no mac-address aging-time
```

seconds is a number from 300 to 1,000,000.

Default(s)

300 (seconds, or 5 minutes)

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

The MAC address aging time is defined in IEEE 802.1D. Use the show mac-address-table summary command to see the current aging time.

Samples

```
bstnA(cfg-stp)# mac-address aging-time 600
```

sets the aging time to 600 (seconds), or 10 minutes.

```
bstnA(cfg-stp)# no mac-address aging-time
```

reverts the aging time to the default.

Related Commands

spanning-tree
show mac-address-table summary

max-age

Purpose

The *Max Age* is the time (in seconds) to keep BPDU information from a neighboring bridge before declaring the port information "stale." If the Max Age is reached for a port, it is considered disconnected by the other bridges in the spanning tree. Use the max-age command to set the Max Age.

Use the no form to revert to the default Max Age.

Mode

cfg-stp

Security Role(s)

network-engineer or crypto-officer

Syntax

```
max-age seconds
no max-age
```

seconds is a number from 6 to 40.

Default(s)

20 (seconds)

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

The Spanning-Tree protocol is defined in IEEE 802.1D.

The Max Age is typically three times the Hello Time; it *must* be at least 2 * (Hello Time + 1). Use the hello-time command to set the Hello Time.

Samples

```
bstnA(cfg-stp)# max-age 15
```

sets the Max Age to 15 (seconds).

```
bstnA(cfg-stp)# no max-age
```

reverts the Max Age to the default.

Related Commands

spanning-tree
hello-time

members (cfg-channel)

Purpose

From cfg-channel mode, use the members command to add a single port or a range of ports to the current channel.

Use the no members command to remove a port(s) from the channel.

Mode

cfg-channel

Security Role(s)

network-engineer or crypto-officer

Syntax

```
members slot/port [to slot/port ]
no members slot/port [to slot/port ]
```

slot/port (2/1-14 on ARX-4000; 1/1-4 or 2/1-2 on ARX-2500; 1/1-12 on ARX-2000; 1/1-8 on ARX-1500) is the first (or only) Ethernet port.

to *slot/port* (optional) is the last port in a range of ports.

Use the show interface summary command to locate the slot(s) for these ports. ARX-4000 and ARX-2500 devices have ten-gigabit interfaces at ports 2/1 and 2/2 and one-gigabit interfaces at the remaining ports.

Default(s)

None

Valid Platforms

ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines

A channel can have up to 8 ports. All of the ports must be shut down before you add them with this command (see shutdown (cfg-if-gig) or shutdown (cfg-if-ten-gig)). The ARX can support different speeds amongst the channel members; confirm that the peer at the other end of the channel can support this before you configure your channel this way.

For a channel to be used as a redundant-pair link, use the redundancy protocol (cfg-channel) command instead of this command.

You cannot use the no members command on the last port in the channel; instead, use no channel to remove the entire channel.

Guidelines: Confirmation Prompt on an ARX-1500 and ARX-2500

The ARX-1500 and ARX-2500 do not support the same VLAN on two physical interfaces, where a "physical interface" is either a port or a channel. Each VLAN can be carried by only one port or channel. Therefore, the system removes all VLANs from any member port you choose before adding the port to the channel. The port then carries the same VLAN(s) as the other ports in the same channel (with vlan (cfg-channel) or vlan-tag). The CLI prompts for confirmation before taking the port off of any VLANs; enter **yes** to continue.

Guidelines: Port 1/1 on an ARX-1500 and ARX-2500

On the ARX-1500 and ARX-2500, port 1/1 is the out-of-band management interface by default, and cannot be included in any client/server channel. If you prefer to use an in-band (VLAN) management interface (see interface vlan) for accessing the CLI or GUI, you can use interface mgmt and shutdown (cfg-mgmt) to stop using port 1/1 for out-of-band management. Then you can use this command to include port 1/1 in a client/server channel.

Samples

```
bstnA(cfg-channel[2])# members 2/7 to 2/10
```
adds ports 2/7-2/10 to channel 2.

```
bstnA(cfg-channel[5])# members 2/2
```
adds port 2/2 (a ten-gigabit port) to channel 5.

```
bstnA(cfg-channel[4])# no members 2/13 to 2/14
```
removes ports 2/13-2/14 from channel 4.

Related Commands

channel
shutdown (cfg-if-gig)
shutdown (cfg-if-ten-gig)
redundancy protocol (cfg-channel)
show interface **summary**

members (cfg-vlan)

Purpose

Use the members command to add a single port or a range of ports to the current VLAN.

Use no members to remove a port(s).

Mode

cfg-vlan

Security Role(s)

network-engineer or crypto-officer

Syntax

```
members slot/port [to slot/port ]
no members slot/port [to slot/port ]
```

slot/port (2/1-14 on ARX-4000; 1/1-4 or 2/1-2 on ARX-2500; 1/1-12 on ARX-2000; 1/1-8 on ARX-1500) is the first (or only) Ethernet port.

to *slot/port* (optional) is the last port in a range of ports. You cannot use this option on the ARX-2500 or ARX-1500, which allow only a single port to carry each VLAN.

Use the show interface summary command to locate the slot(s) for these ports.

Default(s)

None

Valid Platforms

ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines

IEEE 802.1Q defines VLANs.

This command adds ports to the VLAN with *tagging* disabled. A port with tagging disabled does not tag any outgoing frames with the VLAN ID (VID) for this VLAN. Use the tag command to add ports with tagging enabled, or to enable tagging for existing VLAN ports.

On the ARX-1500 and the ARX-2500, only a single port (interface gigabit or interface ten-gigabit) or channel can carry any given VLAN. That is, you can only select a single member port with this command. This includes VLAN 1; if multiple channels or ports default to VLAN 1, all but one of them must be disabled (with shutdown (cfg-if-gig), shutdown (cfg-if-ten-gig), or shutdown (cfg-channel)).

To assign a channel to carry the VLAN, use the vlan (cfg-channel) command. On the ARX-2500 or ARX-1500, this is the method for carrying a VLAN on multiple ports; aggregate the ports into a channel, then assign the desired VLAN to that channel.

Guidelines: Removing the Last Port from the Redundancy Link

In the following circumstances, the no members command causes the backup ARX peer to reboot:

- the interface vlan command establishes an in-band (VLAN) management address for this VLAN,
- the redundancy (cfg-if-vlan) command establishes the above management address as the local end of the redundancy link,
- redundancy is active between the ARX peers, and
- the no members command is removing the last port(s) from the VLAN.

The reboot does not disrupt any storage services, but the ARX peers cannot function as a redundant pair while the link is shut down. Additionally, a quorum-disk failure or disconnection would cause the active peer to reboot, too. If you proceed with shutting down the link, you should establish a new one as soon as possible: use the redundancy protocol, redundancy protocol (cfg-channel), or redundancy (cfg-if-vlan) command on another port, channel, or VLAN interface to establish a new redundant-pair link.

Samples

```
bstnA(cfg-vlan[1])# members 2/3 to 2/6
```

adds ports 2/3-2/6 to VLAN 1.

```
bstnA(cfg-vlan[5])# members 2/1
```

adds port 2/1 to VLAN 5.

```
bstnA(cfg-vlan[1])# no members 2/5 to 2/6
```

removes ports 2/5-2/6 from VLAN 1.

```
bstnA(cfg-vlan[7])# no members 2/11
```

removes port 2/11 from VLAN 7.

Related Commands

vlan
tag

priority (cfg-channel)

Purpose

LACP (IEEE 802.1AX-2008) defines a System Priority parameter for the stations at either end of a link-aggregation channel. From cfg-channel mode, use the priority command to set the System Priority for the current channel.

Use no priority to revert to the default System Priority.

Mode

cfg-channel

Security Role(s)

network-engineer or crypto-officer

Syntax

priority *number*
no priority

number (0-65536) is System Priority that you choose for this channel.

Default(s)

32,767

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

This command is only relevant in a channel with lacp passive enabled.

A lower number indicates higher priority. The station with the highest priority initiates all changes to member-port status in the channel. As an example of a change in member-port status: a port could change from a “standby” state to active use in the channel due to a configuration change.

Each channel can have a different priority setting.

You cannot run this command on the ARX-1500 or ARX-2500; the System Priority on those platforms is fixed at 65535 (lowest priority).

Samples

```
bstnA(cfg-channel[1])# priority 1
```

sets the System Priority to 1 (a very high priority) at channel 1.

```
bstnA(cfg-channel[10])# no priority
```

sets the default System Priority at channel 10.

Related Commands

channel
show channel
lacp passive

priority (cfg-stp)

Purpose

In a spanning tree topology, the bridge with the lowest *Bridge Priority* is elected as the spanning-tree root. From cfg-stp mode, use the priority command to set the Bridge Priority for the ARX.

Use no priority to revert to the default Bridge Priority.

Mode

cfg-stp

Security Role(s)

network-engineer or crypto-officer

Syntax

`priority` ***`number`***
`no priority`

number (0-61440) is Bridge Priority that you choose for the ARX. Use a multiple of 4096 (such as 0, 4096, 8192, or 12288).

Default(s)

61,440

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

The Spanning-Tree protocol and Bridge Priority are defined in IEEE 802.1D.

A lower number indicates higher priority. If all bridges have the same priority, the bridge with the lowest MAC address is elected as the root bridge.

Samples

```
bstnA(cfg-stp)# priority 0
```

sets the Bridge Priority to 0. The ARX will likely be elected as the root bridge.

```
bstnA(cfg-stp)# no priority
```

sets the default Bridge Priority.

Related Commands

spanning-tree

protocol (cfg-stp)

Purpose From cfg-stp mode, use the protocol command to choose the spanning-tree protocol: the original Spanning Tree Protocol (STP), or Rapid Spanning Tree (RST) protocol.

Use no protocol to revert to the default.

Mode cfg-stp

Security Role(s) network-engineer or crypto-officer

Syntax

```
protocol {dot1d | rst}
no protocol
```

dot1d | rst is a required choice:

- **dot1d** runs the original STP (IEEE 802.1D), ignoring RST-based BPDUs.
- **rst** runs RST from IEEE 802.1w, but is compatible with bridges that run the original STP (above).

Default(s) RST

Valid Platforms ARX-2000 and ARX-4000

Guidelines RST improves on the original STP implementation by more-rapidly converging on a new spanning-tree topology after a bridge or port failure. Conforming bridges (including the ARX) agree on the most-modern supported version of the protocol and use that version for all spanning-tree communication.

For an RST implementation, use the cfg-if-gig spanning-tree edgeport command to identify all the Edge Ports on the ARX.

Samples

```
bstnA(cfg-stp)# protocol dot1d
```

selects the earliest version of spanning-tree.

```
bstnA(cfg-stp)# no protocol
```

sets the protocol to the default, RST.

Related Commands spanning-tree
spanning-tree edgeport

redundancy protocol

Purpose

Use the redundancy protocol command to designate the current interface as one end of a redundant-pair link.

Use the no redundancy protocol command to remove support for a redundant-pair link. This causes the standby peer to reboot; see

Modes

cfg-if-gig

cfg-if-ten-gig

Security Role(s)

network-engineer or crypto-officer

Syntax

```
redundancy protocol
no redundancy protocol
```

Default(s)

None

Valid Platforms

ARX-2000 and ARX-4000

In cfg-if-ten-gig mode: ARX-4000 only.

Guidelines

At the layer-2 level, this establishes one end of the link between redundant peers. For best performance, a gigabit or ten-gigabit connection is strongly recommended; use the speed (cfg-if-gig) command to set the speed on a single-gigabit interface. We also recommend that the connection be direct (without any intervening bridges or routers), and that the switches are co-located. If the latency is low, an intervening Gigabit L2 switch is permissible.

Alternatively, you can configure a multi-port channel as the redundant-pair link. Use channel to create the channel, then use redundancy protocol (cfg-channel) to add member ports to it. (The ARX-1500 and ARX-2500 use different commands to establish the redundancy link over a channel, described in the documentation for redundancy protocol (cfg-channel).)

For cases where low latency between the peers is impossible, you may need to increase a timeout value when you set up the redundant pair. When you set up redundancy between the peers later, you can use the resilver-timeout command to increase this timeout value. You can also use the show redundancy metalog command to monitor the latency between the peers.

Guidelines: Shutting Down the Redundant-Pair Link

The no redundancy protocol command disables the redundant-pair link (see redundancy protocol), and therefore causes the standby peer to reboot if redundancy is enabled (enable (cfg-redundancy)). The reboot does not disrupt any storage services, but the ARX peers cannot function as a redundant pair while the link is shut down. Additionally, a quorum-disk failure or disconnection would cause the active peer to reboot, too. If you proceed with shutting down the link, you should establish a new one as soon as possible: use the redundancy protocol or redundancy protocol (cfg-channel) command on another port or channel to establish a new redundant-pair link.

The CLI prompts for confirmation before shutting down a redundant-pair link; enter **yes** to proceed with the interface shutdown and the reboot.

Guidelines: ARX-1500 and ARX-2500

The ARX-1500 and ARX-2500 use layer-3 (IP) networking software for exchanging heartbeats and important metalog data between the peers, so they require layer-3 configuration for their redundancy link. Instead of the layer-2 redundancy protocol command described here, you can use the following commands to set up a layer-3 redundancy link over the current interface:

- Establish a new VLAN for this link. Use the vlan command to create a new VLAN, then use members (cfg-vlan) to assign the current interface to that VLAN.
- Use the interface vlan to create a management-IP interface on the VLAN; this puts you into cfg-if-vlan mode.
  - From cfg-if-vlan mode, use the ip address (cfg-if-vlan) command to establish an in-band (VLAN) IP address. You later use this VLAN-management IP address to identify this ARX to its peer, as described below.
  - From the same mode, use redundancy (cfg-if-vlan) to designate the interface for exchanging metalog data and heartbeats.
  - From the same mode, use no shutdown (cfg-if-vlan) to enable the management interface.

Samples

```
prtlndA(cfg-if-gig[2/1])# redundancy protocol
```

uses interface 2/1 as one end of a redundant-pair link.

```
bstnA(cfg-if-gig[2/13])# no redundancy protocol
Removing redundancy protocol for this interface will cause the peer to
reboot.

Are you sure? [yes/no] yes
```

stops using interface 2/13 as one end of a redundant-pair link. If the redundant pair has already formed (see the enable (cfg-redundancy) documentation), this causes the peer to reboot.

```
bstnA(cfg-if-ten-gig[2/2])# redundancy protocol
```

uses a ten-gigabit interface, 2/2, as one end of a redundant-pair link.

Related Commands

interface gigabit
interface ten-gigabit
speed (cfg-if-gig)
redundancy
resilver-timeout
show redundancy metalog

redundancy protocol (cfg-channel)

Purpose

From cfg-channel mode, use the redundancy protocol command to add a single port or a range of ports to the current redundancy-link channel. This command performs two tasks at once: it adds ports to the channel, and prepares the channel for use as a redundant-pair link.

Use the no redundancy protocol command to remove redundancy-link support along with one or more ports.

Mode

cfg-channel

Security Role(s)

network-engineer or crypto-officer

Syntax

`redundancy protocol` ***`slot/port`*** **`[to`** ***`slot/port`*** **`]`**
`no redundancy protocol` ***`slot/port`*** **`[to`** ***`slot/port`*** **`]`**

slot/port (2/1-14 on ARX-4000; 1/1-12 on ARX-2000) is a single port or the first port in a range.

to ***slot/port*** (optional) is the last port in a range of ports.

Use the show interface summary command to locate the slot(s) for these ports.

Default(s)

None

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

A channel can have up to 8 ports. All ports must belong to the same VLAN(s) and they must run at the same speed.

Redundancy requires a reliable and fast channel for best performance. We recommend that you enable LACP on the channel to increase its reliability in high packet traffic: use lacp passive or lacp active to enable LACP, depending on your platform. For best performance, a gigabit (or higher-bandwidth) connection is strongly recommended. Use the cfg-if-gig speed (cfg-if-gig) command to set the speed on each port. We also recommend that the connection be direct (without any intervening bridges or routers), and that the switches are co-located. If the latency is low, an intervening Gigabit L2 switch is permissible.

For cases where low latency between the peers is impossible, you may need to increase a timeout value when you set up the redundant pair. When you set up redundancy between the peers later, you can use the resilver-timeout command to increase this timeout value. You can also use the show redundancy metalog command to monitor the latency between the peers.

A port cannot be a member of a channel if it has an “auto” speed.

Use the members (cfg-channel) command to add ports to a standard (non-redundancy-link) channel.

You cannot use the no redundancy protocol command on the last port in the channel; instead, use no channel to remove the entire channel.

Guidelines: ARX-1500 and ARX-2500

The ARX-1500 and ARX-2500 use layer-3 (IP) networking software for exchanging heartbeats and important metalog data between the peers, so they require layer-3 configuration for their redundancy link. Instead of the layer-2 redundancy protocol command described here, you can use the following commands to set up a layer-3 redundancy link over the current channel:

- Establish a new VLAN for this link. Use the vlan (cfg-channel) command to assign the channel to the VLAN.
- Use the interface vlan to create a management-IP interface on the VLAN; this puts you into cfg-if-vlan mode.
  - From cfg-if-vlan mode, use the ip address (cfg-if-vlan) command to establish an in-band (VLAN) IP address. You later use this VLAN-management IP address to identify this ARX to its peer, as described below.
  - From the same mode, use redundancy (cfg-if-vlan) to designate the interface for exchanging metalog data and heartbeats.
  - From the same mode, use no shutdown (cfg-if-vlan) to enable the management interface.

Samples

```
bstnA(cfg-channel[2])# redundancy protocol 2/1 to 2/2
```

adds ports 2/1-2/2 to channel 2, a channel to be used as a link between redundant peers.

```
prtlndA(cfg-channel[4])# no redundancy protocol 2/3
```

removes port 2/3 from channel 4.

Related Commands

channel
show interface **summary**
speed (cfg-if-gig)
members (cfg-channel)
redundancy
resilver-timeout
show redundancy metalog

show channel

Purpose A *channel* is a group of Ethernet ports aggregated into a single flow, as specified in IEEE 802.3ad. Use the show channel command to show the configuration of one channel.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show channel {summary | load-balance | channel-id [lacp]}`

summary | load-balance | *channel-id* is a required choice.

- **summary** shows a summary for all channels.
- **load-balance** shows the load-balancing algorithm for all channels.
- ***channel-id*** (1-8) identifies one channel; use this option to show detailed parameters for one channel.

  lacp (optional, if you choose a *channel-id*) displays the configuration and state of the Link Aggregation Control Protocol (LACP) on the given channel.

Valid Platforms ARX-1500, ARX-2000, ARX-2500, or ARX-4000

Guidelines This command shows the configuration and running state of a channel. Its various outputs are described below. For channel statistics (such as packet counts), use show channel ... stats.

Guidelines: Summary The show channel summary command outputs one line per channel. Each line contains the following fields:

Ch Id identifies the channel.

* indicates that the channel is used as a redundancy link with the switch's redundant peer (see redundancy protocol (cfg-channel)).

Admin State is set by the shutdown (cfg-channel) command.

Oper Status is Up if the channel has at least one operational port. This is Down if all of the member ports are down.

Speed is the speed of each port in the channel. Before ports can be aggregated into a channel, their speeds must match (see speed (cfg-if-gig) to set a port's speed).

Load-Balancing Algorithm shows the IP address(es) (source and/or destination) that are hashed to choose a port for an outbound packet. This is set by the load-balance command.

LACP is "Active," "Passive," or "Disabled," depending on the lacp active setting (on the ARX-1500 or ARX-2500), or the lacp passive setting (on other platforms that support channels).

Description is set by the description (cfg-channel) command.

Guidelines: Load-Balance The show channel load-balance command shows the load-balancing algorithm for each channel.

Guidelines: *channel-id* (Detailed)

The show channel *channel-id* command shows details about the given channel:

Channel Id identifies the channel.

Description is set by the description (cfg-channel) command.

Load Balancing Algorithm shows the IP address(es) or MAC address(es) (source and/or destination) that are hashed to choose a port for an outbound packet. This is set by the load-balance command.

LACP State is "Active," "Passive," or "Disabled." LACP (Link Aggregation Control Protocol) is a control protocol for dynamically adapting member usage to topology changes. On some platforms, you can enable LACP "Passive" mode with the lacp passive command; on other platforms, you can use the lacp active command to enable LACP "Active" mode.

LACP Rate is either "Short Timeout" or "Long Timeout." The rate should be the same at both ends of the channel, or the channel may periodically drop out of service. You can set this with the lacp rate command.

Members(Slot/Interface) lists the interfaces (ports) in this channel. Use the members (cfg-channel) command or the redundancy protocol (cfg-channel) command to add members to the channel.

Number of Members counts the ports from the above field.

Admin State is set by the shutdown (cfg-channel) command.

Channel Oper Status is Up if the channel has at least one operational port. This is Down if all of the member ports are down.

Trap Status is set by the trap shutdown command.

These spanning-tree fields only appear on systems that support the spanning-tree command.

- Spanning-Tree Forwarding State is "discard," "forward," or "disabled." The channel is typically in "discard" state when the Spanning-Tree Role is "alternate," and it is typically in "forward" state when the role is "designated" or "root."
- Spanning-Tree State is enabled if the channel participates in the spanning tree. You can use the spanning-tree shutdown command to disable spanning tree for the channel.
- Spanning-Tree Role is the channel's Port Role in the spanning tree: "root," "designated," or "alternate."
- Accept Frames is "All" or "Tagged Only." If this is the latter, the channel rejects all ingress frames unless they are tagged for one of the channel's VLANs. "Tagged Only" appears if (and only if) the channel is used in a link between redundant peers; see redundancy protocol (cfg-channel).

Total Vlans Configured counts all VLANs in which this channel participates.

A small table of VLANs appears next, with one row for each VLAN on this channel. The table contains two columns:

Members VLAN ID is the VLAN (if any) where this channel is a member. The channel does not *tag* any outgoing frames with the VLAN ID (VID) for this VLAN.

Tag VLAN ID lists one or more VLANs (if any) where this channel is a tagging member; that is, this channel *tags* its outgoing frames with the VLAN ID (VID) for the destination VLAN.

Guidelines: *channel-id* (Detailed), Cont.

The next table contains the status of each member port, one per row. Slot/Port identifies each member port, and Link Status its status (up or down).

Spanning-Tree Statistics is a table of counters for the Bridge Protocol Data Units (BPDUs) transmitted and received. The counters are shown for two versions of spanning-tree BPDU: original spanning tree (STP) or rapid spanning-tree (RST).

Guidelines: lacp

If you use the optional lacp argument in the command, the output shows LACP parameters and status for the chosen channel.

Three summary fields appear at the top of the output:

Channel ID identifies this channel

LACP is "Active," "Passive," or "Disabled," depending on the lacp active setting (on the ARX-1500 or ARX-2500), or the lacp passive setting (on other platforms that support channels).

Time since last state change does not appear on the ARX-1500 or ARX-2500. This shows the time that has passed since the last change in channel membership or status.

This is followed by a table of LACP Channel Parameters. This table shows the channel-level configuration for LACP. It is divided into two columns: Local (for the ARX end of the channel) and Peer (for the remote end of the channel):

Admin Key is the numeric key for the channel that was set by its administrative configuration. A channel's *key* is a number used by LACP software to identify the relevant configuration parameters of the channel. This is the same as the Oper Key, below, until or unless a configuration and/or topology change triggers a change in member-port usage. This field only appears in the Local column.

Oper Key is the numeric key that is currently in use for the channel. This key appears in both columns. The peer's operational key is the one that was reported in the most-recent Link Aggregation Control Protocol Data Unit (*LACPDU*) from the peer.

System Priority does not appear on the ARX-1500 or ARX-2500. This shows the System Priority of the ARX and its peer. A lower number is considered a higher priority. The system with the higher priority initiates all port-membership changes in the channel, such as putting a port in "standby" status due to a configuration change. You can use the priority (cfg-channel) command to set the priority for the ARX end of the channel.

System ID is the MAC address used to identify the ARX and its peer.

The final section of the output, LACP Port Parameters, contains one table per channel member. Each table describes one port with its port-level LACP parameters and status. As above, these tables contain one Local column (for the ARX port) and a Peer column (for the corresponding port at the remote end of the channel). On the ARX-1500 and ARX-2500, this table contains the following fields:

Slot/Port identifies the ARX port. This information does not appear for the Peer port.

Oper Key is the key that is currently in use for this port. This number is a code that is only meaningful to the LACP software. It appears for both the local port and the peer port.

Link Status is either "up" or "down." If this link is down, it is not being used for the channel's traffic.

Link Failure Count shows the number of transmissions failures on this link, if any.

Guidelines: LACP Output (Cont.)

On all other platforms, the LACP Port Parameters table contains the following fields:

Slot/Port identifies the ARX port. This information does not appear for the Peer port.

Admin Status shows the administratively-set status of the member port. This only appears for an ARX port. Each port may have one or more of the following codes to signify its status:

- A or P - A (Active) means that the port is actively running LACP, and P (Passive) indicates that the port is sending LACPDUs but has not yet received proper responses from the peer.
- T or L - indicates the timeout between LACPDU transmissions. T is a short timeout (typically, 1 second) and L is a long timeout (typically 30 seconds).
- a - means that the port is eligible for Aggregation, or active use in the channel. Ports without this flag cannot be used in the channel; check the configuration at both ends for possible differences that make them incompatible.
- S - shows that the LACP process considers this port "in Sync" with its usage in the channel. If this flag is missing, it may indicate that the LACP software is in the process of changing the port from a "standby" state to an "active" one, or from active to standby.
- C - indicates that the port is Collecting packets. That is, it can accept incoming traffic.
- D - means that the port is Distributing packets. That is, it can send outbound packets to its peer.
- d - indicates that the port is using Default information for its peer's operational key. This means that the peer's configuration information on the ARX does not contradict the latest learned information from LACPDUs.
- E - indicates that the latest LACPDU has expired, and the LACP process is waiting for the next one.

Oper State is the current status of the port. This appears for both the Local port and its Peer. This has the same possible values as the Admin Status, described above.

Admin Key is a numeric key used by the LACP software as a code to represent the port's capabilities. This key is based on administrative (CLI) settings. The LACP software may choose a different operational key (described below) for the port based on L2-topology changes or configuration changes at the peer port. This only appears for the ARX port, where administrative parameters are known.

Oper Key is the key that is currently in use for this port. This number is a code that is only meaningful to the LACP software. It appears for both the local port and the peer port.

Port Priority is a number that represents the port's eligibility for use in the channel. A lower priority number represents a higher priority. If ports are excluded from active use in the channel, the LACP software prefers low-priority ports (that is, ports with higher-priority numbers). This appears for both the local and peer ports, and it may be different at both ends.

Samples

`stoweA# show channel summary`

shows a one-line summary for each channel. See Figure 9.1 for sample output.

`prtlndA# show channel load-balance`

shows the load-balancing algorithm for all channels on the "prtlndA" switch. See Figure 9.2 on page 9-48 for sample output.

`prtlndA# show channel 1`

shows the configuration for channel 1 on the same switch. See Figure 9.3 on page 9-48 for sample output.

`bstnA# show channel 1 lacp`

shows the LACP configuration for channel 1 on the "bstnA" switch. See Figure 9.4 on page 9-49 for sample output.

Related Commands

description (cfg-channel)
members (cfg-channel)
redundancy protocol (cfg-channel)
shutdown (cfg-channel)
lacp passive
lacp rate
priority (cfg-channel)
trap shutdown

Figure 9.1 *Sample Output: show channel summary*

```
stoweA# show channel summary

Ch  Admin    Oper   Speed   Load-Balance LACP     Description
Id  State    Status         Algorithm
--- -------- ------ ------- ------------ -------- ----------------------
1   Enabled  Up     10 Gb/s src-dst-ip   Passive  default

stoweA#
```

Figure 9.2 *Sample Output: show channel load-balance*

```
prtlndA# show channel load-balance

Channel Id  Load-Balance Algorithm
----------  -------------------------
1           Source and Destination Ip
2           Source and Destination Ip

prtlndA#
```

Figure 9.3 *Sample Output: show channel*

```
prtlndA# show channel 1

Channel ID                    : 1
Description                   : default
Load Balancing Algorithm      : Source and Destination Ip
LACP State                    : Passive
LACP Rate                     : Short Timeout
Members(Slot/Interface)       : 2/3,2/4
```

```
Number of Members             : 2
Admin State                   : Enabled
Channel Oper Status           : Up
Trap Status                   : Disabled
Spanning-Tree Forwarding State: Manual Forwarding
Spanning-Tree State           : Disabled
Spanning-Tree Role            : Disabled
Accept Frames                 : All
Total VLANs Configured        : 2


Members  Tag
VLAN ID  VLAN ID
-----------------

1         N/A
405       N/A



Slot/Port    Link Status
-------------------------
2/3            Up
2/4            Up



Spanning Tree Statistics
------------------------

STP BPDU Transmitted    0
STP BPDU Received       0
RST BPDU Transmitted    0
RST BPDU Received       7975
MST BPDU Transmitted    0
MST BPDU Received       0
prtlndA#
```

Figure 9.4 *Sample Output: show channel 1 lacp*

```
bstnA# show channel 1 lacp

Channel ID                    : 1
LACP                          : Passive
Time since last state change  : 13:32:38 05/20/2008

LACP Channel Parameters :

         Local                                  Peer
------------------------------- ---------------------------------------
Admin Key:       86
Oper Key:        86                 Oper Key:        86
System Priority: 100                System Priority: 32768
System ID:       00:0a:49:17:70:40  System ID:       00:0a:49:17:72:40

LACP Port Parameters:

         Local                                  Peer
------------------------------- ---------------------------------------
Slot/Port:     2/7
Admin Status:* A,T,a
```

```
Oper State:  * A,T,a,S,C,D       Oper State:  * A,T,a,S,C,D
Admin Key:     86
Oper Key:      86                Oper Key:      86
Port Priority: 100               Port Priority: 0

Slot/Port:     2/8
Admin Status:* A,T,a
Oper State:  * A,T,a,S,C,D       Oper State:  * A,T,a,S,C,D
Admin Key:     86
Oper Key:      86                Oper Key:      86
Port Priority: 100               Port Priority: 0

Slot/Port:     2/9
Admin Status:* A,T,a
Oper State:  * A,T,a,S,C,D       Oper State:  * A,T,a,S,C,D
Admin Key:     86
Oper Key:      86                Oper Key:      86
Port Priority: 100               Port Priority: 0

Slot/Port:     2/10
Admin Status:* A,T,a
Oper State:  * A,T,a,S,C,D       Oper State:  * A,T,a,S,C,D
Admin Key:     86
Oper Key:      86                Oper Key:      86
Port Priority: 100               Port Priority: 0


* A - Active,        a - Aggregating, C - Collecting,   D - Distributing,
  d - Defaulted,     E - Expired,     L - Long Timeout, P - Passive,
  T - Short Timeout, S - Synchronizing
```

show channel ... stats

Purpose

A *channel* is a group of Ethernet ports aggregated into a single flow, as specified in IEEE 802.3ad. Use the show channel ... stats command to show traffic statistics for one channel.

Mode

(any)

Security Role(s)

crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

`show channel channel-id [lacp] stats`

channel-id (1-8) identifies the channel to show.

lacp (optional) focuses the output on the Link Aggregation Control Protocol (LACP) statistics for the channel.

stats is a required keyword.

Valid Platforms

ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines

This command shows traffic statistics or Link Aggregation Control Protocol (LACP) statistics for a given channel. The sections below describe the two outputs from this command.

For configuration information about the channel, use show channel. To clear the channel statistics in the default output, use clear counters channel. To clear the LACP statistics, use clear counters lacp.

Guidelines: Traffic Statistics

If you omit the optional lacp keyword, the output focuses on frame counts for the overall channel.

Channel Id identifies the channel.

The statistics are divided into three tables:

- Basic Ingress and Egress statistics (primarily packet counts and error counts.)
- Statistics for MIB values from RFC 1493.
- Ethernet statistics.

Guidelines: LACP Statistics

If you use the lacp keyword, a table appears with LACP statistics. To enable LACP on a channel, you use the lacp active command on the ARX-1500 or ARX-2500, or the lacp passive command on other platforms that support channels.

Each table has a separate row for every member port in the channel, with the following columns:

S/P identifies the channel member in *slot*/*port* format.

LACP Packets are the numbers of Link Aggregation Control Protocol Data Units (*LACPDUs*) transmitted from and received on this member port.

Marker Response counts the LACP-marker frames transmitted and received. The LACP software sometimes injects *marker* frames to find the ends of one or more frame "conversations." A *marker response* frame from the peer indicates that the conversation(s) is/are finished. Once the LACP software receives this marker, it can migrate future conversations to another link in the channel. For details on the Marker Protocol, see EEE802.3ad, Section 43.5.

Illegal is the number of illegal Slow-Protocol PDUs (see IEEE802.3ad, Section 43B.4) received on this member port.

Unknown is the number of unknown Slow-Protocol PDUs received on this member port.

Samples

`prtlndA# show channel 1 stats`

shows the statistics for channel 1. See Figure 9.5 for sample output.

`bstnB# show channel 1 lacp stats`

shows the LACP statistics for channel 1 on a different switch. See Figure 9.6 on page 9-53 for sample output.

Related Commands

show channel
clear counters lacp
clear counters channel

Figure 9.5 *Sample Output: show channel 1 stats*

```
prtlndA# show channel 1 stats


Channel Id  : 1
-------------------------------------------------------------------
                               Ingress                  Egress
-------------------------------------------------------------------

Octets                      1564310788               1906062907
Total Frames                   3246593                  3149307
Unicast Frames                 3229634                  3138606
Multicast Frames                  9709                        0
Broadcast Frames                  7250                    10701
PAUSE Frames                         0                        0
If Discards                          0                        0
If Errors                            0                        0
Int Mac Errors                       0                        0
If Unknown Protocol                  0
Alignment Errors                     0
CRC Errors                           0
Single Collisions                                             0
Multiple Collisions                                           0
```

```
Late Collisions                                          0
Excessive Collisions                                     0
Frames Too Long                      0

------------------------------------------------------------------
                               RFC 1493
------------------------------------------------------------------

TpPortDelayExceed                    0
TpPortMTUExceed                      0
TpPortInDis                          0

------------------------------------------------------------------
                               Ether Stats
------------------------------------------------------------------

Packet Size:64                 1229163
            65-127             1988693
            128-255             947001
            256-511             193302
            512-1023             49683
            1024-1518          1988058
            1519-2047                0
            2048-4095                0
            4096-9216                0
Multicast Packets                 9709
Broadcast Packets                17951
Total Octets                3470373695
Good Oversize Frames                 0
Drop Events                          0
Total Discards                       0
Undersize Packets                    0
Fragments                            0
Jabbers                              0
Total Collisions                     0
CRC+Alignment Errors                 0
prtlndA#
```

Figure 9.6 *Sample Output: show channel 1 lacp stats*

```
bstnB# show channel 1 lacp stats

LACP Statistics:

S/P       LACP  Packets      Marker Response     Illegal    Unknown
        Tx         Rx           Tx       Rx         Rx         Rx
---- ---------- ---------- ---------- ---------- ---------- ----------
2/7  1336       1315       0          0          0          0
2/8  1339       1317       0          0          0          0
2/9  1334       1313       0          0          0          0
2/10 1338       1318       0          0          0          0
```

show interface gigabit

Purpose Use the show interface gigabit command to show the configuration of one Gigabit interface. Add the stats keyword to the end of the command to show the interface's traffic statistics.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show interface gigabit`** ***`slot/port`*** **`[stats]`**

slot/port (2/3-14 on ARX-4000; 1/1-4 on ARX-2500; 1/1-12 on ARX-2000; 1/1-8 on ARX-1500; 1/2 on ARX-500; 1/1 on ARX-VE) specifies the interface.

stats (optional) displays statistics for this interface.

Use the show interface summary command to locate slots and ports.

Guidelines The default command shows a table of configuration settings.

The stats output is a table of counters, separated into Ingress and Egress counts. Use the clear counters gigabit command to clear and restart the statistics count.

PAUSE Frames are notices to control traffic flow: the flowcontrol command determines whether the interface sends or receives these.

These outputs are abbreviated for the ARX-VE platform. For example, on an ARX-VE the output does not include any references to flow control or storm control, which are managed at the hypervisor's VNIC.

Samples

`bstnA>` **`show interface gigabit 2/6`**

shows the configuration for slot 2, port 6. See Figure 9.7 for sample output.

`bstnA>` **`show interface gigabit 2/6 stats`**

shows the statistics for the same slot and port. See Figure 9.8 on page 9-55 for sample output.

`stoweA>` **`show interface gigabit 1/2`**

shows the configuration for port 1/2 on an ARX-2500. See Figure 9.9 on page 9-56 for sample output.

`stkbrgA#` **`show interface gigabit 1/1 stats`**

shows the statistics for the only interface on an ARX-VE. As mentioned above, the statistics are abbreviated on this chassis type. See Figure 9.9 on page 9-56 for sample output.

Related Commands clear counters gigabit
show interface **`summary`**

Figure 9.7 *Sample Output: show interface gigabit*

```
bstnA> show interface gigabit 2/6

Slot                      2
Interface                 6
Description               Default
Type                      Copper
Mode                      Normal
Admin State               Enabled
Link Status               Up
Speed                     1 Gb/s
Duplex                    Full
Auto Negotiation(Admin)   Enabled
Flow Control(Admin)
             Receive      Off
             Send         Off
MAC Address               00:0a:49:17:92:37
Storm Control:Broadcast   1000 packets/sec
              Multicast   1000 packets/sec
              Unknown DA 1000 packets/sec
Port VLAN ID

Accept Frames             Admit All
```

Figure 9.8 *Sample Output: show interface gigabit stats*

```
bstnA> show interface gigabit 2/6 stats

Slot                           2
Interface Id                   6
-------------------------------------------------------------------
                                Ingress                  Egress
-------------------------------------------------------------------

Octets                       730594224                648324230
Total Frames                   1045187                  1153090
Unicast Frames                 1037623                  1152780
Multicast Frames                  2479                        0
Broadcast Frames                  25 5                      310
PAUSE Frames                         0                        0
CRC Errors                           0
Total Discards                       0
Alignment Errors                     0
Single Collisions                                             0
Multiple Collisions                                           0
Late Collisions                                               0
Excessive Collisions                                          0

-------------------------------------------------------------------
                                Ether Stats
-------------------------------------------------------------------

Packet Size:64                  108120
            65-127               26871
            128-255            1224687
            256-511              42220
            512-1023            104267
            1024-1518           692112
            > 1519                   0
Total Collisions                     0
```

Figure 9.9 *Sample Output: show interface gigabit on ARX-2500*

```
stoweA> show interface gigabit 1/2

Slot                    1
Interface               2
Description             Default
Type                    Copper
Mode                    Normal
Admin State             Enabled
Link Status             Up
Speed                   1 Gb/s
Duplex                  Full
Auto Negotiation(Admin) Enabled
Flow Control(Admin)
             Receive    Off
             Send       Off
MAC Address             00:90:fb:33:2e:1d

Accept Frames           Admit All
```

Figure 9.10 *Sample Output: show interface gigabit stats on ARX-VE*

```
stkbrgA# show interface gigabit 1/1 stats

Slot                          1
Interface Id                  1
-------------------------------------------------------------------
                                Ingress                  Egress
-------------------------------------------------------------------

Octets                       48624230                34815281
Total Frames                   503544                  302010
Unicast Frames                      0                       0
Multicast Frames                    0                       0
Broadcast Frames                    0                       0
CRC Errors                          0
Total Discards                      0
Alignment Errors                    0
Single Collisions                                           0
Multiple Collisions                                         0
Late Collisions                                             0
Excessive Collisions                                        0

-------------------------------------------------------------------
                               Ether Stats
-------------------------------------------------------------------

Total Collisions                    0
```

show interface ten-gigabit

Purpose Use the show interface ten-gigabit command to show the configuration of one ten-Gigabit interface. Add the stats keyword to the end of the command to show the interface's traffic statistics.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show interface ten-gigabit`** ***`slot/port`*** **`[stats]`**

slot/port (2/1-2) specifies the interface. Use the show interface summary command to locate all slots and ports.

stats (optional) displays statistics for this interface.

Valid Platforms ARX-2500 and ARX-4000 only

Guidelines The default command shows a table of configuration settings.

The stats output is a table of counters, separated into Ingress and Egress counts. Use the clear counters ten-gigabit command to clear and restart the statistics count.

PAUSE Frames are notices to control traffic flow: the flowcontrol command determines whether the interface sends or receives these.

Samples `bstnA(cfg)#` **`show interface ten-gigabit 2/2`**

shows the configuration for the ten-gigabit interface in slot 2, port 2. See Figure 9.11 for sample output.

`bstnA(cfg)#` **`show interface ten-gigabit 2/1 stats`**

shows the statistics for the interface at 2/1. See Figure 9.12 for sample output.

Related Commands clear counters ten-gigabit
show interface **`summary`**

Figure 9.11 *Sample Output: show interface ten-gigabit*

```
bstnA(cfg)# show interface ten-gigabit 2/2

Slot                    2
Interface               2
Description             Default
Type                    10GBASE-SR X2
Mode                    Normal
Admin State             Enabled
Link Status             Up
Speed                   10 Gb/s
Duplex                  Full
Auto Negotiation(Admin) Disabled
Flow Control(Admin)
             Receive    Off
             Send       Off
MAC Address             00:0a:49:17:92:33
Storm Control:Broadcast 1000 packets/sec
```

```
                                    Multicast  1000 packets/sec
                                    Unknown DA 1000 packets/sec
                          Port VLAN ID             25

                          Accept Frames            Admit All
```

Figure 9.12 *Sample Output: show interface ten-gigabit stats*

```
bstnA(cfg)# show interface ten-gigabit 2/1 stats

Slot                          2
Interface Id                  1
-------------------------------------------------------------------
                               Ingress                 Egress
-------------------------------------------------------------------

Octets                                0                       0
Total Frames                          0                       0
Unicast Frames                        0                       0
Multicast Frames                      0                       0
Broadcast Frames                      0                       0
PAUSE Frames                          0                       0
CRC Errors                            0
Total Discards                        0
Alignment Errors                      0
Single Collisions                                             0
Multiple Collisions                                           0
Late Collisions                                               0
Excessive Collisions                                          0

-------------------------------------------------------------------
                                Ether Stats
-------------------------------------------------------------------

Packet Size:64                        0
            65-127                    0
            128-255                   0
            256-511                   0
            512-1023                  0
            1024-1518                 0
            > 1519                    0
Total Collisions                      0
```

show load-balancing

Purpose A *channel* is a group of Ethernet ports aggregated into a single flow, as specified in IEEE 802.3ad. Each channel uses a hash algorithm to balance the traffic load between its member ports; for each packet, the hash uses some combination of the source and destination IPs to choose a port. Use show load-balancing to show which port is used for a given source and destination IP.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show load-balancing source-ip`** ***`src`*** **`destination-ip`** ***`dest`*** **`channel`** ***`chnl-id`***

src is the source-IP address.

dest is the destination-IP address.

chnl-id (1-8) identifies the channel.

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines Use the load-balance command to change the algorithm for load-balancing a channel. Use this command to test the channel's current hash.

The output shows the Slot Id and Interface (port) that would be chosen for the source and destination IP that you provided.

Sample `prtlndA#` **`show load-balancing source-ip 172.16.100.98 destination-ip 192.168.25.23 channel 1`**

shows the port that channel 1 uses for a packet with the given source and destination IPs. See Figure 9.13 for sample output.

Related Commands load-balance

Figure 9.13 *Sample Output: show load-balancing*

```
prtlndA# show load-balancing source-ip 172.16.100.98 destination-ip 192.168.25.23 channel 1

Report for source-ip 172.16.100.98 and destination Ip 192.168.25.23
Channel Id :1
Slot Id    :4
Interface  :2
```

show mac-address-table

Purpose Use the show mac-address-table command for a list of MAC addresses used by the ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show mac-address-table`

Valid Platforms ARX-2000 and ARX-4000

Guidelines Slot and

Port show the port where the MAC address is used.

MAC Address is address.

VLAN ID is the VLAN for this MAC, if any.

Channel ID is the 802.3ad channel for this MAC, if any.

Mode is the method by which the MAC was added to the table:

- Learned is an address learned from a neighboring bridge.
- Management is one of two addresses:
  - an address associated with the out-of-band interface (refer to the documentation for interface mgmt), or
  - the management address advertised through the spanning-tree protocol (spanning-tree).
- Inband is associated with an in-band (VLAN) management interface, created with the interface vlan command.
- Self is an internally-assigned address.

Related Commands show chassis
interface mgmt
interface vlan

Figure 9.14 *Sample Output: show mac-address-table*

```
bstnA> show mac-address-table

Slot Port  MAC Address         VLAN ID    Channel ID    Mode
-----------------------------------------------------------------
         00:0a:49:17:78:ff        25                    Inband
         00:0a:49:17:78:fe        25                    Inband
         00:0a:49:17:78:40         1                    Management
   2   5 00:01:e8:5e:ea:1f        25                    Learned
   2   5 00:0a:49:17:74:c0        25                    Learned
   2   5 00:0a:49:17:74:ff        25                    Learned
   2   5 00:0a:49:17:7c:c0        25                    Learned
   2   5 00:0a:49:17:7c:ff        25                    Learned
   2   5 00:0a:49:17:80:c0        25                    Learned
   2   5 00:0a:49:17:80:ff        25                    Learned
   2   5 00:0a:49:17:81:c0        25                    Learned
   2   5 00:0a:49:17:81:ff        25                    Learned
```

```
   2   5 00:0a:49:17:86:c0       25                   Learned
   2   5 00:0a:49:17:86:ff       25                   Learned
   2   5 00:0a:49:17:8c:c0       25                   Learned
   2   5 00:0a:49:17:8c:ff       25                   Learned
   2   5 00:0a:49:17:92:ff       25                   Learned
   2   5 00:0a:49:17:a1:ff       25                   Learned
   2   5 00:0a:49:17:aa:ff       25                   Learned
   2   5 00:0a:49:17:af:c0       25                   Learned
...
   2   6 00:0a:49:17:cb:fe       25                   Learned
   2   6 00:0d:ec:d0:10:1b       25                   Learned
   2   6 00:1e:2a:3b:9f:a7       25                   Learned
```

show mac-address-table summary

Purpose Use the show mac-address-table summary command for a high-level view of the MAC-address table.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show mac-address-table summary`**

Valid Platforms ARX-2000 and ARX-4000

Guidelines Mac Address Learning Mode is always IVL.

MAC Address High Count and:

Active MAC Addresses in FDB are counts of addresses in the Forwarding DataBase (FDB), the table that holds all MAC addresses.

Maximum MAC Address Supported is fixed at 12,000.

Configured Aging Time is the maximum time a learned MAC address is kept in the FDB without any updates; if a MAC address is not re-learned for this many seconds, it is deleted. Use the mac-address aging-time command to set the aging time.

Use the show mac-address-table command to see all the MAC addresses in the FDB.

Related Commands mac-address aging-time
show mac-address-table

Figure 9.15 *Sample Output: show mac-address-table summary*

```
bstnA> show mac-address-table summary

Mac Address Learning Mode      IVL (Independent Vlan Learning)
MAC Address High Count           183
Active MAC Addresses in FDB      183
Maximum MAC Addresses Supported  12000
Configured Aging Time            300 secs
```

show redundancy network

Purpose Use the show redundancy network command to show the layer-2 status of the redundant-pair link.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show redundancy network`

Valid Platforms ARX-2000 and ARX-4000

Guidelines This command shows an error unless at least one port has the redundancy protocol setting, or one channel has ports added with the redundancy protocol (cfg-channel) command.

Network is "Client," "Server," "Private," or "Metalog." The redundant-pair link carries the Private and Metalog networks between the redundant peers.

VLAN is the VLAN number for this port. This does not appear on the ARX-500.

Port(s) are in *slot/port* format (for example, 2/9).

Admin State is "Enabled" or "Disabled." You can set this with the no shutdown (cfg-if-gig) or the no shutdown (cfg-if-ten-gig) command.

Link Status is "Up" or "Down:" this is the link's operational state.

Spanning-Tree Status is "Discard," "Learning," "Forward," "Disabled," "Manual Forwarding," or "Not Participating." This is the port's current role in the spanning tree.

The Link Transitions table tracks any and all state/status transitions for redundancy-related ports:

Count is the total number of transitions. Use the clear counters redundancy network command to clear this counter.

Last is the date and time of the last transition.

Reason explains the nature and cause of the last transition.

Last Cleared is the last time someone used the clear counters redundancy network command.

Related Commands clear counters redundancy network

Figure 9.16 *Sample Output: show redundancy network*

```
prtlndA# show redundancy network


Network  VLAN  Port(s)   Admin    Link   Spanning-Tree
                         State    Status Status
-------- ----- --------- -------- ------ -----------------
External 1     1/1       Enabled  Down   Disabled
External 1     1/2       Enabled  Down   Disabled
External 1     1/3       Enabled  Down   Disabled
External 1     1/4       Enabled  Down   Disabled
External 1     1/5       Enabled  Up     Manual Forwarding
```

```
External 1     1/6       Enabled  Up     Manual Forwarding
External 1     1/7       Disabled Down   Disabled
External 1     1/8       Disabled Down   Disabled
External 1     1/9       Disabled Down   Disabled
External 1     1/10      Disabled Down   Disabled
External 1     1/11      Disabled Down   Disabled
External 74    1/5       Enabled  Up     Manual Forwarding
External 74    1/6       Enabled  Up     Manual Forwarding
Private  1008  1/12      Enabled  Up     Manual Forwarding
Metalog  1009  1/12      Enabled  Up     Manual Forwarding


Link Transitions:
  Count:          3
  Last:           07:38:58 03/05/2010
  Reason:         Port 1/5 link up
  Last Cleared:   Never
```

show spanning-tree detailed

Purpose Use the show spanning-tree detailed command for a detailed view of the spanning-tree configuration.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show spanning-tree detailed`**

Valid Platforms ARX-2000 and ARX-4000

Guidelines This command shows detailed spanning-tree configuration parameters along with some BPDU traffic statistics.

Use the show spanning-tree summary command to show a summary of this data.

Related Commands spanning-tree
show spanning-tree summary

Figure 9.17 *Sample Output: show spanning-tree detailed*

```
bstnA(cfg)# show spanning-tree detailed

Bridge is Executing the IEEE compatible IEEE_Dot1d Spanning Tree protocol

Switch STP Admin State     Disabled
Bridge Priority            61440
Bridge Address             00:0a:49:17:78:40
Bridge Max Age             20 sec
Bridge Hello Time          2 sec
Bridge Forward Delay       15 sec
Bridge Hold Time           3 sec

Designated Root: Priority 240
                 Address  00:0a:49:17:78:40
Root Path Cost             0
Root Port Max Age          20
Root Port Fwd Delay        15
Time Since Topology Change 2719 sec
Topology Change Count      0
Topology Change            0


Slot/     Spanning Tree                  Port     BPDU Packets
Port    Forwarding    Admin    Path Cost Type    Rx          Tx
---- --------------- -------- ---------- ---- ---------- -----------
2/5  Manual Forward  Disabled 20000      gbe  0          0
2/6  Manual Forward  Disabled 20000      gbe  0          0
```

show spanning-tree interface

Purpose Use the show spanning-tree interface command to show the spanning-tree configuration for a particular port.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show spanning-tree interface`** ***`slot/port`***

slot/port (2/1-14 on ARX-4000 or 1/1-12 on ARX-2000) is the slot and port number. Use the show interface summary command to show all ports in all slots.

Valid Platforms ARX-2000 and ARX-4000

Guidelines This command shows the spanning-tree configuration parameters, the port's role in the spanning-tree topology, and some port statistics.

Related Commands spanning-tree

Figure 9.18 *Sample Output: show spanning-tree interface*

```
bstnA(cfg)# show spanning-tree interface 2/3

Bridge is Executing the IEEE compatible Spanning Tree protocol

Slot                    2
Interface               3
Port: SNMP ID           32785
      STP  ID           17
      Priority          128
      Forwarding State  Disabled
      STP State         Disabled
      Role              N/A
      Path cost         20000
Designated: SNMP Port ID 0
            STP Port ID  0
            Priority     0
            Address      00:0a:49:17:78:40
Edge Port: Admin Status       Configured
           Operational Status Operational
Point to Point Mac Status     Auto
Topology Change Ack     0
Port Up Time            2719 sec
STP BPDU Transmitted    0
STP BPDU Received       0
RST BPDU Transmitted    0
RST BPDU Received       0
MST BPDU Transmitted    0
MST BPDU Received       0
```

show spanning-tree summary

Purpose Use the show spanning-tree summary command to show high-level information about the spanning-tree configuration.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show spanning-tree summary`**

Valid Platforms ARX-2000 and ARX-4000

Guidelines Use the show spanning-tree detailed command to show details.

Sample

```
bstnA(cfg)# show spanning-tree summary

Spanning Tree Admin State is Enabled
Configuration: Protocol          : IEEE_Dot1w
               Revision Level    : 0
               Format Selector   : 0
Default      : Name              : 00-0A-49-00-13-02
Total MST Instances Created      : 0
```

Related Commands

spanning-tree
show spanning-tree detailed

show vlan

Purpose Use the show vlan command to show the configuration of one VLAN.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show vlan`** ***`vlanId`***

vlanId (1-4009) identifies the VLAN to display.

Guidelines Information from the show vlan command includes:

Vlan Id identifies the VLAN.

Description is set by the description (cfg-vlan) command.

Frame/MTU is the Ethernet packet size: Use jumbo mtu to change this.

Members S/P lists the ports in this VLAN. Use the members (cfg-vlan) or tag command to add members to the VLAN.

Non-Members (Blocked) S/P lists any ports identified as non-members with the no members command.

Tag S/P lists the ports set to *tag* outgoing packets with a VLAN ID. Use the tag command to enable tagging for one or more ports.

Use the show vlan summary command to list all configured VLANs.

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Samples

```
bstnA# show vlan 1
```

shows the configuration for VLAN 1. Figure 9.19 shows sample output.

```
stoweA# show vlan 1
```

shows the configuration for VLAN 1 on an ARX-2500 named "stoweA." Figure 9.20 on page 9-69 shows sample output.

Related Commands

show vlan summary
description (cfg-vlan)
members (cfg-vlan)
tag

Figure 9.19 *Sample Output: show vlan*

```
bstnA# show vlan 1

Vlan Id     : 1
Description : Default VLAN.
Frame/MTU   : Standard

Members  Non-Members   Tag
         (Blocked)
  S/P       S/P        S/P
```

```
----------------------------
2/1       N/A         N/A
2/2       N/A         N/A
2/3       N/A         N/A
2/4       N/A         N/A
2/5       N/A         N/A
2/6       N/A         N/A
2/7       N/A         N/A
2/8       N/A         N/A
2/9       N/A         N/A
2/10      N/A         N/A
2/11      N/A         N/A
2/12      N/A         N/A
2/13      N/A         N/A
2/14      N/A         N/A
```

Figure 9.20 *Sample Output: show vlan (ARX-2500)*

```
stoweA# show vlan 1

Vlan Id     : 1
Description : Default VLAN.
Frame/MTU   : Standard

Members  Non-Members   Tag
         (Blocked)
  S/P       S/P        S/P
----------------------------
1/1       N/A         N/A
1/2       N/A         N/A
1/3       N/A         N/A
1/4       N/A         N/A
```

show vlan summary

Purpose Use the show vlan summary command to list all configured VLANs.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show vlan summary`**

Valid Platforms any *except* ARX-VE

Guidelines The sample output shows a table of configured VLANs:

Vlan ID identifies each VLAN.

Usage is: 'External' for VLANs used outside the box, 'Internal' for private VLANs inside the box, 'Private' for internal communications and communications with redundant peers, and 'Metalog' for metadata exchange between redundant peers.

Channel is set by assigning a channel to this VLAN (with vlan (cfg-channel) or vlan-tag). This is "N/A" if not set.

Frame/MTU is the packet size, set by jumbo mtu.

Description is set by the description (cfg-vlan) command.

Use the show vlan command to show details about a VLAN.

Sample

```
bstnA# show vlan summary

Total VLANs Configured: 5

VLAN ID      Usage     Channel      Frame/MTU     Description
-------------------------------------------------------------
1           External    N/A         Standard      Default VLAN.
25          External    N/A         Standard      personnel dept.
1010        Private     N/A         Standard      Private Subnet VLAN.
1011        Metalog     N/A         Standard      Private Metalog VLAN.
```

Related Commands description (cfg-vlan)
show vlan

shutdown (cfg-channel)

Purpose

A *channel* is a group of aggregated Ethernet ports (IEEE 802.3ad). From cfg-channel mode, use the shutdown command to stop traffic on the current channel.

Use no shutdown to restart traffic on the channel.

Mode

cfg-channel

Security Role(s)

network-engineer or crypto-officer

Syntax

```
shutdown
no shutdown
```

Default(s)

no shutdown: link aggregation is enabled on a new channel by default.

Valid Platforms

ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines

If this channel is used as a redundant-pair link (see redundancy protocol (cfg-channel)) and redundancy is enabled (enable (cfg-redundancy)), this command causes the standby peer to reboot. The reboot does not disrupt any storage services, but the ARX peers cannot function as a redundant pair while the link is shut down. Additionally, a quorum-disk failure or disconnection would cause the active peer to reboot, too. If you proceed with shutting down the link, you should establish a new one as soon as possible: use the redundancy protocol or redundancy protocol (cfg-channel) command on another port or channel to establish a new redundant-pair link.

The CLI prompts for confirmation before shutting down a redundant-pair link; enter **yes** to proceed with the interface shutdown and the reboot.

Samples

```
bstnA(cfg-channel[4])# shutdown
```

shuts down channel 4.

```
bstnA(cfg-channel[6])# no shutdown
```

restarts channel 6.

Related Commands

channel
redundancy protocol (cfg-channel)

shutdown (cfg-if-gig)

Purpose

Use no shutdown to start traffic on the current port.

Use the shutdown command to stop the port.

Mode

cfg-if-gig

Security Role(s)

network-engineer or crypto-officer

Syntax

```
shutdown
no shutdown
```

Default(s)

```
shutdown
```

Guidelines

You cannot use this command on a channel member; use the shutdown (cfg-channel) command to stop traffic on the channel.

The ARX-1500 and ARX-2500 allow only a single logical interface (port or channel) to carry any given VLAN. That is, any VLAN can only be assigned to one port or channel (interface gigabit, interface ten-gigabit, or channel). The CLI therefore prevents a no shutdown for any port on those chassis types if it carries a VLAN that is already assigned to an active channel or another active port. The CLI returns an error in this case, displaying the VLAN that has the conflict along with the active port(s) that already carry the VLAN. You can use show vlan *n* to see the port or channel that carries VLAN *n*.

Guidelines: Shutting Down the Redundant-Pair Link

If this interface is the only one used as a redundant-pair link (through redundancy protocol or redundancy (cfg-if-vlan)) and redundancy is enabled (enable (cfg-redundancy)), this command causes the standby peer to reboot. The reboot does not disrupt any storage services, but the ARX peers cannot function as a redundant pair while the link is shut down. Additionally, a quorum-disk failure or disconnection would cause the active peer to reboot, too. If you proceed with shutting down the link, you should establish a new one as soon as possible: use the redundancy protocol or redundancy protocol (cfg-channel) command on another port or channel to establish a new redundant-pair link.

The CLI prompts for confirmation before shutting down a redundant-pair link; enter **yes** to proceed with the interface shutdown and the reboot.

Samples

```
bstnA(cfg-if-gig[2/4])# shutdown
```

shuts down the current port, 2/4.

```
bstnA(cfg-if-gig[2/4])# no shutdown
```

restarts the same port.

Related Commands

interface gigabit
shutdown (cfg-channel)
redundancy protocol
redundancy (cfg-if-vlan)

shutdown (cfg-if-ten-gig)

Purpose

Use no shutdown to start traffic on the current ten-gigabit port.

Use the shutdown command to stop the port.

Mode

cfg-if-ten-gig

Security Role(s)

network-engineer or crypto-officer

Syntax

```
shutdown
no shutdown
```

Default(s)

```
shutdown
```

Valid Platforms

ARX-2500 and ARX-4000 only

Guidelines

You cannot use this command on a channel member; use the shutdown (cfg-channel) command to stop traffic on the channel.

The ARX-2500 allows only a single logical interface (port or channel) to carry any given VLAN. That is, any VLAN can only be assigned to one port or channel (interface gigabit, interface ten-gigabit, or channel). The CLI therefore prevents a no shutdown for any port on those chassis types if it carries a VLAN that is already assigned to an active channel or another active port. The CLI returns an error in this case, displaying the VLAN that has the conflict along with the active port(s) that already carry the VLAN. You can use show vlan *n* to see the port or channel that carries VLAN *n*.

Guidelines: Shutting Down the Redundant-Pair Link

If this interface is used as a redundant-pair link (see redundancy protocol) and redundancy is enabled (enable (cfg-redundancy)), this command causes the standby peer to reboot. The reboot does not disrupt any storage services, but the ARX peers cannot function as a redundant pair while the link is shut down. Additionally, a quorum-disk failure or disconnection would cause the active peer to reboot, too. If you proceed with shutting down the link, you should establish a new one as soon as possible: use the redundancy protocol or redundancy protocol (cfg-channel) command on another port or channel to establish a new redundant-pair link.

The CLI prompts for confirmation before shutting down a redundant-pair link; enter **yes** to proceed with the interface shutdown and the reboot.

Samples

```
bstnA(cfg-if-ten-gig[2/1])# shutdown
```

shuts down the current ten-gigabit port, 2/1.

```
bstnA(cfg-if-ten-gig[2/1])# no shutdown
```

restarts the same port.

Related Commands

interface ten-gigabit
shutdown (cfg-channel)
redundancy protocol

shutdown (cfg-stp)

Purpose From cfg-stp mode, use the shutdown command to stop all spanning-tree processing.

Use no shutdown to restart spanning tree.

Mode cfg-stp

Security Role(s) network-engineer or crypto-officer

Syntax

```
shutdown
no shutdown
```

Default(s) **shutdown**: spanning tree is disabled (along with switch forwarding; see below) by default.

Valid Platforms ARX-2000 and ARX-4000

Guidelines Before you can run no shutdown to start spanning-tree processing, you must use the switch-forwarding enable command. This permits the ARX to forward packets from one client/server port to another. By default, the ARX behaves as an end station rather than a bridge. The switch-forwarding enable command runs no shutdown as a side-effect; conversely, no switch-forwarding enable runs shutdown as a side-effect.

If you shut down spanning tree, neighboring bridges must adjust the spanning-tree topology as though the ARX had been removed from the network.

Important

Spanning tree is activated along with switch forwarding to protect against possible network loops. You have the option to disable spanning tree while switch forwarding is active (using this command), but this is dangerous. Do not disable spanning tree unless you are confident that the ARX cannot create a loop.

Samples

```
bstnA(cfg-stp)# shutdown
```

shuts down spanning tree.

```
bstnA(cfg-stp)# no shutdown
```

restarts the spanning tree.

Related Commands spanning-tree
switch-forwarding enable

spanning-tree

Purpose

The Spanning-Tree Protocol (STP) creates a loop-free topology in bridged networks. Use the spanning-tree command to configure the spanning-tree parameters on the ARX.

Use the no form of this command to revert the spanning-tree parameters back to their defaults.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
spanning-tree
no spanning-tree
```

Default(s)

None

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

The Spanning-Tree protocol is defined in IEEE 802.1D.

Before the ARX can use STP, you must use the switch-forwarding enable command. This permits the ARX to forward packets from one client/server port to another. By default, the ARX behaves as an end station rather than a bridge.

The spanning-tree command puts you into cfg-stp mode, where you must set the protocol (STP or Rapid STP) with the protocol (cfg-stp) command. There are several spanning-tree options you can set from cfg-stp mode, and there are some port-level options you can set from cfg-if-gig mode; see *Related Commands*, below.

The no form of the command resets all global parameters to their respective defaults. The global parameters are the protocol, bridge Priority, Hello Time, Max Age, and Forward Delay.

Samples

```
bstnA(cfg)# spanning-tree
bstnA(cfg-stp)#
```

enters cfg-stp mode.

```
bstnA(cfg)# no spanning-tree
bstnA(cfg)#
```

sets all global spanning-tree parameters to their default values.

Related Commands

switch-forwarding enable
protocol (cfg-stp)

Cfg-stp commands for setting global 802.1D parameters:

priority (cfg-stp)
hello-time
max-age
forward-delay

A cfg-if-gig command to identify each RSTP Edge Port:

spanning-tree edgeport

Cfg-if-gig commands for setting Port Cost and Port Priority:

spanning-tree cost
spanning-tree priority

Show commands:

show spanning-tree summary
show spanning-tree detailed
show spanning-tree interface

spanning-tree cost

Purpose

IEEE 802.1D defines *Port Cost* as the relative cost to relay a frame to the root bridge. A lower cost is preferred. The port with the lowest cost is the *root port* for the ARX. Neighboring bridges compare the Port Costs of spanning-tree ports to *designate* one of them for their traffic. Use the spanning-tree cost command to set the Port Cost for the current port.

Use no spanning-tree cost to reset the Port Cost to the default.

Mode

cfg-if-gig

Security Role(s)

network-engineer or crypto-officer

Syntax

`spanning-tree cost` ***`port-cost`***
`no spanning-tree cost`

port-cost (1-200,000) is the Port Cost.

Default(s)

20,000

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

The Spanning-Tree protocol is defined in IEEE 802.1D.

This influences the election of a *designated port* by neighboring bridges. For installations with multiple ports on the same LAN segment, set the Port Cost lower for the fastest ports.

Use the spanning-tree priority command to set the Port Priority.

Samples

```
bstnA(cfg-if-gig[2/4])# spanning-tree cost 1
```

sets the lowest-possible Port Cost for port 2/4. This port will be the root port for the ARX. If the port is connected to upstream bridges in the tree, it will likely be *designated* by those bridges, too.

```
bstnA(cfg-if-gig[2/5])# no spanning-tree cost
```

reverts port 2/5 to the default Port Cost.

Related Commands

interface gigabit
spanning-tree priority

spanning-tree edgeport

Purpose

Rapid Spanning Tree Protocol (RSTP) defines an *Edge Port* as a port that connects to only one other port, as opposed to several ports on a LAN segment. This applies to RSTP and MSTP configurations only. Use the spanning-tree edgeport command to declare that the current port is an Edge Port.

Use the no form to declare that the current port is not an Edge Port.

Mode

cfg-if-gig

Security Role(s)

network-engineer or crypto-officer

Syntax

```
spanning-tree edgeport
no spanning-tree edgeport
```

Default(s)

no; new ports are not Edge Ports by default.

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

The Spanning-Tree protocol is defined in IEEE 802.1D. RSTP is defined in IEEE 802.1w.

Whenever there is a spanning-tree topology change, RSTP updates its Edge Ports faster than the others. Edge ports cannot possibly create bridge loops, so RSTP can move an Edge Port from *discarding* state directly to *forwarding* state, skipping the *learning* state.

Samples

```
bstnA(cfg-if-gig[2/4])# spanning-tree edgeport
```

sets port 2/4 as an Edge Port. In the event of a spanning-tree topology change, this port will convert to the new topology immediately.

```
bstnA(cfg-if-gig[2/6])# no spanning-tree edgeport
```

designates port 2/6 as one that connects to multiple ports.

Related Commands

interface gigabit

spanning-tree priority

Purpose

IEEE 802.1D defines *Port Priority* as the relative priority between ports with equal Port Costs. Neighboring bridges compare the Port Costs of spanning-tree ports to *designate* one of them for their traffic; if two or more Port Costs are the same, the protocol uses the Port Priority to break the tie. A lower number represents a higher priority. Use the spanning-tree priority command to set the Port Priority for the current port.

Use no spanning-tree cost to reset the Port Priority to the default.

Mode

cfg-if-gig

Security Role(s)

network-engineer or crypto-officer

Syntax

`spanning-tree priority` ***`port-priority`***
`no spanning-tree priority`

port-priority (0-240) is the Port Priority. 0 (zero) is the highest priority, 240 is the lowest.

Default(s)

128

Valid Platforms

ARX-2000 and ARX-4000

Guidelines

The Spanning-Tree protocol is defined in IEEE 802.1D.

This influences the election of a *designated port* by neighboring bridges. For installations with multiple ports on the same LAN segment, set the Port Priority higher (that is, to a lower number) for the ports that you want to be designated by their neighbors.

Use the spanning-tree cost command to set the Port Cost, which has a greater influence on the election of designated ports.

Samples

```
bstnA(cfg-if-gig[2/3])# spanning-tree priority 0
```

sets the lowest-possible Port Priority for port 3.

```
bstnA(cfg-if-gig[2/4])# no spanning-tree priority
```

reverts port 2/4 to the default Port Cost.

Related Commands

interface gigabit
spanning-tree cost

spanning-tree shutdown

Purpose

You can shut down spanning-tree processing on an individual port, thus removing the port from the spanning-tree topology. This stops the port from relaying frames from one LAN segment to another. You can also shutdown spanning tree for an aggregated channel. Use the spanning-tree shutdown command to shut down spanning tree at the current port or channel.

Use no spanning-tree shutdown to put the current port back into the spanning tree.

Modes

cfg-if-gig

cfg-if-ten-gig

cfg-channel

Security Role(s)

network-engineer or crypto-officer

Syntax

```
spanning-tree shutdown
no spanning-tree shutdown
```

Default(s)

All ports and channels are part of the spanning tree by default.

Valid Platforms

In cfg-if-gig mode or cfg-channel: ARX-2000 and ARX-4000.

In cfg-if-ten-gig mode, ARX-4000 only.

Guidelines

The Spanning-Tree protocol is defined in IEEE 802.1D.

Use the spanning-tree command to configure spanning-tree on the ARX.

You cannot use this command on a port or channel used in a redundant-pair link (see the documentation for the redundancy protocol and redundancy protocol (cfg-channel) commands).

The ARX must be allowed to forward packets (with the switch-forwarding enable command) before you can use this command to enable spanning tree.

Samples

```
bstnA(cfg-if-gig[2/3])# spanning-tree shutdown
```
removes port 2/3 from the spanning tree.

```
bstnA(cfg-channel[1])# spanning-tree shutdown
```
removes channel 1 from the spanning tree.

```
bstnA(cfg-if-gig[2/3])# no spanning-tree shutdown
```
returns port 2/3 to the spanning tree.

```
bstnA(cfg-if-ten-gig[2/2])# spanning-tree shutdown
```
removes port 2/2, a ten-gigabit port, from the spanning tree.

Related Commands

interface gigabit
interface ten-gigabit
channel
spanning-tree
switch-forwarding enable

speed (cfg-if-gig)

Purpose From cfg-if-gig mode, use the speed command to set the speed, line-type, and duplex configuration on a specified gigabit port.

Mode cfg-if-gig

Security Role(s) network-engineer or crypto-officer

Syntax **`speed {auto | 100-tx-half | 100-tx-full | 100-fx-full | 1000-full }`**

auto | 100-tx-half | 100-tx-full | 100-fx-full | 1000-full is a required choice:

auto makes the port auto-negotiate with its peer.

100-tx-half is fast Ethernet, 100 megabits per second (mbps), half-duplex.

100-tx-full is fast Ethernet, 100 mbps, full duplex.

100-fx-full is fiber Ethernet, 100 mbps, full duplex.

1000-full is fiber or copper Ethernet, 1000 mbps, full duplex.

Default(s) auto

Valid Platforms ARX-500, ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines ARX Gigabit Ethernet ports support automatic MDI/MDIX cross-over. This feature automatically corrects the polarity of the attached CAT5 cable, regardless if it is a cross-over or straight-through type.

The speed must be set manually (not to auto) for the port to be a member of a channel (see the channel command).

Sample

```
bstnA(cfg-if-gig[2/5])# speed 100-fx-full
```

For the interface at slot 2, port 5, sets the speed to 100 mbps, the line type to fiber Ethernet, and the duplex configuration to full duplex.

Related Commands interface gigabit
show interface gigabit

switch-forwarding enable

Purpose By default, the ARX does not forward packets between its client/server ports; it behaves as an end station instead of a MAC bridge. Use the switch-forwarding enable command to enable packet forwarding and start using the ARX's bridging features.

Use the no form of this command to stop all packet forwarding.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax
```
switch-forwarding enable
no switch-forwarding enable
```

Default(s) disabled

Valid Platforms ARX-2000 and ARX-4000

Guidelines If switch forwarding is disabled, you cannot enable the spanning-tree protocol. When you enable switch forwarding with this command, spanning tree is enabled as a side-effect. Conversely, when you disable switch forwarding then spanning tree is disabled. The CLI warns you of these side-effects and prompts for confirmation; enter **yes** to proceed.

Important

Spanning tree is enabled along with switch forwarding to protect against possible network loops. You have the option to disable spanning tree while switch forwarding is active (using shutdown (cfg-stp)), but this is dangerous. Do not disable spanning tree unless you are confident that the ARX cannot create a loop.

Use the show version command to see the current setting for this ARX.

Samples
```
bstnA(cfg)# switch-forwarding enable
Warning: Enable switch forwarding between all Ethernet Interfaces?
This will also enable the Spanning Tree Protocol. [yes/no] yes
bstnA(cfg)#
```
enables switch forwarding and spanning tree. This ARX can now behave as a MAC bridge.

```
bstnA(cfg)# no switch-forwarding enable
Warning: Disable switch forwarding between all Ethernet Interfaces?
This will also disable the Spanning Tree Protocol. [yes/no] yes
bstnA(cfg)#
```
returns the "bstnA" switch to its default status. It no longer forwards packets from one client/server port to another.

Related Commands spanning-tree
show version

tag

Purpose

A port with VLAN *tagging* enabled adds the VLAN ID (VID) to outbound frames, and only accepts ingress packets that are specifically tagged for the current VLAN. This is required to support multiple external VLANs. Use the tag command to add a single tag-enabled port or a range of tag-enabled ports to the current VLAN.

Use the no form of this command to disable tagging on a port(s).

Mode

cfg-vlan

Security Role(s)

network-engineer or crypto-officer

Syntax

tag *slot/port* [to *slot/port*]
no tag *slot/port* [to *slot/port*]

slot/port (2/1-14 on ARX-4000; 1/1-4 or 2/1-2 on ARX-2500; 1/1-12 on ARX-2000; 1/1-8 on ARX-1500) is a single port or the first port in a range of ports.

to *slot/port* (optional) is the last port in a range of ports. You cannot use this option on the ARX-2500 or ARX-1500, which allow only a single port to carry each VLAN.

Use the show interface summary command to locate all slots and ports.

Default(s)

None

Valid Platforms

ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines

IEEE 802.1Q defines VLANs.

If you select ports that were already configured as members of the VLAN, this command simply enables tagging on the selected ports. This means that incoming packets on those ports are not accepted as part of the current VLAN unless they are explicitly tagged with the VLAN's ID.

Use the members (cfg-vlan) command to add ports with tagging disabled.

A port can be a member of multiple VLANs, as long as each port is tagged for the VLANs. Use vlan to go into cfg-vlan mode for a VLAN, then use tag *slot/port ...* to tag some ports for that VLAN. You can repeat this for multiple VLANs, tagging the same set of ports for each.

To assign a channel to the VLAN with tagging enabled, use the vlan-tag command.

On the ARX-1500 and the ARX-2500, only a single port (interface gigabit or interface ten-gigabit) or channel can carry any given VLAN. That is, if you use this command to tag a port with a particular VLAN ID, you cannot tag any other port or channel with the same VLAN ID. This includes VLAN 1; if multiple channels or ports default to VLAN 1, all but one of them must be disabled (with shutdown (cfg-if-gig), shutdown (cfg-if-ten-gig), or shutdown (cfg-channel)).

Guidelines: Removing the Last Port from the Redundancy Link

In the following circumstances, the no tag command causes the backup ARX peer to reboot:

- the interface vlan command establishes an in-band (VLAN) management address for this VLAN,
- the redundancy (cfg-if-vlan) command establishes the above management address as the local end of the redundancy link,
- redundancy is active between the ARX peers, and
- the no tag command is removing the last port(s) from the VLAN.

The reboot does not disrupt any storage services, but the ARX peers cannot function as a redundant pair while the link is shut down. Additionally, a quorum-disk failure or disconnection would cause the active peer to reboot, too. If you proceed with shutting down the link, you should establish a new one as soon as possible: use the redundancy protocol, redundancy protocol (cfg-channel), or redundancy (cfg-if-vlan) command on another port, channel, or VLAN interface to establish a new redundant-pair link.

Samples

```
bstnA(cfg-vlan[1])# tag 2/3 to 2/5
```

adds ports 2/3-2/5, tagging enabled, to VLAN 1.

```
bstnA(cfg-vlan[1])# tag 2/6
```

adds port 2/6, tagging enabled, to VLAN 1.

```
bstnA(cfg-vlan[2])# no tag 2/7
```

disables tagging for port 2/7, VLAN 2. The port will not tag any outbound frames destined for a MAC on VLAN 2.

Related Commands

vlan
members (cfg-vlan)

trap shutdown

Purpose From cfg-channel mode, use the no trap shutdown command to activate SNMP traps for the current channel.

Use the affirmative form, trap shutdown, to stop issuing traps.

Mode cfg-channel

Security Role(s) network-engineer or crypto-officer

Syntax

```
trap shutdown
no trap shutdown
```

Default(s) trap shutdown

Valid Platforms ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines This activates link traps (link up/down) for the channel as a whole, as opposed to its member ports.

Samples

```
bstnA(cfg-channel[1])# no trap shutdown
```

activates SNMP traps from channel 1.

```
bstnA(cfg-channel[7])# trap shutdown
```

shuts off SNMP traps from channel 7.

Related Commands channel

vlan

Purpose

A Virtual Bridged Local Area Network (*VLAN*) is a group of physically-separated MAC addresses that appear as a single LAN segment. Devices on the same VLAN appear to be physically co-located even though some of the devices may be on different floors or different buildings. VLAN membership is often driven by human factors like departmental membership in a company. A VLAN often carries a single IP subnet. Use the vlan command to begin configuring a VLAN.

Use the no form of the command to remove a VLAN configuration.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`vlan` *`vlan-id`*
`no vlan` *`vlan-id`*

vlan-id (1-4009) is an ID you choose for the VLAN. If the VLAN is already configured, this command edits its configuration.

Default(s)

On all platforms except the ARX-1500, ARX-2500, and ARX-VE, the following VLANs are configured by default

- VLAN 1, which carries all client and server subnets.
- another VLAN to carry *private* IP traffic amongst the internal processes of the ARX.
- a third VLAN to carry *metalog* IP traffic amongst the internal processes of the ARX.

The ARX-1500 and ARX-2500 do not use internal VLANs for their private or metalog traffic.

Valid Platforms

any *except* ARX-VE

Guidelines

IEEE 802.1Q defines VLANs.

This command puts you into cfg-vlan mode, where you must configure at least one member port with either the members (cfg-vlan) command or the tag command. You can edit the ingress options for the VLAN members through some cfg-if-gig commands. See *Related Commands*, below, for a complete list of CLI options.

The ARX-1500 and ARX-2500 support only one port per VLAN. If you require multiple ports to carry a given VLAN on those chassis types, aggregate them into a channel and assign the VLAN to that channel with vlan (cfg-channel) or vlan-tag.

Samples

```
bstnA(cfg)# vlan 2
bstnA(cfg-vlan[2])#
```

creates a configuration for VLAN 2.

```
bstnA(cfg)# no vlan 18
```

removes the configuration for VLAN 18.

Related Commands

members (cfg-vlan)

adds ports to the VLAN with tagging disabled.

tag

adds ports to the VLAN with tagging enabled, or enables tagging for existing member ports.

description (cfg-vlan)

names the VLAN, for show commands.

Show commands:

show vlan summary

lists all configured VLANs.

show vlan

shows details for one VLAN.

vlan (cfg-channel)

Purpose

From cfg-channel mode, use the vlan command to change the VLAN for the current channel.

Use no vlan to revert the current channel to the default VLAN.

Mode

cfg-channel

Security Role(s)

network-engineer or crypto-officer

Syntax

vlan *vlan-id*
no vlan *vlan-id*

vlan-id (1-4095) identifies the VLAN to add or remove.

Default(s)

VLAN 1

tagging disabled

Valid Platforms

ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines

Before you assign a channel to a VLAN, the channel must have at least one member port (see members (cfg-channel)), and all members must be disabled (shutdown (cfg-if-gig) or shutdown (cfg-if-ten-gig)).

On the ARX-1500 and the ARX-2500, only a single port (interface gigabit or interface ten-gigabit) or channel can carry any given VLAN. That is, if you use this command to assign a VLAN to the current channel, you cannot assign the same VLAN to any other port or channel. This includes VLAN 1; if multiple channels or ports default to VLAN 1, all but one of them must be disabled (with shutdown (cfg-if-gig), shutdown (cfg-if-ten-gig), or shutdown (cfg-channel)).

A channel with VLAN *tagging* enabled adds a VLAN ID to every outgoing frame. Use the vlan-tag command to enable VLAN tagging for the current channel. If VLAN tagging is disabled for the channel, the channel can belong to only one VLAN.

Samples

```
bstnA(cfg-channel[1])# vlan 5
```

assigns channel 1 to VLAN 5, untagged.

```
bstnA(cfg-channel[7])# no vlan 2
```

removes channel 7 from VLAN 2. If channel 7 did not have VLAN tagging set, it reverts to VLAN 1.

Related Commands

channel
members (cfg-channel)
shutdown (cfg-if-gig)
shutdown (cfg-if-ten-gig)
vlan-tag

vlan-tag

Purpose

You can configure a channel to tag its outgoing packets with a VLAN ID, and to start accepting only explicitly-tagged ingress packets into the current VLAN. Channels with tagging enabled can support multiple VLANs; channels without tagging enabled can support only one. From cfg-channel mode, use the vlan-tag command to add a tagged VLAN to the current channel.

Use no vlan-tag to remove a tagged VLAN.

Mode

cfg-channel

Security Role(s)

network-engineer or crypto-officer

Syntax

vlan-tag *vlan-id*
no vlan-tag *vlan-id*

vlan-id (1-4095) identifies the VLAN to add or remove.

Default(s)

channels belong to VLAN 1 by default.

tagging disabled

Valid Platforms

ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines

A channel with VLAN *tagging* enabled adds a VLAN ID to every outgoing frame. Repeat this command with different VLAN IDs to carry multiple VLANs on the channel. Once you tag the channel for one or more VLANs, it only accepts ingress frames that are explicitly tagged for one of the VLANs.

If you remove the last VLAN with the no form of this command, tagging is disabled for the channel and the channel is assigned to VLAN 1.

On the ARX-1500 and the ARX-2500, only a single port (interface gigabit or interface ten-gigabit) or channel can carry any given VLAN. That is, if you use this command to tag a channel with a particular VLAN ID, you cannot tag any other port or channel with the same VLAN ID. This includes VLAN 1; if multiple channels or ports default to VLAN 1, all but one of them must be disabled (with shutdown (cfg-if-gig), shutdown (cfg-if-ten-gig), or shutdown (cfg-channel)).

Samples

```
bstnA(cfg-channel[3])# vlan-tag 6
```

assigns channel 3 to VLAN 6 with tagging enabled.

```
bstnA(cfg-channel[7])# no vlan-tag 2
```

removes channel 7 from VLAN 2. If channel 7 did not have another VLAN set for tagging, it reverts to VLAN 1 with no tagging.

Related Commands

channel
jumbo mtu

10

Layer 3 (Network Layer)

arp

Purpose The Address-Resolution Protocol (ARP) maps IP addresses to MAC addresses in an ARP table. Each network processor in the ARX has its own ARP table. Use the arp command to add a static entry to all ARP tables on the switch.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
arp ip-address mac-address [vlan vlan-id]
no arp ip-address
```

ip-address (for example, 10.125.16.3) is the IP address.

mac-address (for example, 12:34:56:78:9a:bc) is the MAC address you are statically mapping to the *ip-address*.

vlan *vlan-id* (optional; 0-4096) applies the mapping to a single VLAN.

Default(s) None

Guidelines The network processors are the ones behind the client/server interfaces, as well as the one behind the out-of-band MGMT interface.

Use the show arp command to show all ARP-table entries. Use the clear arp command to clear all dynamic-ARP entries, learned from neighboring equipment.

Samples

```
bstnA(cfg)# arp 192.168.25.38 11:54:d6:2a:95:f2
```

adds a static entry to the ARP table.

```
bstnA(cfg)# arp 10.1.1.159 11:df:45:b3:95:36 vlan 4
```

adds a static entry to the ARP table and applies it to VLAN 4.

```
bstnA(cfg)# no arp 172.16.209.55
```

removes an entry from the ARP table.

Related Commands

show arp
clear arp

arp gratuitous

Purpose The Address-Resolution Protocol (ARP) maps IP addresses to MAC addresses in an ARP table. Each processor in the ARX has its own ARP table. Use the arp gratuitous command to issue gratuitous ARP entries for a single IP address or all IP addresses on the switch (including virtual IP addresses (VIPs), management IP addresses (MIPs), and proxy IP addresses (XIPs)).

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `arp gratuitous {`***`ip-address`*** `| yes}`

ip-address (0.0.0.0. to 255.255.255.255) is the IP address for which you want to add a gratuitous ARP entry. If there is an error, the system displays the error message on the CLI console.

yes specifies to add gratuitous ARP entries for all IP addresses owned by the switch. If there is an error, the system logs the error message (failed GARP) in the syslog.

Default(s) None.

Guidelines If you do not enter an IP address or **yes** at the end of this command, the switch sends a gratuitous ARP for all of its publicly-visible IP addresses (such as VIPs). Before sending all of this traffic, the CLI prompts for confirmation: enter **yes** to proceed.

Use the show arp command to show all ARP-table entries. Use the clear arp command to clear all dynamic-ARP entries, learned from neighboring equipment.

Samples

```
bstnA# arp gratuitous 192.168.25.38
```

sends a gratuitous ARP entry for one IP address.

```
bstnA# arp gratuitous
Send a gratuitous ARP for all of this switch's IP addresses? [yes/no]
yes
```

sends a gratuitous ARP for all IP addresses on the switch (VIPs, MIPs, and Proxy IPs).

```
bstnA# arp gratuitous yes
```

also sends gratuitous ARPs for all IPs, but skips the prompt.

Related Commands

show arp
clear arp
arp

clear arp

Purpose The Address-Resolution Protocol (ARP) maps IP addresses to MAC addresses in an ARP table. ARP-table entries are either learned from neighbors (*dynamic*), set through the CLI (*static*), or set by internal software (*local*). Use the clear arp command to clear all dynamic entries form the ARP table.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `clear arp [from slot.processor]`

from *slot.processor* (optional) specifies one module processor. This option is not supported on the ARX-1500, ARX-2500, or ARX-VE. Each network-connected processor (the ones behind the client/server ports and the one behind the MGMT port) has its own ARP table.

slot (1-2 for an ARX-4000;1 for all others) is the slot number.

processor is the processor number. Use the show processors command to show all processors and their associated *slot.processor* IDs.

Default(s) None

Guidelines Use the show arp command to show all ARP-table entries. Use the arp command to create a static-ARP entry.

Samples

```
bstnA# clear arp
```

clears all dynamic-ARP entries from the ARX.

```
bstnA# clear arp from 2.2
```

clears the dynamic-ARP entries from processor 2.2 only.

Related Commands

show processors
show arp
arp

clear counters mgmt

Purpose Administrators can log into the CLI or GUI through the out-of-band (OOB) management interface. Port 1/1 is used as the OOB management interface on the ARX-1500 and ARX-2500. The show interface mgmt stats command shows packet counters for this management interface. On the ARX-1500 or the ARX-2500, you can use the clear counters mgmt command to reset these counters to 0 (zero).

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `clear counters mgmt`

Default(s) None

Platforms ARX-1500 or ARX-2500

Guidelines Use the show interface mgmt stats command to show all ingress and egress counters for the out-of-band management interface.

The ARX-1500 and the ARX-2500 can operate without using port 1/1 as an out-of-band management interface; you can use the no interface mgmt command to designate port 1/1 for client/server traffic instead of management traffic. If the port is not used for out-of-band management statistics, this command is unnecessary. You can use the show interface mgmt command (with or without the stats option) to confirm that port 1/1 is designated as an out-of-band management interface.

Sample

```
stoweA# clear counters mgmt
```

clears all out-of-band-management counters from the current ARX.

Related Commands show interface mgmt
interface mgmt

description (cfg-if-vlan)

Purpose

You can configure one in-band management interface per VLAN. From cfg-if-vlan mode, use the optional description command to create a descriptive string for an in-band-management interface.

Use the no form of the command to delete the interface description.

Mode

cfg-if-vlan

Security Role(s)

network-engineer or crypto-officer

Syntax

`description` ***`text`***
`no description`

text (up to 128 characters) is your description. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s)

No description

Guidelines

The description appears in the show interface vlan command.

Sample

`bstnA(cfg-if-vlan[1])#` **`description “management for vlan 1”`**
specifies a description for the in-band-management interface on VLAN 1.

Related Commands

show interface vlan

description (cfg-mgmt)

Purpose

An ARX (except the ARX-VE) can have one out-of-band management interface, on a separate IP network from all clients and servers. From cfg-mgmt mode, use the optional description command to create a descriptive string for the out-of-band management interface.

Use the no form of the command to delete the interface description.

Mode

cfg-mgmt

Security Role(s)

network-engineer or crypto-officer

Syntax

`description` ***`text`***
`no description`

text (up to 128 characters) is your description. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s)

`no description`

Platforms

any *except* ARX-VE

Guidelines

The description appears in the show interface mgmt command.

Sample

`bstnA(cfg-mgmt)#` **`description “oob management”`**

specifies a description for the management interface.

Related Commands

show interface mgmt

interface mgmt

Purpose

The out-of-band management port is on the front panel of the ARX, typically labeled MGMT. You configure this as part of the initial-boot process. Use the interface mgmt command to modify the management-interface configuration.

On an ARX-1500 or ARX-2500, you can use no interface mgmt to stop using port 1/1 for out-of-band management. You can then use other commands to use the port for client/server traffic.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
interface mgmt
no interface mgmt
```

Default(s)

None.

Platforms

any *except* ARX-VE

Guidelines

The ARX software keeps a separate IP routing table for its out-of-band management interface; this network is designed to be completely separate and disjoint from any client or server subnets. This interface is designed for installations with a separate subnet exclusively for management use.

This command puts you into cfg-mgmt mode, where you can set several configuration parameters for the management interface. Use the ip address (cfg-mgmt) command to change the IP address. Use the description (cfg-mgmt) command to set an optional description for the interface, for show commands. Use the shutdown (cfg-mgmt) command to shut down the interface.

To show the configuration of the management interface, use show interface mgmt.

To set up an in-band management interface for one or more VLANs, use the interface vlan command.

You cannot use no interface mgmt if you are logged into the CLI through the out-of-band management interface; this would abruptly end your CLI session. The CLI also prevents no interface mgmt if the ARX has redundancy configured; on many platforms, this interface is used for important redundancy-related traffic.

Guidelines: ARX-1500 and ARX-2500

On the ARX-1500 and ARX-2500, port 1/1 is set up as the out-of-band management interface by default. For installations with no separate management subnet, you can use the no interface mgmt command to stop using port 1/1 for out-of-band management. You can then edit the port as a standard client/server interface with the interface gigabit 1/1 command.

Guidelines: No Out-of-Band Management on the ARX-VE

The ARX-VE is a Virtual Appliance (VA) that only uses a single VNIC, so it only has an in-band management interface. This command is therefore unavailable in the ARX-VE CLI.

Sample

```
bstnA(cfg)# interface mgmt
bstnA(cfg-mgmt)#
```

Related Commands

ip address (cfg-mgmt)
description (cfg-mgmt)
shutdown (cfg-mgmt)
show interface mgmt
interface vlan

interface vlan

Purpose

You can configure one in-band management interface per supported VLAN. Use the interface vlan command to begin configuring the management interface for a VLAN.

Use the no form to remove the in-band-management interface for a VLAN.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
interface vlan vlan-id
no interface vlan vlan-id
```

vlan-id (1-4096) identifies the VLAN. Use the show vlan summary command for a list of all configured VLANs.

Default(s)

None

Guidelines

This command puts you into cfg-if-vlan mode, where you can set several configuration parameters for the in-band management interface. Use the ip address (cfg-if-vlan) command to set the IP address. Use the description (cfg-if-vlan) command to set an optional description for the interface, for show commands. Use the shutdown (cfg-if-vlan) command to shut down the interface.

You can re-use this interface as a connection to the switch's redundant peer and/or to multiple switches on the same RON. The redundancy (cfg-if-vlan) command makes the interface eligible for the initial rendezvous with a redundant peer; this command is required for the ARX-1500 and ARX-2500, which use this layer-3 connection for exchanging heartbeats and metalog data. The ron tunnel command enters a sub-mode for configuring a RON tunnel to another ARX.

To show the configuration of the management interface, use show interface vlan.

Sample

```
bstnA(cfg)# interface vlan 9
bstnA(cfg-if-vlan[9])#
```

Related Commands

show vlan summary
ip address (cfg-if-vlan)
description (cfg-if-vlan)
shutdown (cfg-if-vlan)
show interface vlan
redundancy (cfg-if-vlan)
ron tunnel

ip address (cfg-if-vlan)

Purpose

Administrators can use the in-band management address to log into the CLI or GUI from a client or server VLAN. From cfg-if-vlan mode, use the ip address command to set the address for the current VLAN's in-band management interface.

Use the no form of this command to remove the IP address and disable the interface.

Important

In a redundant pair of ARXes, the network software uses an in-band (VLAN) management address as a home address for its communication with the quorum-disk. Without an in-band-management address and an ip route to the quorum disk, a failover is impossible. Additionally, any ron tunnels that use this address will fail if you remove it; a shadow-copy-rule depends on RON tunnels to communicate with other ARXes in the network. Use the no form of the command only on the advice of F5 Support.

Mode

cfg-if-vlan

Security Role(s)

network-engineer or crypto-officer

Syntax

```
ip address address mask
no ip address
```

address is the IP address you choose for the VLAN-management interface (for example, 192.168.108.223).

mask defines the network part of the address (for example, 255.255.255.0).

Default(s)

None

Guidelines

An ARX in a redundant pair requires an in-band (VLAN) management address with an IP route to its quorum disk. This is required so that the ARX can reach its quorum disk while in the backup role. Refer to the quorum-disk documentation for more details. Do not delete this address if this ARX has a quorum disk on the current VLAN.

A redundant pair of ARX-2500 or ARX-1500 devices requires an in-band (VLAN) management interface for each end of the redundancy link. The redundancy link is the channel (or possibly single link) that connects the redundant ARX devices. The IP address that you assign to the redundancy link (with this command) is used in the peer command when you join the redundant pair.

If the current in-band (VLAN) interface is the end point for one or more RON tunnels, those tunnels also depend in this IP address.

Sample

```
bstnA(cfg-if-vlan[9])# ip address 192.168.25.28 255.255.255.0
```

sets an in-band management IP for VLAN 9.

Related Commands

interface vlan
quorum-disk
peer
ron tunnel

ip address (cfg-mgmt)

Purpose Administrators can use the out-of-band management address to log into the CLI. This is configured as part of the initial-boot process. From cfg-mgmt mode, use the ip address command to change the address for the out-of-band management interface.

Use the no form of this command to remove the IP address and disable the interface.

Mode cfg-mgmt

Security Role(s) network-engineer or crypto-officer

Syntax

```
ip address address mask
no ip address
```

address is the IP address you choose for the management interface (for example, 10.1.1.10).

mask defines the network part of the address (for example, 255.255.255.0).

Default(s) None

Platforms any *except* ARX-VE

Guidelines This address must belong to a management network that is entirely distinct from any client subnet (established with virtual server) or the proxy-IP subnet (created by ip proxy-address). The MGMT interface uses a separate IP-routing table; use the ip route ... mgmt command to specify a default route (or any other static route) for the out-of-band management network.

Samples

```
bstnA(cfg-mgmt)# ip address 10.1.1.7 255.255.255.0
```

sets an out-of-band management IP for the ARX.

```
bstnA(cfg-mgmt)# no ip address
```

removes the IP configuration from the out-of-band management interface.

Related Commands interface mgmt
ip route

ip domain-list

Purpose

You can create an optional *search list* of domain names for the ARX to use in its DNS lookups. Whenever the switch needs to perform a DNS lookup for a hostname (for example, "fs5"), it appends a domain name (for example, "mycompany.com") and tries a DNS lookup; on failure, it appends the next domain name in the list; and so on. Use the ip domain-list command to add one domain name to the search list.

Use the no form of this command to remove a domain name from the search list.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
ip domain-list name
no ip domain-list name
```

name (1-255 characters) is a name for one local domain (for example, "myorg.org").

Default(s)

None

Guidelines

The search list is analogous to the "search" list in an /etc/resolv.conf file. You can enter this command multiple times, once for each search domain. The search domains are concatenated together, separated by spaces: the total length of this domain list cannot exceed 256 characters.

To identify a DNS server, use the ip name-server command. Use the show ip domain command to view the current DNS-lookup configuration.

Samples

```
bstnA(cfg)# ip domain-list estorage.com
bstnA(cfg)# ip domain-list enet.com
```

adds two domain names to the search list. The ARX will try "estorage.com" first when it looks up short names.

```
bstnA(cfg)# no ip domain-list enet.com
```

removes one domain name from the search list.

Related Commands

ip name-server
show ip domain

ip name-server

Purpose

This command identifies a DNS server that the ARX can use for DNS lookups. You can enter up to three DNS servers. Use the ip name-server command to add a DNS server.

Use the no form of this command to remove a DNS server from the list.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`ip name-server` ***`ip-address`***
`no ip name-server` ***`ip-address`***

ip-address identifies a DNS server (for example, "172.16.98.36"). This address must be on the server/proxy-IP subnet (see ip proxy-address) or reachable through a static route (see ip route).

Default(s)

None

Guidelines

You can enter this command multiple times, once for each DNS server; the switch supports a maximum of three. The servers are queried in the order that you enter them. The ARX switches from one DNS server to another if (and only if) the server is unreachable.

To support lookups of hostnames (for example, "myserver" instead of "myserver.mycompany.com"), you can declare one or more local domains (such as "mycompany.com") for the switch with the ip domain-list command.

All of the DNS servers should provide service for the same set of networks and domains.

Use the show ip domain command to view the current DNS-lookup configuration.

Samples

```
bstnA(cfg)# ip name-server 192.168.25.201
```

identifies a DNS server.

```
bstnA(cfg)# no ip name-server 192.168.25.212
```

removes one DNS server from the list.

Related Commands

ip proxy-address
ip route
ip domain-list
show ip domain

ip proxy-address

Purpose

Every NSM processor requires a *proxy IP* address to communicate with back-end devices. Use the ip proxy-address command to add a range of proxy IPs.

Use the no form of the command to remove a range of unused proxy-IP addresses.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
ip proxy-address address mask [vlan vlan-id] [count number]
[processor slot.proc]
no ip proxy-address address
```

address is the starting IP address for a range of proxy IPs (for example, 192.168.25.0).

mask is the subnet mask (for example, 255.255.255.0).

vlan *vlan-id* (optional, 1-65535) is the VLAN for this subnet, if there is one. Use show vlan summary for a complete list of configured VLANs.

count *number* (optional, 1-64) is the number of contiguous IP addresses to assign to the proxy pool.

slot.proc (optional: for example, 1.4) assigns the proxy-IP address to a particular NSM processor. You can only use this option if you specify a single IP.

Default(s)

vlan-id defaults to VLAN 1.

number defaults to 1.

Guidelines

You must configure one proxy-IP address per network processor, where

- the ARX-VE and ARX-500 each have 1,
- the ARX-1500 has 2,
- the ARX-2000 has 4,
- the ARX-2500 has 3 (or 4, depending on the setting for resource-profile), and
- the ARX-4000 has 12.

The proxy IP addresses must all belong to the same subnet, but do not need to be contiguous. You can use this command multiple times to define multiple ranges, as long as there are enough proxy IPs for all of the network processors.

Be sure to assign the correct proxy-IP addresses the first time. Once the proxy IP is assigned to an NSM processor, it is difficult to change. To change an assigned proxy IP, you must save your configuration (with the priv-exec copy startup-config command), remove it from the switch (delete startup-config), reboot (reload), edit the saved configuration with the correct proxy-IP addresses, and replay it (that is, copy it and paste it into the CLI). The CLI prompts for confirmation before making any change to the proxy-IP addresses; please examine your proxy-IP change carefully before you enter **yes** to proceed.

Sample

```
bstnA(cfg)# ip proxy-address 192.168.25.31 255.255.255.0 vlan 25 count
4
%WARNING: The IP proxy address changes will take effect,
and cannot be modified without clearing your startup-config.

The ip proxy-address configuration will be 4 proxy-address(es)
starting at IP 192.168.25.31 on VLAN 25.

Continue? [yes/no] yes
bstnA(cfg)# ip proxy-address 192.168.25.141 255.255.255.0 vlan 25 count
8
%WARNING: The IP proxy address changes will take effect,
and cannot be modified without clearing your startup-config.

The ip proxy-address configuration will be 8 proxy-address(es)
starting at IP 192.168.25.141 on VLAN 25.

Continue? [yes/no] yes
```

configures two ranges of proxy-IP addresses, for a total of 12.

Related Commands show vlan summary

ip route

Purpose Use the ip route command to configure a static IP route.

Use the no form of this command to remove a static route.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
ip route ip-subnet ip-mask gateway [distance] [mgmt]
no ip route ip-subnet ip-mask gateway [distance] [mgmt]
```

ip-subnet is the IP address of a remote subnet (for example, 172.16.151.0).

ip-mask defines the network part of the subnet (for example, 255.255.255.0).

gateway identifies the gateway to the subnet (for example, 192.168.25.1).

distance (optional; 1-255) is an arbitrary distance metric; if you configure two routes to the same subnet, the route with the lowest distance is used.

mgmt (optional) is a flag that adds the route to a separate routing table for the out-of-band management network. This option is not available on the ARX-VE, which does not have a separate out-of-band management interface. Also, the option does not apply to any ARX-1500 or an ARX-2500 where port 1/1 is being used for client/server traffic.

Default(s) ***distance*** - 128

Guidelines Configure a static route for every IP subnet with clients or servers that is outside any client subnet (defined by the virtual server command) or the proxy-IP subnet (see ip proxy-address). For a remote client subnet, the next-hop gateway must be in the subnet where their VIP resides. Similarly, a route to a remote server network must go through the proxy-IP subnet.

The ARX keeps a separate routing table for the out-of-band management interface. This management interface connects to a separate IP network, and therefore requires a routing table that is tailored to its network. Use the **mgmt** flag to add or remove a route from this table. The gateways for these routes must be on the management subnet defined by the ip address (cfg-mgmt) command.

You can also use multiple static routes to the same destination, each with different next hops. You can use different ***distance*** costs with each route to indicate your route preferences.

If the ARX has a redundant peer, you can use the critical route command to designate that a route is critical. If a critical route fails, the ARX may fail over to its peer.

Use the show ip route command to list all static routes, including the routes in the separate table for management routes.

Samples

```
bstnA(cfg)# ip route 172.16.231.0 255.255.255.0 192.168.25.1
```
creates a static route to a client subnet.

```
bstnA(cfg)# ip route 172.16.231.0 255.255.255.0 192.168.25.2 255
```
creates another static route to the same subnet. The distance is set very high, so this route would not be chosen unless the gateway fails for the previous route.

```
bstnA(cfg)# ip route 10.16.10.0 255.255.255.0 10.1.1.1 mgmt
```
creates a static route for the out-of-band management network.

```
bstnA(cfg)# no ip route 10.16.165.0 255.255.255.0 10.1.1.1 mgmt
```
removes a static route for the out-of-band management network.

Related Commands

virtual server
ip proxy-address
ip address (cfg-mgmt)
show ip route

ip route ... per-vlan

Purpose

Some installations have a firewall between the ARX and its clients, and require VIPs on multiple client VLANs. In those situations, the ARX's single default route (created with the ip route command) causes the ARX to send all response packets over the default route's VLAN. If that VLAN is not the same as the VIP's VLAN, the firewall may drop the response packet. For example, if there are VIPs on each of VLANs A, B, and C, the single default route can only go over one of those VLANs (for example, VLAN A). Clients from the other VLANs (B and C) would send requests to those VLANs and get responses from VLAN A. If the firewall is connected to each VLAN through different interfaces, the response packet arrives on a different interface than the request packet. A firewall drops such packets, with different source and destination interfaces. To solve this specific problem, on the advice of F5 Support, you can use the ip route ... per-vlan command to make a separate default route for each client VLAN.

Use the no form of this command to remove a VLAN-specific-default route.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
ip route 0.0.0.0 0.0.0.0 gateway [distance] per-vlan vlan-id
no ip route 0.0.0.0 0.0.0.0 gateway [distance] per-vlan vlan-id
```

0.0.0.0 0.0.0.0 is the IP subnet and mask for a default route. You cannot define a subnet-specific route on a per-VLAN basis.

gateway identifies the gateway to use for this VLAN (for example, 192.168.30.1).

distance (optional; 1-255) is an arbitrary distance metric; if you configure two routes to the same subnet, the route with the lowest distance is used.

per-vlan *vlan-id* (0-4096) identifies the VLAN for this default route. You must choose a VLAN that is already defined on the ARX; use show vlan summary for a list of defined VLANs, and use vlan to define a new one.

Default(s)

distance - 128

Platforms

ARX-500, ARX-2000, ARX-4000, and ARX-VE

Guidelines

Use this command only on the advice of F5 Support. A standard ip route has some features (listed below) that are not supported for the per-VLAN route.

Because of this route's limitations, the CLI prompts for confirmation before it accepts the per-VLAN route; enter **yes** to continue.

If a client request arrives over a VLAN without a per-VLAN-default route, the response goes over the ARX's default route (defined with the standard ip route command).

This command is not available on the ARX-1500 or ARX-2500. The ip route ... source-ip command performs a similar function for those platforms.

Guidelines: Limitations

As mentioned above, this type of default route has several limitations.

- You cannot create a per-VLAN route to a specific subnet. This is always a default route.
- You cannot designate a per-VLAN route as a critical route, as you can with a standard IP route. If you lose connectivity to the ***gateway***, the ARX does not fail over to its peer (which may have a better connection).
- Unlike a standard default route, you are limited to a single ***gateway*** for each VLAN. For a standard default route, you can enter multiple ip route commands with different gateways and different values for ***distance***. The ARX attempts to reach the gateway with the lowest distance value and tries a higher-distance gateway if the first is unreachable. This form of the command does not support a ***distance*** value, so you cannot establish redundant gateways for the VLAN's default route.
- This type of default route is used only for outbound IP packets where a VIP is the source address. (The virtual server command creates a VIP.)

Samples

```
bstnA(cfg)# ip route 0.0.0.0 0.0.0.0 192.168.76.1 1 per-vlan 38

This static route overrides the default route for the specified VLAN.
The route does not support health checks.

Do you want to continue? [yes/no] yes
```

creates a default route for VLAN 38 with the lowest possible distance metric, 1. Responses to packets from this VLAN go back out over the same VLAN.

```
prtlndB(cfg)# no ip route 0.0.0.0 0.0.0.0 per-vlan 99

Removing this route will result in the global default route being used
for this VLAN.

Do you want to proceed? [yes/no] yes
```

removes the default route for VLAN 99. Responses to packets from this VLAN use the default route designated by the standard ip route command.

Related Commands

ip route
ip route ... source-ip
show ip route
vlan
show vlan summary
critical route
virtual server

ip route ... source-ip

Purpose Some installations have a firewall between the ARX and its clients, and require VIPs on multiple client VLANs. In those situations, the ARX's single default route (created with the ip route command) causes the ARX to send all response packets over the default route's VLAN. If that VLAN is not the same as the VIP's VLAN, the firewall may drop the response packet. For example, if there are VIPs on each of VLANs A, B, and C, the single default route can only go over one of those VLANs (for example, VLAN A). Clients from the other VLANs (B and C) would send requests to those VLANs and get responses from VLAN A. If the firewall is connected to each VLAN through different interfaces, the response packet arrives on a different interface than the request packet. A firewall typically drops such packets, with different source and destination interfaces. To solve this specific problem on an ARX-1500 or ARX-2500, you can use the ip route ... source-ip command to make a separate default route for each VIP; any packet received at the VIP uses this default route with the same VIP as its source IP.

Use the no form of this command to remove a VIP-specific-default route.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
ip route 0.0.0.0 0.0.0.0 gateway [distance] source-ip vip
no ip route 0.0.0.0 0.0.0.0 [gateway] [distance] source-ip vip
```

0.0.0.0 0.0.0.0 is the IP subnet and mask for a default route. You cannot define a subnet-specific route on a per-VIP basis.

gateway identifies the gateway to use for this route (for example, 192.168.30.1).

distance (optional; 1-255) is an arbitrary distance metric; if you configure two routes to the same subnet, the route with the lowest distance is used.

source-ip *vip* identifies the VIP for this default route. You must choose a VIP that is already defined on the ARX; use show global server for a list of defined VIPs, and use the virtual server command to define a new one.

Default(s) ***distance*** - 128

Platforms ARX-1500 or ARX-2500

Guidelines This command only functions on the ARX-1500 or ARX-2500. The ip route ... per-vlan command performs a similar function on all other platforms.

Any VIP that you designate with this command needs a management interface on its VLAN. You can use the show global server *fqdn* command to find the VIP and VLAN for a given *fqdn*, then use interface vlan *vlan-id* to create a management interface for that VLAN.

If a client request arrives at a VIP without a per-VIP-default route, the response goes over the ARX's default route (defined with the standard ip route command).

Samples

```
canbyA(cfg)# ip route 0.0.0.0 0.0.0.0 192.168.121.1 1 source-ip
192.168.121.76
```

creates a default route for VIP 192.168.121.76 with the lowest possible distance metric, 1. Responses to packets from this VIP go back out over the same VIP.

```
stoweA(cfg)# no ip route 0.0.0.0 0.0.0.0 192.168.90.1 source-ip
192.168.90.29
```

removes the default route for VIP 192.168.90.29. Responses to packets from this VIP use the default route designated by the standard ip route command.

Related Commands

ip route
ip route ... per-vlan
show ip route
vlan
show vlan summary
critical route
virtual server

ntp server

Purpose

The ARX can synchronize its internal clock with an external Network Time Protocol (NTP) server. Use the ntp server command to identify an NTP server.

Use the no form of this command to disconnect from an NTP server.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
ntp server ip-address [version {3 | 4}]
no ntp server ip-address
```

ip-address (1-128 characters) identifies the external NTP server by its IP address.

> **Note**
>
> *Windows Domain Controllers (DCs) support NTP and can be used as NTP servers. You should only use a DC as your NTP server if all the back-end CIFS filers also use the DC in this way.*

version {3 | 4} (optional) is the NTP version to use, NTPv3 or SNTPv4.

Default(s)

version - 4 (SNTPv4)

Guidelines

Use the same NTP servers for the ARX that you use for the back-end filers and front-end clients. For namespace policy it is vital that the ARX is synchronized with the back-end servers; it is also required for installations that use Kerberos. Note that synchronization is more important than time accuracy; if all the servers agree on the wrong time, policy and Kerberos continue to function.

The default protocol for external NTP communication is Simple NTP (SNTP) v4, defined in RFC 2030. If you set the version to v3, the switch uses NTPv3 defined in RFC 1305.

You can repeat this command to configure up to 8 NTP servers. The switch selects the best server as specified by the NTP protocol.

Use show ntp servers to see the current NTP configuration. Use show clock to see the current time/date setting on the ARX.

You can set the internal clock manually with the clock set command. If this setting conflicts with NTP-server time, the NTP time overrides the manual setting.

Samples

```
bstnA(cfg)# ntp server 192.168.25.201
```

selects an NTP server at 192.168.25.201.

```
bstnA(cfg)# ntp server 192.168.25.202 version 3
```

selects another NTP server, but uses NTPv3 to communicate with it.

Related Commands

clock set
show ntp servers
show clock

redundancy (cfg-if-vlan)

Purpose

You can connect two identical ARXes as a redundant pair. To initially join, or *rendezvous*, each switch communicates over one of its management interfaces. This command makes it possible to use the current in-band (VLAN) management interface as the rendezvous interface.

An ARX-1500 or ARX-2500 also use this type of management interface as one end of a *redundant-pair link*. They use this link to exchange heartbeat messages and metalog data.

Use the no form to disallow the current interface from being used for a redundant-pair rendezvous.

Mode

cfg-if-vlan

Security Role(s)

network-engineer or crypto-officer

Syntax

```
redundancy
no redundancy
```

Default(s)

no redundancy

Platforms

any *except* ARX-VE

Guidelines

To select this management interface for rendezvous, go to this switch's redundant peer and use the peer command.

Guidelines: ARX-1500 and ARX-2500

As mentioned above, the ARX-1500 or ARX-2500 can also use this management interface as one end of a redundant-pair link. The redundant pair depends on this link for ongoing heartbeats and metalog updates.

For the best failover performance in this case, we strongly recommend a gigabit or ten-gigabit connection to the redundant peer. That is, the current VLAN's channel (or member) should connect to the redundant peer and have speeds of 1-gigabit or more. We also recommend that the connection be direct (without any intervening bridges or routers), and that the switches are co-located. If the latency is low, an intervening Gigabit L2 switch is permissible.

When redundancy is enabled and the ARX-1500 or ARX-2500 is using this interface as a redundant-pair link, the cfg-if-vlan no redundancy command causes the backup peer to reboot. The reboot does not disrupt any storage services, but the ARX peers cannot function as a redundant pair while the link is shut down. Additionally, a quorum-disk failure or disconnection would cause the active peer to reboot, too. If you proceed with removing the link, you should establish a new one as soon as possible: use the redundancy protocol, redundancy protocol (cfg-channel), or this command on another port, channel, or VLAN interface to establish a new redundant-pair link.

Samples

```
bstnA(cfg-if-vlan[555])# redundancy
```

allows the in-band-management interface on VLAN 555 to be used for redundancy rendezvous.

```
bstnA(cfg-if-vlan[8])# no redundancy
```

disallows rendezvous for the VLAN-8 interface.

Related Commands

interface vlan
peer
redundancy protocol
redundancy protocol (cfg-channel)

show arp

Purpose Address-Resolution Protocol (ARP) maps IP addresses to MAC addresses. Every network-connected processor on the ARX keeps a separate ARP table with its known IP/MAC entries. Use the show arp command to show one or more ARP tables.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show arp [all]
show arp from slot.processor [type {dynamic | static | local}]
```

all (optional) specifies all entries on the switch.

from *slot.processor* (optional) focuses on the ARP table at one processor. This option is not supported on the ARX-1500, ARX-2500, or ARX-VE:

slot (1-2 for ARX-4000; 1 for all others) is the slot number.

processor is the processor number. Use the show processors command to show all processors on the ARX, along with their associated module name(s) and status.

type {dynamic | static | local} (optional, if you use the **from** clause) selects one type of ARP-table entry, based on how the entry was learned. You must choose one of the following:

dynamic - learned from neighboring equipment.

static - specified by the arp command.

local - set internally by the switch.

Guidelines Proc is the processor (in *slot.processor* format). This only appears if you use an option, **all** or **from**, on a platform other than the ARX-1500, ARX-2500, or ARX-VE.

IP Address and

MAC Address are mapped together. If the MAC address is all zeros, the ARX could not find the IP address through ARP; likely the IP address does not exist in this case.

VLAN is the VLAN where the IP address was learned or specified.

Type is

- dynamic if the entry was learned from neighboring equipment,
- static if the entry was specified by the arp command, or
- local if the entry is defined by the switch software.

The summary output (from show arp, without any additional arguments) shows only dynamic-ARP entries.

Age (sec) is the time the entry has been in the ARP table, shown in seconds.

Use the arp command to create a static ARP entry. Use the clear arp command to clear all dynamic-ARP entries.

Samples

`bstnA(cfg)# show arp`

shows a summary of the ARP-table entries on the switch. This includes only dynamic-ARP entries. See Figure 10.1 for sample output.

`bstnA(cfg)# show arp all`

shows all ARP-table entries from all processors. See Figure 10.2 on page 10-29 for sample output.

`bstnA(cfg)# show arp from 1.1`

shows the ARP table for processor 1.1 only. See Figure 10.3 on page 10-31 for sample output.

`bstnA(cfg)# show arp from 1.1 type local`

shows the local ARP entries for processor 1.1. See Figure 10.4 on page 10-32 for sample output.

See Figure 10.5 on page 10-32 for sample output on the ARX-2000.

Related Commands

show processors
arp
clear arp

Figure 10.1 *Sample Output: show arp (ARX-4000)*

```
bstnA(cfg)# show arp

IP Address         MAC Address      VLAN    Type     Age(sec)
-------------------------------------------------------------
192.168.25.1     00:01:e8:5e:ea:1f   25    dynamic    0
192.168.25.2     00:01:e8:5e:ea:1f   25    dynamic    0
10.46.28.218     00:00:00:00:00:00   25    dynamic    0
10.1.1.1         00:01:e8:5e:ea:1f   mgmt  dynamic    0
169.254.42.64    00:0a:49:17:92:46   mgmt  dynamic    0
169.254.42.65    00:0a:49:17:92:47   mgmt  dynamic    0
169.254.42.66    00:0a:49:17:92:48   mgmt  dynamic    0
169.254.42.67    00:0a:49:17:92:49   mgmt  dynamic    0
169.254.42.68    00:0a:49:17:92:44   mgmt  dynamic    0
169.254.42.69    00:0a:49:17:92:45   mgmt  dynamic    0
169.254.42.70    00:0a:49:17:92:46   mgmt  dynamic    0
169.254.42.71    00:0a:49:17:92:47   mgmt  dynamic    0
169.254.42.72    00:0a:49:17:92:48   mgmt  dynamic    0
169.254.42.73    00:0a:49:17:92:49   mgmt  dynamic    0
169.254.42.74    00:0a:49:17:92:4a   mgmt  dynamic    0
169.254.42.75    00:0a:49:17:92:4b   mgmt  dynamic    0
169.254.42.82    00:0a:49:17:92:45   mgmt  dynamic    0
```

Figure 10.2 *Sample Output: show arp all (ARX-4000)*

```
bstnA(cfg)# show arp all

Proc  IP Address         MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------
1.1  169.254.68.74    00:0a:49:17:73:4a   mgmt  dynamic    0
```

```
1.1  169.254.68.66    00:0a:49:17:73:48   mgmt  dynamic    0
1.1  10.1.1.1         00:01:e8:5e:ea:1f   mgmt  dynamic    0
1.1  169.254.68.75    00:0a:49:17:73:4b   mgmt  dynamic    0
1.1  169.254.68.71    00:0a:49:17:73:47   mgmt  dynamic    0
1.1  169.254.68.73    00:0a:49:17:73:49   mgmt  dynamic    0
1.1  169.254.68.67    00:0a:49:17:73:49   mgmt  dynamic    0
1.1  169.254.68.69    00:0a:49:17:73:45   mgmt  dynamic    0
1.1  169.254.68.82    00:0a:49:17:73:45   mgmt  dynamic    0
1.1  169.254.68.68    00:0a:49:17:73:44   mgmt  dynamic    0
1.1  169.254.68.64    00:0a:49:17:73:46   mgmt  dynamic    0
1.1  169.254.68.65    00:0a:49:17:73:47   mgmt  dynamic    0
1.1  169.254.68.70    00:0a:49:17:73:46   mgmt  dynamic    0
1.1  169.254.68.72    00:0a:49:17:73:48   mgmt  dynamic    0
1.1  10.1.1.7         00:15:17:47:15:e9   mgmt  local      -
1.1  169.254.68.32    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.33    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.76    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.77    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.78    00:0a:49:17:73:09   mgmt  local      -


Proc  IP Address          MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
1.1  169.254.68.79    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.80    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.81    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.83    00:0a:49:17:73:09   mgmt  local      -
2.1  10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.1  192.168.25.5     00:0a:49:17:73:ff   25    local      -
2.2  10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.2  192.168.25.5     00:0a:49:17:73:ff   25    local      -
2.3  10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.3  192.168.25.5     00:0a:49:17:73:ff   25    local      -
2.4  10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.4  192.168.25.5     00:0a:49:17:73:ff   25    local      -
2.5  192.168.25.1     00:01:e8:5e:ea:1f   25    dynamic    0
2.5  192.168.25.2     00:01:e8:5e:ea:1f   25    dynamic    0
2.5  10.46.16.218     00:00:00:00:00:00   25    dynamic    0
2.5  192.168.25.15    00:0a:49:17:73:c0   25    local      -
2.5  192.168.25.12    00:0a:49:17:73:c0   25    local      -
2.5  192.168.25.141   00:0a:49:17:73:84   25    local      -
2.5  10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.5  192.168.25.5     00:0a:49:17:73:ff   25    local      -


Proc  IP Address          MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
2.6  192.168.25.15    00:0a:49:17:73:c0   25    local      -
2.6  192.168.25.12    00:0a:49:17:73:c0   25    local      -
2.6  192.168.25.142   00:0a:49:17:73:85   25    local      -
2.6  10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.6  192.168.25.5     00:0a:49:17:73:ff   25    local      -
2.7  192.168.25.15    00:0a:49:17:73:c0   25    local      -
2.7  192.168.25.12    00:0a:49:17:73:c0   25    local      -
2.7  192.168.25.31    00:0a:49:17:73:80   25    local      -
2.7  192.168.25.143   00:0a:49:17:73:86   25    local      -
2.7  10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.7  192.168.25.5     00:0a:49:17:73:ff   25    local      -
2.8  192.168.25.15    00:0a:49:17:73:c0   25    local      -
2.8  192.168.25.12    00:0a:49:17:73:c0   25    local      -
```

```
2.8  192.168.25.32    00:0a:49:17:73:81   25    local      -
2.8  192.168.25.144   00:0a:49:17:73:87   25    local      -
2.8  10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.8  192.168.25.5     00:0a:49:17:73:ff   25    local      -
2.9  192.168.25.15    00:0a:49:17:73:c0   25    local      -
2.9  192.168.25.12    00:0a:49:17:73:c0   25    local      -
2.9  192.168.25.33    00:0a:49:17:73:82   25    local      -


Proc  IP Address         MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
2.9  192.168.25.145   00:0a:49:17:73:88   25    local      -
2.9  10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.9  192.168.25.5     00:0a:49:17:73:ff   25    local      -
2.10 192.168.25.15    00:0a:49:17:73:c0   25    local      -
2.10 192.168.25.12    00:0a:49:17:73:c0   25    local      -
2.10 192.168.25.34    00:0a:49:17:73:83   25    local      -
2.10 192.168.25.146   00:0a:49:17:73:89   25    local      -
2.10 10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.10 192.168.25.5     00:0a:49:17:73:ff   25    local      -
2.11 192.168.25.15    00:0a:49:17:73:c0   25    local      -
2.11 192.168.25.12    00:0a:49:17:73:c0   25    local      -
2.11 192.168.25.147   00:0a:49:17:73:8a   25    local      -
2.11 10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.11 192.168.25.5     00:0a:49:17:73:ff   25    local      -
2.12 192.168.25.15    00:0a:49:17:73:c0   25    local      -
2.12 192.168.25.12    00:0a:49:17:73:c0   25    local      -
2.12 192.168.25.148   00:0a:49:17:73:8b   25    local      -
2.12 10.46.16.253     00:0a:49:17:73:fe   25    local      -
2.12 192.168.25.5     00:0a:49:17:73:ff   25    local      -
```

Figure 10.3 *Sample Output: show arp from 1.1 (ARX-4000)*

```
bstnA(cfg)# show arp from 1.1

Proc  IP Address         MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
1.1  169.254.68.74    00:0a:49:17:73:4a   mgmt  dynamic    0
1.1  169.254.68.66    00:0a:49:17:73:48   mgmt  dynamic    0
1.1  10.1.1.1         00:01:e8:5e:ea:1f   mgmt  dynamic    0
1.1  169.254.68.75    00:0a:49:17:73:4b   mgmt  dynamic    0
1.1  169.254.68.71    00:0a:49:17:73:47   mgmt  dynamic    0
1.1  169.254.68.73    00:0a:49:17:73:49   mgmt  dynamic    0
1.1  169.254.68.67    00:0a:49:17:73:49   mgmt  dynamic    0
1.1  169.254.68.69    00:0a:49:17:73:45   mgmt  dynamic    0
1.1  169.254.68.82    00:0a:49:17:73:45   mgmt  dynamic    0
1.1  169.254.68.68    00:0a:49:17:73:44   mgmt  dynamic    0
1.1  169.254.68.64    00:0a:49:17:73:46   mgmt  dynamic    0
1.1  169.254.68.65    00:0a:49:17:73:47   mgmt  dynamic    0
1.1  169.254.68.70    00:0a:49:17:73:46   mgmt  dynamic    0
1.1  169.254.68.72    00:0a:49:17:73:48   mgmt  dynamic    0
1.1  10.1.1.7         00:15:17:47:15:e9   mgmt  local      -
1.1  169.254.68.32    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.33    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.76    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.77    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.78    00:0a:49:17:73:09   mgmt  local      -
```

```
Proc  IP Address         MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
1.1  169.254.68.79    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.80    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.81    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.83    00:0a:49:17:73:09   mgmt  local      -
```

Figure 10.4 *Sample Output: show arp from 1.1 type local (ARX-4000)*

```
bstnA(cfg)# show arp from 1.1 type local

Proc  IP Address         MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
1.1  10.1.1.7         00:15:17:47:15:e9   mgmt  local      -
1.1  169.254.68.32    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.33    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.76    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.77    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.78    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.79    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.80    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.81    00:0a:49:17:73:09   mgmt  local      -
1.1  169.254.68.83    00:0a:49:17:73:09   mgmt  local      -
```

Figure 10.5 *Sample Output: show arp (ARX-2000)*

```
prtlndA(cfg)# show arp all

Proc  IP Address         MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
1.1  10.1.23.1         00:01:e8:5e:ea:1f   mgmt  dynamic    0
1.1  169.254.104.65    00:0a:49:17:b1:41   mgmt  dynamic    0
1.1  169.254.104.64    00:0a:49:17:b1:40   mgmt  dynamic    0
1.1  169.254.104.66    00:0a:49:17:b1:42   mgmt  dynamic    0
1.1  169.254.104.67    00:0a:49:17:b1:43   mgmt  dynamic    0
1.1  10.1.23.11        00:15:17:6b:a7:d9   mgmt  local      -
1.1  169.254.104.32    00:0a:49:17:b1:09   mgmt  local      -
1.1  169.254.104.33    00:0a:49:17:b1:09   mgmt  local      -
1.1  169.254.104.68    00:0a:49:17:b1:09   mgmt  local      -
1.1  169.254.104.69    00:0a:49:17:b1:09   mgmt  local      -
1.1  169.254.104.70    00:0a:49:17:b1:09   mgmt  local      -
1.1  169.254.104.71    00:0a:49:17:b1:09   mgmt  local      -
1.2  192.168.74.1      00:01:e8:5e:ea:1f   74    dynamic    0
1.2  192.168.74.2      00:01:e8:5e:ea:1f   74    dynamic    0
1.2  10.46.26.218      00:00:00:00:00:00   74    dynamic    0
1.2  192.168.74.92     00:0a:49:17:b1:c0   74    local      -
1.2  192.168.74.91     00:0a:49:17:b1:c0   74    local      -
1.2  192.168.74.41     00:0a:49:27:84:80   74    local      -
1.2  192.168.74.21     00:0a:49:17:b1:80   74    local      -
1.2  10.46.26.253      00:0a:49:17:b1:fe   74    local      -


Proc  IP Address         MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
1.2  192.168.74.66     00:0a:49:17:b1:ff   74    local      -
1.3  192.168.25.5      00:0a:49:17:73:ff   74    dynamic  1712
1.3  192.168.74.1      00:01:e8:5e:ea:1f   74    dynamic    0
1.3  192.168.74.2      00:01:e8:5e:ea:1f   74    dynamic    0
1.3  192.168.74.92     00:0a:49:17:b1:c0   74    local      -
1.3  192.168.74.91     00:0a:49:17:b1:c0   74    local      -
```

```
1.3  192.168.74.42    00:0a:49:27:84:81   74    local      -
1.3  192.168.74.22    00:0a:49:17:b1:81   74    local      -
1.3  10.46.26.253     00:0a:49:17:b1:fe   74    local      -
1.3  192.168.74.66    00:0a:49:17:b1:ff   74    local      -
1.4  192.168.74.92    00:0a:49:17:b1:c0   74    local      -
1.4  192.168.74.91    00:0a:49:17:b1:c0   74    local      -
1.4  192.168.74.43    00:0a:49:27:84:82   74    local      -
1.4  192.168.74.23    00:0a:49:17:b1:82   74    local      -
1.4  10.46.26.253     00:0a:49:17:b1:fe   74    local      -
1.4  192.168.74.66    00:0a:49:17:b1:ff   74    local      -
1.5  192.168.74.92    00:0a:49:17:b1:c0   74    local      -
1.5  192.168.74.91    00:0a:49:17:b1:c0   74    local      -
1.5  192.168.74.44    00:0a:49:27:84:83   74    local      -
1.5  192.168.74.24    00:0a:49:17:b1:83   74    local      -


Proc  IP Address         MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
1.5  10.46.26.253     00:0a:49:17:b1:fe   74    local      -
1.5  192.168.74.66    00:0a:49:17:b1:ff   74    local      -


prtlndA(cfg)# show arp from 1.1

Proc  IP Address         MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
1.1  10.1.23.1        00:01:e8:5e:ea:1f   mgmt  dynamic    0
1.1  169.254.104.65   00:0a:49:17:b1:41   mgmt  dynamic    0
1.1  169.254.104.64   00:0a:49:17:b1:40   mgmt  dynamic    0
1.1  169.254.104.66   00:0a:49:17:b1:42   mgmt  dynamic    0
1.1  169.254.104.67   00:0a:49:17:b1:43   mgmt  dynamic    0
1.1  10.1.23.11       00:15:17:6b:a7:d9   mgmt  local      -
1.1  169.254.104.32   00:0a:49:17:b1:09   mgmt  local      -
1.1  169.254.104.33   00:0a:49:17:b1:09   mgmt  local      -
1.1  169.254.104.68   00:0a:49:17:b1:09   mgmt  local      -
1.1  169.254.104.69   00:0a:49:17:b1:09   mgmt  local      -
1.1  169.254.104.70   00:0a:49:17:b1:09   mgmt  local      -
1.1  169.254.104.71   00:0a:49:17:b1:09   mgmt  local      -


prtlndA(cfg)# show arp from 1.1 type dynamic

Proc  IP Address         MAC Address      VLAN  Type     Age(sec)
-------------------------------------------------------------------
1.1  10.1.23.1        00:01:e8:5e:ea:1f   mgmt  dynamic    0
1.1  169.254.104.65   00:0a:49:17:b1:41   mgmt  dynamic    0
1.1  169.254.104.64   00:0a:49:17:b1:40   mgmt  dynamic    0
1.1  169.254.104.66   00:0a:49:17:b1:42   mgmt  dynamic    0
1.1  169.254.104.67   00:0a:49:17:b1:43   mgmt  dynamic    0
```

show interface

Purpose Use the show interface command to show the full configuration for all interfaces. Use show interface summary to see a single status line for each interface.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show interface [summary]`

summary (optional) reduces the output to a one-line summary for each interface.

Guidelines The default command, show interface, displays all interface configurations. The output is the same for each individual show command: show interface mgmt, show interface gigabit, show interface vlan, and show ron.

The show interface summary command shows the following fields for each interface:

Type is mgmt (the out-of-band MGMT interface, if it exists on this chassis), gbe (GigaBit Ethernet, an external port), or 10gbe (ten-Gigabit Ethernet, a faster external port on the ARX-4000).

Slot/Port shows the location of the interface. An asterisk (*) indicates that the interface is used as a redundant-pair link.

Admin State is Enabled or Disabled (also called "shut down"), as set by the administrator. To change this, use [no] shutdown in the interface's configuration mode: see shutdown (cfg-if-gig), shutdown (cfg-if-ten-gig), shutdown (cfg-mgmt), shutdown (cfg-if-vlan), or shutdown (cfg-if-vlan-ron-tnl). If the interface is a member of a channel, its administrative state is controlled by its channel: therefore, this shows "Ch *n*," where *n* is the channel number. You can use show channel to see the administrative state of a channel.

Link Status is the actual state of the interface (up or down).

Speed and

Duplex are both set by the speed (cfg-if-gig) command.

Description is set by the description command in the interface's config mode: description (cfg-mgmt) for fe, description (cfg-if-gig) for gbe, or description (cfg-if-ten-gig) for 10gbe.

Samples

`bstnA> show interface`

shows the configuration of every interface on the current ARX. See Figure 10.6 for sample output.

`prtlndA> show interface summary`

shows summaries of all interfaces on the "prtlndA" chassis. For sample output, see Figure 10.7 on page 10-40.

`canbyA> show interface summary`

shows summaries of all interfaces on the "canbyA" chassis, an ARX-1500. For sample output, see Figure 10.8 on page 10-41.

Related Commands

show interface gigabit
show interface ten-gigabit
show interface mgmt
show interface vlan
show ron
speed (cfg-if-gig)
description (cfg-mgmt)
description (cfg-if-gig)
description (cfg-if-ten-gig)

Figure 10.6 *Sample Output: show interface*

```
bstnA> show interface

Interface Type               management
    Slot                     1
    Interface                1
    Description
    Admin State              Enabled
    Link Status              Up
    Speed                    1  Gb/s
    Duplex                   Full
    Auto Negotiation         Enabled
    MAC Address     00:04:23:e2:bb:01
    MTU Size                 1500
Interface Type             10-gigabit
    Slot                     2
    Interface                1
    Description              Default
    Type                     10GBASE-SR X2
    Mode                     Normal
    Admin State              Enabled
    Link Status              Down
    Speed                    10 Gb/s
    Duplex                   Unknown
    Auto Negotiation(Admin)  Disabled
    Flow Control(Admin)
                 Receive     Off
                 Send        Off
    MAC Address              00:0a:49:17:78:32
    LACP Priority            32768
    Storm Control:Broadcast  1000 packets/sec
                  Multicast  1000 packets/sec
```

```
                Unknown DA 1000 packets/sec
    Port VLAN ID            0
    Accept Frames           Admit All
Interface Type            10-gigabit
    Slot                    2
    Interface               2
    Description             Default
    Type                    10GBASE-SR X2
    Mode                    Normal
    Admin State             Enabled
    Link Status             Down
    Speed                   10 Gb/s
    Duplex                  Unknown
    Auto Negotiation(Admin) Disabled
    Flow Control(Admin)
                Receive     Off
                Send        Off
    MAC Address             00:0a:49:17:78:33
    LACP Priority           32768
    Storm Control:Broadcast 1000 packets/sec
                Multicast   1000 packets/sec
                Unknown DA 1000 packets/sec
    Port VLAN ID            0
    Accept Frames           Admit All
Interface Type       gigabit
    Slot                    2
    Interface               3
    Description             Default
    Type                    Copper
    Mode                    Normal
    Admin State             Enabled
    Link Status             Down
    Speed                   1 Gb/s
    Duplex                  Unknown
    Auto Negotiation(Admin) Enabled
    Flow Control(Admin)
                Receive     Off
                Send        Off
    MAC Address             00:0a:49:17:78:34
    LACP Priority           32768
    Storm Control:Broadcast 1000 packets/sec
                Multicast   1000 packets/sec
                Unknown DA 1000 packets/sec
    Port VLAN ID            0
    Accept Frames           Admit All
Interface Type       gigabit
    Slot                    2
    Interface               4
    Description             Default
    Type                    Copper
    Mode                    Normal
    Admin State             Enabled
    Link Status             Down
    Speed                   1 Gb/s
    Duplex                  Unknown
    Auto Negotiation(Admin) Enabled
    Flow Control(Admin)
                Receive     Off
                Send        Off
    MAC Address             00:0a:49:17:78:35
    LACP Priority           32768
    Storm Control:Broadcast 1000 packets/sec
```

```
                Multicast  1000 packets/sec
                Unknown DA 1000 packets/sec
    Port VLAN ID            0
    Accept Frames           Admit All
Interface Type       gigabit
    Slot                    2
    Interface               5
    Description             Default
    Type                    Copper
    Mode                    Normal
    Admin State             Enabled
    Link Status             Up
    Speed                   1 Gb/s
    Duplex                  Full
    Auto Negotiation(Admin) Enabled
    Flow Control(Admin)
                Receive     Off
                Send        Off
    MAC Address             00:0a:49:17:78:36
    LACP Priority           32768
    Storm Control:Broadcast 1000 packets/sec
                Multicast  1000 packets/sec
                Unknown DA 1000 packets/sec
    Port VLAN ID            25
    Accept Frames           Admit All
Interface Type       gigabit
    Slot                    2
    Interface               6
    Description             Default
    Type                    Copper
    Mode                    Normal
    Admin State             Enabled
    Link Status             Up
    Speed                   1 Gb/s
    Duplex                  Full
    Auto Negotiation(Admin) Enabled
    Flow Control(Admin)
                Receive     Off
                Send        Off
    MAC Address             00:0a:49:17:78:37
    LACP Priority           32768
    Storm Control:Broadcast 1000 packets/sec
                Multicast  1000 packets/sec
                Unknown DA 1000 packets/sec
    Port VLAN ID            25
    Accept Frames           Admit All
Interface Type       gigabit
    Slot                    2
    Interface               7
    Description             Default
    Type                    Copper
    Mode                    Normal
    Admin State             Enabled
    Link Status             Down
    Speed                   1 Gb/s
    Duplex                  Unknown
    Auto Negotiation(Admin) Enabled
    Flow Control(Admin)
                Receive     Off
                Send        Off
    MAC Address             00:0a:49:17:78:38
    LACP Priority           32768
```

```
    Storm Control:Broadcast  1000 packets/sec
                  Multicast  1000 packets/sec
                  Unknown DA 1000 packets/sec
    Port VLAN ID             0
    Accept Frames            Admit All
Interface Type       gigabit
    Slot                     2
    Interface                8
    Description              Default
    Type                     Copper
    Mode                     Normal
    Admin State              Enabled
    Link Status              Down
    Speed                    1 Gb/s
    Duplex                   Unknown
    Auto Negotiation(Admin)  Enabled
    Flow Control(Admin)
                 Receive     Off
                 Send        Off
    MAC Address              00:0a:49:17:78:39
    LACP Priority            32768
    Storm Control:Broadcast  1000 packets/sec
                  Multicast  1000 packets/sec
                  Unknown DA 1000 packets/sec
    Port VLAN ID             0
    Accept Frames            Admit All
Interface Type       gigabit
    Slot                     2
    Interface                9
    Description              Default
    Type                     Copper
    Mode                     Normal
    Admin State              Enabled
    Link Status              Down
    Speed                    1 Gb/s
    Duplex                   Unknown
    Auto Negotiation(Admin)  Enabled
    Flow Control(Admin)
                 Receive     Off
                 Send        Off
    MAC Address              00:0a:49:17:78:3a
    LACP Priority            32768
    Storm Control:Broadcast  1000 packets/sec
                  Multicast  1000 packets/sec
                  Unknown DA 1000 packets/sec
    Port VLAN ID             0
    Accept Frames            Admit All
Interface Type       gigabit
    Slot                     2
    Interface                10
    Description              Default
    Type                     Copper
    Mode                     Normal
    Admin State              Enabled
    Link Status              Down
    Speed                    1 Gb/s
    Duplex                   Unknown
    Auto Negotiation(Admin)  Enabled
    Flow Control(Admin)
                 Receive     Off
                 Send        Off
    MAC Address              00:0a:49:17:78:3b
```

```
    LACP Priority            32768
    Storm Control:Broadcast  1000 packets/sec
                  Multicast  1000 packets/sec
                  Unknown DA 1000 packets/sec
    Port VLAN ID             0
    Accept Frames            Admit All
Interface Type       gigabit
    Slot                     2
    Interface                11
    Description              Default
    Type                     Copper
    Mode                     Normal
    Admin State              Enabled
    Link Status              Down
    Speed                    1 Gb/s
    Duplex                   Unknown
    Auto Negotiation(Admin)  Enabled
    Flow Control(Admin)
                 Receive     Off
                 Send        Off
    MAC Address              00:0a:49:17:78:3c
    LACP Priority            32768
    Storm Control:Broadcast  1000 packets/sec
                  Multicast  1000 packets/sec
                  Unknown DA 1000 packets/sec
    Port VLAN ID             0
    Accept Frames            Admit All
Interface Type       gigabit
    Slot                     2
    Interface                12
    Description              Default
    Type                     Copper
    Mode                     Normal
    Admin State              Enabled
    Link Status              Down
    Speed                    1 Gb/s
    Duplex                   Unknown
    Auto Negotiation(Admin)  Enabled
    Flow Control(Admin)
                 Receive     Off
                 Send        Off
    MAC Address              00:0a:49:17:78:3d
    LACP Priority            32768
    Storm Control:Broadcast  1000 packets/sec
                  Multicast  1000 packets/sec
                  Unknown DA 1000 packets/sec
    Port VLAN ID             0
    Accept Frames            Admit All
Interface Type       gigabit
    Slot                     2
    Interface                13
    Description              Default
    Type                     Copper
    Mode                     Normal
    Admin State              Disabled
    Link Status              Down
    Speed                    1 Gb/s
    Duplex                   Unknown
    Auto Negotiation(Admin)  Enabled
    Flow Control(Admin)
                 Receive     Off
                 Send        Off
```

```
    MAC Address              00:0a:49:17:78:3e
    LACP Priority            32768
    Storm Control:Broadcast  1000 packets/sec
                  Multicast  1000 packets/sec
                  Unknown DA 1000 packets/sec
    Port VLAN ID             0
    Accept Frames            Admit All
Interface Type       gigabit
    Slot                     2
    Interface                14
    Description              Default
    Type                     Copper
    Mode                     Normal
    Admin State              Disabled
    Link Status              Down
    Speed                    1 Gb/s
    Duplex                   Unknown
    Auto Negotiation(Admin)  Enabled
    Flow Control(Admin)
                 Receive     Off
                 Send        Off
    MAC Address              00:0a:49:17:78:3f
    LACP Priority            32768
    Storm Control:Broadcast  1000 packets/sec
                  Multicast  1000 packets/sec
                  Unknown DA 1000 packets/sec
    Port VLAN ID             0
    Accept Frames            Admit All




Interface Type               vlan

 Vlan   Admin      IP Address        Subnet Mask       Description
 ----   --------   ---------------   ---------------   -------------------
  25    Enabled    192.168.25.5      255.255.0.0



bstnA>
```

Figure 10.7 *Sample Output: show interface summary (ARX-2000)*

```
prtlndA> show interface summary

Type   Slot/ Admin    Link   Speed   Duplex  Description
       Port  State    Status
-----  ----  -------  ------ ------- ------- ----------------------------------
mgmt   1/1   Enabled  Up     1  Gb/s Full
gbe    1/1   Enabled  Down   1  Gb/s Unknown
gbe    1/2   Enabled  Down   1  Gb/s Unknown
gbe    1/3   Enabled  Down   1  Gb/s Unknown
gbe    1/4   Enabled  Down   1  Gb/s Unknown
gbe    1/5   Enabled  Up     1  Gb/s Full
gbe    1/6   Enabled  Up     1  Gb/s Full
gbe    1/7   Disabled Down   1  Gb/s Unknown
gbe    1/8   Disabled Down   1  Gb/s Unknown
gbe    1/9   Disabled Down   1  Gb/s Unknown
gbe    1/10  Disabled Down   1  Gb/s Unknown
gbe    1/11  Disabled Down   1  Gb/s Unknown
```

```
gbe    1/12* Enabled  Up     1  Gb/s Full

* Redundancy-Interface

prtlndA>
```

Figure 10.8 *Sample Output: show interface summary (ARX-1500)*

```
canbyA> show interface summary

Type   Slot/ Admin    Link   Speed   Duplex  Description
       Port  State    Status
-----  ----  -------  ------ ------- ------- -----------------------------------

mgmt   1/1   Enabled  Up     Auto-N  Full
gbe    1/2   Enabled  Down   1  Gb/s Unknown Default
gbe    1/3   Disabled Down   1  Gb/s Unknown Default
gbe    1/4   Disabled Down   1  Gb/s Unknown Default
gbe    1/5   Disabled Down   1  Gb/s Unknown Default
gbe    1/6   Disabled Down   1  Gb/s Unknown Default
gbe    1/7   Enabled  Up     100Mb/s Full    Default
gbe    1/8   Enabled  Up     100Mb/s Full    Default
```

show interface mgmt

Purpose Administrators can log into the CLI or GUI through the out-of-band management interface, typically labeled "MGMT." Use the show interface mgmt command to show the configuration and status of the out-of-band management interface.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show interface mgmt [stats]`

stats (optional) displays statistics for the interface.

Platforms any *except* ARX-VE

Guidelines The default output contains the following fields:

Slot is always "1," and

Interface is also always "1." The 1/1 interface is the MGMT interface for all platforms.

Description is an optional description, set by the description (cfg-mgmt) command.

Admin Status shows whether or not the interface is administratively enabled. You can disable (or restart) this interface with the shutdown (cfg-mgmt) command.

Link Status is the actual state of the interface (up or down).

Speed,

Duplex, and

Auto Negotiation are all set by the speed (cfg-mgmt) command.

MAC Address is the MAC for the MGMT interface.

MTU Size shows the Maximum Transmission Unit, or maximum packet size, for this interface.

IP Address, and

Subnet Mask define the management address and subnet. You can set these with the ip address (cfg-mgmt) command.

The stats output is a table of counters, separated into Ingress and Egress counts. These counts restart when the ARX reboots; use the reload command to reboot the ARX. On an ARX-1500 or an ARX-2500, you can use the clear counters mgmt command to clear the stats, without performing a full reboot.

Use the show interface vlan command to list all of the VLAN-based in-band management interfaces.

Samples

`bstnA# show interface mgmt`

shows the configuration and status of the out-of-band management interface. See Figure 10.9 for sample output.

`bstnA# show interface mgmt stats`

shows the statistics for the same interface. Sample output appears in Figure 10.10.

Related Commands

description (cfg-mgmt)
ip address (cfg-mgmt)
speed (cfg-mgmt)
shutdown (cfg-mgmt)
show interface vlan
clear counters mgmt

Figure 10.9 *Sample Output: show interface mgmt*

```
bstnA# show interface mgmt

Slot                    1
Interface               1
Description
Admin Status            Enabled
Link Status             Up
Speed                   1  Gb/s
Duplex                  Full
Auto Negotiation        Disabled
MAC Address             00:04:23:e2:9f:95
MTU Size                1500
IP Address              10.1.1.7
Subnet Mask             255.255.255.0
```

Figure 10.10 *Sample Output: show interface mgmt stats*

```
bstnA# show interface mgmt stats

Slot                         1
Interface                    1
-----------------------------------------------------------------
                           Ingress                  Egress
-----------------------------------------------------------------

Octets                     3048207                 6532008
Total Frames                 39131                   35497
Dropped Frames                   0                       0
Error Frames                     0                       0
FIFO Errors                      0                       0
Multicast Frames                38
CRC Errors                       0
Symbol Errors                    0
Oversize Errors                  0
Frame Errors                     0
Length Errors                    0
Alignment Errors                 0
Missed Frames                    0
Collision  Frames                                        0
```

show interface vlan

Purpose You can configure one in-band-management interface per VLAN. Administrators on the VLAN can log into the CLI through this interface. Use the show interface vlan command to show the configuration for all in-band management interfaces.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show interface vlan`**

Guidelines Vlan is the VLAN ID for the interface.

Admin shows whether or not the interface is administratively enabled. You can disable (or restart) an in-band-management interface with the shutdown (cfg-if-vlan) command.

IP Address, and

Subnet Mask define the management address and subnet. You can set these with the ip address (cfg-if-vlan) command.

Description is an optional description, set by the description (cfg-if-vlan) command.

Use the show interface mgmt command to show the configuration for the single out-of-band management interface.

Samples `bstnA(cfg)# show interface vlan`
shows a summary of configured VLANs. See the sample output below.

Related Commands description (cfg-if-vlan)
ip address (cfg-if-vlan)
shutdown (cfg-if-vlan)

Figure 10.11 *Sample Output: show interface vlan*

```
bstnA(cfg)# show interface vlan

 Vlan   Admin      IP Address        Subnet Mask       Description
 ----   --------   ---------------   ---------------   -------------------
  25    Enabled    192.168.25.5      255.255.255.0
```

show ip address

Purpose Use the show ip address command to show configuration details for a Proxy IP, private IP, VIP, or some other IP address on the ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show ip address`** ***`ip-address`***

ip-address is the desired IP address (for example, 192.168.25.56).

Guidelines Slot ID shows the location of the module that processes all packets for the address.

Processor is the network processor that serves the address. Every IP address is assigned to a single network processor. Use the show processors command to show all processors.

Type is any of the following:

- External is an IP address not owned by the switch, for example, the IP address of a router or a back-end filer.
- Proxy is a proxy-IP address. Use the show ip proxy-addresses command to show all proxy IPs.
- VIP is a virtual-IP address. Use the show virtual service command to show all Virtual IPs.
- management is the out-of-band-management interface. Use the show interface mgmt command to show this interface.
- VLAN is an in-band (VLAN-based) management interface. Use the show interface vlan command to list these interfaces.

MAC Address is the MAC for the IP.

VLAN ID is the VLAN for IP's subnet.

Sample

```
prtlndA(cfg)# show ip address 192.168.74.91

Report for 0.0.0.0
Slot ID    :1
Processor  :2
Type       :VIP
MAC Address:00:0a:49:17:b1:c0
VLAN ID    :74
```

Related Commands

show processors
show ip proxy-addresses
show virtual service
show interface mgmt
show interface vlan

show ip domain

Purpose The ARX can perform DNS lookups to translate IP addresses (for example, "172.16.36.55") into FQDNs (for example, "www.mycompany.com"). Use the show ip domain command to show the current configuration for DNS lookups.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show ip domain`**

Guidelines Domain List is analogous to the *search* list in resolv.conf. To edit this, use the ip domain-list command.

Name Servers is analogous to the *nameserver* list in resolv.conf. To edit this, use the ip name-server command.

Sample

```
bstnA(cfg)# show ip domain

DNS Server Configuration
   Domain List:    wwmed.com medarch.org bigorg.org
   Name Servers:   192.168.25.201 192.168.25.202 192.168.25.209
```

shows the current configuration for DNS lookups.

Related Commands ip domain-list
ip name-server

show ip proxy-addresses

Purpose Every network processor on the ARX has a *proxy IP* address, used as a home address for communication with filers and servers on the back end. Use the show ip proxy-addresses command to show all configured proxy IPs.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show ip proxy-addresses`**

Guidelines This command displays a table of proxy IP addresses, one per row. For a redundant pair of switches, this shows all proxy IPs for both peers.

Proxy Address is the IP address.

VLAN is the VLAN where the proxy IP resides.

MAC Address is the layer-2-MAC address for the proxy IP.

Owner is the chassis where the proxy IP was configured. This is relevant in a redundant-switch configuration.

In Use By is the chassis that is currently using the proxy IP. In a redundancy failover, the surviving chassis assumes all proxy IPs from the failed chassis.

Proc identifies the network processor that is using the proxy IP, in *slot.processor* format.

Use the ip proxy-address command to add a range of proxy-IP addresses.

Samples `bstnA# show ip proxy-addresses`

See Figure 10.12 on page 10-47 for sample output from a non-redundant peer.

`prtlndA# show ip proxy-addresses`

See Figure 10.12 on page 10-47 for sample output from a redundant peer.

Related Commands ip proxy-address
show processors

Figure 10.12 *Sample Output: show ip proxy-addresses*

```
bstnA# show ip proxy-addresses

   Proxy Address    VLAN     Mac Address       Owner     In Use By    Proc
-----------------  ----  -----------------   -------   -----------   ----
   192.168.25.31/24  25   00:0a:49:17:78:80    bstnA       bstnA       2.1
   192.168.25.32/24  25   00:0a:49:17:78:81    bstnA       bstnA       2.2
   192.168.25.33/24  25   00:0a:49:17:78:82    bstnA       bstnA       2.3
   192.168.25.34/24  25   00:0a:49:17:78:83    bstnA       bstnA       2.4
   192.168.25.141/24 25   00:0a:49:17:78:84    bstnA       bstnA       2.5
   192.168.25.142/24 25   00:0a:49:17:78:85    bstnA       bstnA       2.6
   192.168.25.143/24 25   00:0a:49:17:78:86    bstnA       bstnA       2.7
   192.168.25.144/24 25   00:0a:49:17:78:87    bstnA       bstnA       2.8
   192.168.25.145/24 25   00:0a:49:17:78:88    bstnA       bstnA       2.9
   192.168.25.146/24 25   00:0a:49:17:78:89    bstnA       bstnA       2.10
   192.168.25.147/24 25   00:0a:49:17:78:8a    bstnA       bstnA       2.11
   192.168.25.148/24 25   00:0a:49:17:78:8b    bstnA       bstnA       2.12
```

Figure 10.13 *Sample Output: show ip proxy-addresses (Redundant Pair)*

```
prtlndA# show ip proxy-addresses

   Proxy Address    VLAN     Mac Address       Owner     In Use By    Proc
------------------  ----  -----------------   -------   -----------   ----
   192.168.74.21/24  74   00:0a:49:17:b1:80    prtlndA     prtlndA     1.2
   192.168.74.22/24  74   00:0a:49:17:b1:81    prtlndA     prtlndA     1.3
   192.168.74.23/24  74   00:0a:49:17:b1:82    prtlndA     prtlndA     1.4
   192.168.74.24/24  74   00:0a:49:17:b1:83    prtlndA     prtlndA     1.5
   192.168.74.41/24       00:0a:49:17:a1:80    prtlndB     prtlndA     1.2
   192.168.74.42/24       00:0a:49:17:a1:81    prtlndB     prtlndA     1.3
   192.168.74.43/24       00:0a:49:17:a1:82    prtlndB     prtlndA     1.4
   192.168.74.44/24       00:0a:49:17:a1:83    prtlndB     prtlndA     1.5
```

show ip route

Purpose Use the show ip route command to show the active and static routes on the switch.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show ip route [all | from slot.processor | monitor]`

all (optional) specifies all routes on the switch.

from *slot.processor* (optional) specifies routes from one network processor. This option is not supported on the ARX-1500, ARX-2500, or ARX-VE:

slot (1-2 on an ARX-4000, 1 on any other platform) is the slot number, and

processor is the processor number. Use the show processors command to show a complete list of processors.

monitor (optional) shows the status of the next-hop gateway, and how the ARX is using it.

Guidelines Every network processor on the ARX has its own IP-routing table. (In addition to the processors behind the client/server Ethernet ports, this includes processor 1.1 because it connects to the out-of-band network.) The show ip route all command shows the routing tables for all processors. Many routes are common to all processors; duplicate routes are removed from the summary command, show ip route.

Use the ip route command to create a static route.

All versions of this command except show ip route monitor have the following fields:

Proc is the number of the module processor with the static route, in *slot.processor* format. This only appears if you use an option, **all** or **from**, on a platform other than the ARX-1500, ARX-2500, or ARX-VE. Each processor has a separate routing table. Use the show processors command to show a complete list of processors and their slots.

Destination/Mask defines the subnet for a route in CIDR format.

Gateway is the next-hop router for the route.

Cost is the relative cost of this route, weighed against any other routes to the same destination. A lower-cost metric is preferred.

Interface is the interface or VLAN that carries packets to this subnet. “Mgmt” is the out-of-band management interface, labeled MGMT on the front panel.

Age is the time (in seconds) that the ARX has been in continuous contact with the Gateway. The ARX uses periodic ARP requests to monitor the gateway while the route is active; the route is declared “Offline” if it fails to respond to ARPs. “Direct” means that the route is directly-connected to the interface. “Unacquired” applies to a per-VLAN-default route (created with ip route ... per-vlan), where the gateway has not yet responded to ARP requests. “Static” means that the per-VLAN route has responded to an ARP request and is evidently reachable.

Guidelines: show ip route monitor

The show ip route monitor command has similar output to show ip route, but it focuses on the status of the connections. It contains the following fields:

Destination/Mask defines the subnet for a route in CIDR format.

Type is the general destination for the route. This is either "VLAN" (indicating a client/server VLAN) or "Mgmt" (the out-of-band management subnet). "Mgmt" routes are created with a special flag in the ip route command.

Gateway is the next-hop router for the route.

Cost is the relative cost of this route, weighed against any other routes to the same destination. A lower-cost metric is preferred.

Status and

Details indicate whether or not this route is in use. If it is, this indicates how the switch uses the next-hop gateway. If not, it indicates the problem. Here are the possible conditions in this field:

- Up/Current Gateway indicates that the switch uses this route (as opposed to another static route to the same subnet).
- VLAN/Current Gateway applies to a per-VLAN-default route, created with the ip route ... per-vlan variant of the ip route command. This indicates that VLAN-default route is configured.
- Up/Backup Gateway is a gateway that is available to reach the subnet, but is not being used because some other, preferred gateway is being used instead. Another route should be in the same output with the same Destination/Mask, a different Gateway, and Status/Details of Up/Current Gateway. If this subnet's Current Gateway stops responding to ARP packets, this route is a candidate to take its place.
- Down/No Reply shows that the gateway is reachable through layer 2 and ICMP, but is not replying to IP packets.
- Down/Unreachable says that the gateway is not on a reachable subnet. Use ip route to reset the gateway. A gateway to file servers must be on the same subnet as the proxy-IP addresses (show ip proxy-addresses shows all such addresses), a gateway to clients must be on the same subnet as the clients' VIP (use show virtual service for a list of all VIPs), and a gateway to stations in the out-of-band (OOB) management network must be in the same subnet as the OOB management address (show interface mgmt).

Samples

`bstnA> show ip route`

lists the static routes on the switch by IP address. See Figure 10.14 on page 10-51 for sample output. For sample output on a different platform, see Figure 10.18 on page 10-53.

`bstnA> show ip route monitor`

displays the current status of all of the above routes. See Figure 10.15 on page 10-51 for sample output. For sample output on a different platform, see Figure 10.19 on page 10-53.

`bstnA> show ip route all`

shows all static routes for all processors on the switch. See Figure 10.16 on page 10-51 for sample output. For sample output on a different platform, see Figure 10.20 on page 10-53.

`bstnA> show ip route from 2.1`

shows the static routes from processor 1 on module 2 (the NSM, or Data Plane). See Figure 10.17 on page 10-53 for sample output. For sample output on a different platform, see Figure 10.21 on page 10-54.

Related Commands

show processors
ip route

Figure 10.14 *Sample Output: show ip route*

```
bstnA> show ip route

Destination/Mask    Gateway          Cost Interface        Age
-------------------------------------------------------------------
0.0.0.0/0           192.168.25.1     128  VLAN25          2489
0.0.0.0/0           10.1.1.1         128  Mgmt            2365
192.168.25.0/24     0.0.0.0          0    VLAN25          Direct
192.168.25.0/24     0.0.0.0          128  VLAN            Direct
192.168.78.0/24     0.0.0.0          128  VLAN            Direct
192.168.78.0/24     192.168.25.2     128  VLAN25          2489
10.1.1.0/24         0.0.0.0          0    Mgmt            Direct
```

Figure 10.15 *Sample Output: show ip route monitor*

```
bstnA> show ip route monitor

Destination/Mask    Type  Gateway         Cost  Status  Details
-------------------------------------------------------------------------------
0.0.0.0/0           Mgmt  10.1.1.1        128   Up      Current Gateway
0.0.0.0/0           VLAN  192.168.25.1    128   Up      Current Gateway
```

Figure 10.16 *Sample Output: show ip route all*

```
bstnA> show ip route all

Proc Destination/Mask    Gateway          Cost Interface        Age
-------------------------------------------------------------------
1.1  192.168.78.0/24     0.0.0.0          128  VLAN            Direct
1.1  10.1.1.0/24         0.0.0.0          0    Mgmt            Direct
1.1  192.168.25.0/24     0.0.0.0          128  VLAN            Direct
1.1  192.168.25.0/24     0.0.0.0          128  VLAN            Direct
1.1  10.1.1.0/24         0.0.0.0          0    Mgmt            Direct
```

```
1.1  0.0.0.0/0           10.1.1.1         128  Mgmt           2365
2.1  0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.1  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.1  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.1  192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.2  0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.2  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.2  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.2  192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.3  0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.3  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.3  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.3  192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.4  0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.4  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct


Proc Destination/Mask    Gateway          Cost Interface        Age
-------------------------------------------------------------------
2.4  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.4  192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.5  0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.5  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.5  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.5  192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.6  0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.6  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.6  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.6  192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.7  0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.7  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.7  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.7  192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.8  0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.8  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.8  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.8  192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.9  0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.9  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct


Proc Destination/Mask    Gateway          Cost Interface        Age
-------------------------------------------------------------------
2.9  192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.9  192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.10 0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.10 192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.10 192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.10 192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.11 0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.11 192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.11 192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.11 192.168.78.0/24     192.168.25.2     128  VLAN25         2489
2.12 0.0.0.0/0           192.168.25.1     128  VLAN25         2489
2.12 192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.12 192.168.25.0/24     0.0.0.0          0    VLAN25         Direct
2.12 192.168.78.0/24     192.168.25.2     128  VLAN25         2489
```

Figure 10.17 *Sample Output: show ip route from 2.1*

```
bstnA> show ip route from 2.1

Proc Destination/Mask    Gateway          Cost Interface        Age
------------------------------------------------------------------------
2.1  0.0.0.0/0           192.168.25.1     128  VLAN25           2489
2.1  192.168.25.0/24     0.0.0.0          0    VLAN25           Direct
2.1  192.168.25.0/24     0.0.0.0          0    VLAN25           Direct
2.1  192.168.78.0/24     192.168.25.2     128  VLAN25           2489
```

Figure 10.18 *Sample Output: show ip route (ARX-2000)*

```
prtlndA> show ip route

Destination/Mask    Gateway          Cost Interface        Age
------------------------------------------------------------------------
0.0.0.0/0           192.168.74.1     128  VLAN74           6794
0.0.0.0/0           10.1.23.1        128  VLAN             9230
192.168.74.0/24     0.0.0.0          0    VLAN74           Direct
192.168.74.0/24     0.0.0.0          128  VLAN             Direct
192.168.74.0/24     0.0.0.0          0    VLAN74           Direct
192.168.74.0/24     0.0.0.0          128  VLAN             Direct
192.168.78.0/24     0.0.0.0          128  VLAN             Direct
192.168.78.0/24     192.168.74.2     128  VLAN74           6799
10.1.23.0/24        0.0.0.0          0    VLAN             Direct
```

Figure 10.19 *Sample Output: show ip route monitor (ARX-2000)*

```
prtlndA> show ip route monitor

Destination/Mask    Type  Gateway         Cost  Status  Details
-------------------------------------------------------------------------------
0.0.0.0/0           Mgmt  10.1.23.1       128   Up      Current Gateway
0.0.0.0/0           VLAN  192.168.74.1    128   Up      Current Gateway
192.168.78.0/24     VLAN  192.168.74.2    128   Up      Current Gateway
```

Figure 10.20 *Sample Output: show ip route all (ARX-2000)*

```
prtlndA> show ip route all

Proc Destination/Mask    Gateway          Cost Interface        Age
------------------------------------------------------------------------
1.1  192.168.78.0/24     0.0.0.0          128  VLAN             Direct
1.1  10.1.23.0/24        0.0.0.0          0    VLAN             Direct
1.1  192.168.74.0/24     0.0.0.0          128  VLAN             Direct
1.1  192.168.74.0/24     0.0.0.0          128  VLAN             Direct
1.1  10.1.23.0/24        0.0.0.0          0    VLAN             Direct
1.1  0.0.0.0/0           10.1.23.1        128  VLAN             9230
1.2  0.0.0.0/0           192.168.74.1     128  VLAN74           6794
1.2  192.168.74.0/24     0.0.0.0          0    VLAN74           Direct
1.2  192.168.78.0/24     192.168.74.2     128  VLAN74           6799
1.2  192.168.74.0/24     0.0.0.0          0    VLAN74           Direct
1.3  0.0.0.0/0           192.168.74.1     128  VLAN74           6794
1.3  192.168.74.0/24     0.0.0.0          0    VLAN74           Direct
1.3  192.168.78.0/24     192.168.74.2     128  VLAN74           6799
1.3  192.168.74.0/24     0.0.0.0          0    VLAN74           Direct
1.4  0.0.0.0/0           192.168.74.1     128  VLAN74           6794
1.4  192.168.74.0/24     0.0.0.0          0    VLAN74           Direct
1.4  192.168.78.0/24     192.168.74.2     128  VLAN74           6799
1.4  192.168.74.0/24     0.0.0.0          0    VLAN74           Direct
1.5  0.0.0.0/0           192.168.74.1     128  VLAN74           6794
1.5  192.168.74.0/24     0.0.0.0          0    VLAN74           Direct
```

```
Proc Destination/Mask    Gateway          Cost Interface        Age
------------------------------------------------------------------------
1.5  192.168.78.0/24     192.168.74.2     128  VLAN74           6799
1.5  192.168.74.0/24     0.0.0.0          0    VLAN74           Direct
```

Figure 10.21 *Sample Output: show ip route from 1.1 (ARX-2000)*

```
prtlndA> show ip route from 1.1

Proc Destination/Mask    Gateway          Cost Interface        Age
------------------------------------------------------------------------
1.1  192.168.78.0/24     0.0.0.0          128  VLAN             Direct
1.1  10.1.23.0/24        0.0.0.0          0    Mgmt             Direct
1.1  192.168.74.0/24     0.0.0.0          128  VLAN             Direct
1.1  192.168.74.0/24     0.0.0.0          128  VLAN             Direct
1.1  10.1.23.0/24        0.0.0.0          0    Mgmt             Direct
1.1  0.0.0.0/0           10.1.23.1        128  Mgmt             750
```

show ntp servers

Purpose Use the show ntp servers command to display all configured NTP servers.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show ntp servers`

Guidelines This command displays a table of NTP servers. For each server, the table shows a numeric ID, the NTP protocol for the server (3 for NTPv3 or 4 for SNTPv4), the number of seconds between Polls, and the server's IP address.

Use the ntp server command to add an NTP server to the list. Use the show clock command to verify that the switch is getting the correct time from the NTP server(s).

Related Commands ntp server
show clock

Figure 10.22 *Sample Output: show ntp servers*

```
bstnA(cfg)# show ntp servers

Configured NTP Servers

ID        Proto    Poll       Address
-------------------------------------------------------------
2         4         64        192.168.25.201
1         4         64        192.168.25.209
```

show ntp status

Purpose Use the show ntp status command to display operational status for each configured NTP server.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show ntp status`

Guidelines This command displays operational status information for each configured NTP server, one row per server. The fields match the output of the **pe[er]** command on the Unix *ntpq* program, or **ntpq -p** on Windows:

- (tally code) - the one-character tally code, representing the server's current status (see Table 10.1, "Tally Codes," on page 57),
- remote - the remote NTP server's hostname or IP address,
- refid - the IP address or hostname of the server's reference clock (another NTP server; 0.0.0.0 if unknown),
- st - the stratum of the server (1-16, where 1 is ideal, 15 is worst, and 16 means "unusable"),
- t - the type of the NTP server (local, unicast, multicast, or broadcast) when the last packet was received,
- when - the time of the last received packet,
- poll - the number of seconds between polls,
- reach - the reachability register, in octal,
- delay - the interval (in milliseconds) to be added to requests that require authentication,
- offset - the number of milliseconds between the server's clock and the ARX clock, and
- jitter - the estimated time error of the server clock, measured as an exponential average of RMS time differences.

Use the show ntp servers command to get a list of all configured NTP servers. Use the show clock command to verify that the switch is getting the correct time from the NTP server(s).

Related Commands show ntp servers
show clock

Figure 10.23 *Sample Output: show ntp status*

```
bstnA(cfg)# show ntp status

NTP server status

     remote           refid      st t when poll reach   delay   offset  jitter
------------------------------------------------------------------------------
+lager.wwmed     navobs1.mit.edu  2 -   14   64  377    0.640  -147.94   6.475
*dc1.wwmed.c     lager.wwmed      2 -   15   64  377    0.411  145.091  13.112
```

Table 10.1 *Tally Codes*

*	sys.peer - The local switch has declared the peer the “system peer;” the peer lends its variables to the system variables.
o	pps.peer - The peer has been declared the system peer and lends its variables to the system variables. However, the actual system synchronization is derived from a pulse-per-second (PPS) signal, either indirectly via the PPS reference clock driver or directly via kernel interface.
+	candidate - The peer is a survivor and a candidate for the combining algorithm.
#	selected - The peer is a survivor, but not among the first six peers sorted by synchronization distance. If the association is ephemeral, it may be demobilized to conserve resources.
.	excess - The peer is discarded as not among the first ten peers sorted by synchronization distance. It is probably a poor candidate for further consideration.
<Space>	reject - The peer is discarded as unreachable, synchronized to this server (synch loop) or excessive synchronization distance.
x	falsetick - The peer is discarded by the intersection algorithm as a falseticker.
-	outlyer - The clustering algorithm designated the peer as an outlyer and discarded it.

shutdown (cfg-if-vlan)

Purpose From cfg-if-vlan mode, use the shutdown command to shut down the in-band management interface for a VLAN.

> **Important**
>
> *In a redundant pair of ARXes, the network software uses an in-band (VLAN) management address as a home address for its communication with the quorum-disk. Without an in-band-management address and an ip route to the quorum disk, a failover is impossible. If this interface has redundancy (cfg-if-vlan) enabled, the shutdown command causes the backup to reboot. Additionally, any ron tunnels that use this address will fail if you shut it down; a shadow-copy-rule depends on RON tunnels to communicate with other ARXes in the network.*
>
> *Use the shutdown command only on the advice of F5 Support.*

Use no shutdown to restart the management interface.

Mode cfg-if-vlan

Security Role(s) network-engineer or crypto-officer

Syntax

```
shutdown
no shutdown
```

Default(s) shutdown

Guidelines If you shut down the in-band management interface, administrators are immediately disconnected from the interface.

An ARX in a redundant pair has additional issues. A redundant ARX requires an in-band (VLAN) management interface on the same VLAN as its quorum disk. This is required so that the ARX can reach its quorum disk while in the backup role. Refer to the quorum-disk documentation for more details. If this interface has redundancy (cfg-if-vlan) enabled, the interface is even more vital to redundancy; in this case, a shutdown causes the backup ARX to reboot. Do not shut down this interface if this ARX has a quorum disk on the current VLAN or if the interface is used as the redundancy link.

If the current in-band (VLAN) interface terminates any Resilient Overlay Network (RON) tunnels, this command shuts down communication with all of them. This could potentially disrupt a shadow-copy rule or management access to another ARX.

Samples

```
bstnA(cfg-if-vlan[4])# shutdown
```

shuts down the in-band management interface for VLAN 4.

```
bstnA(cfg-if-vlan[4])# no shutdown
```

restarts the same interface.

Related Commands interface vlan
quorum-disk
ron tunnel

shutdown (cfg-mgmt)

Purpose From cfg-mgmt mode, use the shutdown command to shut down the out-of-band management interface.

Use no shutdown to restart the management interface.

Mode cfg-mgmt

Security Role(s) network-engineer or crypto-officer

Syntax
```
shutdown
no shutdown
```

Default(s) **no shutdown**: the interface is enabled by default after you go through the initial-boot process.

Platforms any *except* ARX-VE

Guidelines If you shut down the out-of-band management interface, administrators are immediately disconnected from the interface.

The ARX-1500 and ARX-2500 use port 1/1 for out-of-band management by default. If you use this command on one of those platforms, you can then dedicate port 1/1 to client/server traffic with the interface gigabit 1/1 command and its sub commands.

Samples
```
bstnA(cfg-mgmt)# shutdown
```
shuts down the out-of-band management interface.

```
bstnA(cfg-mgmt)# no shutdown
```
restarts the interface.

Related Commands interface mgmt

speed (cfg-mgmt)

Purpose From cfg-mgmt mode, use the speed command to set the speed and duplex configuration on the fast Ethernet port for the switch's out-of-band (OOB) management interface.

Mode cfg-mgmt

Security Role(s) network-engineer or crypto-officer

Syntax **`speed {auto | 100-half | 100-full | 10-half | 10-full | 1000-full}`**

auto | 100-half | 100-full | 10-half | 10-full | 1000-full is a required choice:

auto is auto-negotiate (the default) 10/100 megabits-per-second (mbps), half/full duplex. Makes the port auto-negotiate with its peer. Use this setting to enable MDI/MDIX cross-over on the OOB-management port.

100-half is fast Ethernet, 100 mbps, half duplex.

100-full is fast Ethernet, 100 mbps, full duplex.

10-half is fast Ethernet, 10 mbps, half duplex.

10-full is fast Ethernet, 10 mbps, full duplex.

1000-full is Gigabit Ethernet, 1000 mbps, full duplex.

Default(s) auto

Platforms any *except* ARX-VE

Guidelines The ARX has one out-of-band management interface labeled MGMT on the front panel of the switch.

The MGMT port supports automatic MDI/MDIX cross-over. This feature automatically corrects the polarity of the attached CAT5 cable, regardless if it is a cross-over or straight-through type. However, the port must be set to "auto" (auto-negotiate enabled). When the port speed/duplex is forced, automatic MDI/MDIX cross-over is disabled, and you must cable the port using standard cross-over or straight-through cabling. That is, connections between switches and/or routers require a "crossover" cable, and connections between switches/routers to hosts/end-stations require a "straight-through" cable.

Sample

```
bstnA(cfg)# interface mgmt
bstnA(cfg-mgmt)# speed 100-half
```

For the OOB-management interface, sets the fast Ethernet port speed to 100 mbps and the duplex configuration to half duplex.

Related Commands interface mgmt
show interface mgmt

wait-for ip-routes

Purpose Use the wait-for ip-routes command to wait until all of your configured static routes are operational.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`wait-for ip-routes [timeout`** ***`timeout`*****`]`**

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - none, wait indefinitely

Guidelines This command is useful in CLI scripts.

After you have established one or more static routes with the ip route command, you can use the wait-for ip-routes command to wait for all of those routes to come online. This can be useful for CLI scripts, which you can copy onto the switch (with copy ftp, copy scp, copy {nfs|cifs}, or copy tftp), and run.

If you set a ***timeout*** and it expires before the last static route is up, the command exits with a warning. To interrupt the wait-for ip-routes command, press **<Ctrl-C>**. You can use the show ip route monitor command to see the current status of all static routes.

Samples

```
stkbrgA# wait-for ip-routes
```

waits indefinitely for all static routes to come up. The CLI prompt does not return until this command proves that all routes are functional.

```
stkbrgA# wait-for ip-routes timeout 30
```

waits up to 30 seconds for all static routes to come up.

Related Commands ip route
show ip route

11

Licensing

This chapter contains an alphabetical list of commands for activating and using software licenses in the ARX.

clear active-license

Purpose Use the clear active-license command to remove the active license from this system. This renders the current system unlicensed, so that you can no longer use its storage services or enter gbl mode in the CLI. Use the command only on the advice of F5 Support.

Mode priv-exec

Security Role(s) crypto-officer

Syntax **`clear active-license`**

Default None

Guidelines Use this command only on the advice of F5 Support. The CLI prompts for confirmation before it clears the license; enter **yes** to proceed.

If the ARX can connect to the Internet, you can use the license activate command to automatically re-activate the license. Otherwise, you can use a manual activation method, as described in the documentation for the license create license-dossier command.

You can use the show active-license command to view the license-expiration time and all the currently-licensed features on the system.

Sample

```
stkbrgA# clear active-license
Warning: Clearing the license renders the ARX inoperable. Continue with
operation? [yes/no] yes
```

clears the license on an ARX named "stkbrgA."

Related Commands

license activate
license create license-dossier
show active-license

license activate

Purpose

Use the license activate command to activate your software license and begin setting up storage services on the ARX.

Mode

priv-exec

Security Role(s)

crypto-officer

Syntax

```
license activate base-reg-key base-key
license activate add-on-key add-key
```

base-reg-key *base-key* (1-33 characters) is the base registration key for the current ARX. If you purchased hardware with Release 6.00.000 or later, the key is installed before shipment; use the <tab> key to auto-complete your base registration key. To upgrade older hardware to Release 6.00.000 or later, contact your F5 Sales representative to purchase a license and receive the key in an E-mail message. For trial software, you can get a key for a trial license E-mailed to you from the same site where you downloaded the software, https://www.f5.com/trial/.

add-on-key *add-key* (1-15 characters) is a registration key for an add-on license. You can purchase an add-on license to add more capacity or additional features to your ARX. As above, contact your F5 Sales representative to purchase an add-on license and receive the key in an E-mail message.

Guidelines: Trial or Evaluation Licenses in a Redundant Pair

Trial licenses, obtained from https://www.f5.com/trial/, and evaluation licenses have an expiration date built into them. The timer for the license expiration begins when you run the license activate command or manual activation (described below). If you are installing one of these temporary licenses on a redundant pair of ARX devices, the start dates should be no more than 24 hours apart. We recommend activating the license at both peers on the same day. This avoids a large gap in time where one peer is licensed and the other peer is unlicensed, and it also avoids an SNMP trap that warns you of the license mismatch.

Guidelines: Automatic Activation

This command contacts a license server at F5, http://activate.f5.com/, so the ARX must have a static route (ip route) to the Internet and access to a DNS server (ip name-server) before you use it. These network settings should be properly configured during installation, in the initial-boot script. You can use ping license-server to confirm that you have a usable connection to the server. If not, you can use the show ip route command for a list of static routes, and use show ip route monitor to check the status of each route. You can also use show ip domain to see your DNS setup, and/or ip name-server to add a new DNS server to the ARX configuration.

If circumstances or network policy prevents you from connecting the ARX to the Internet, you can use a manual activation method, as described in the documentation for the license create license-dossier command.

In a redundant pair of switches, each ARX has its own license. Activate the license on both peers. If there are any add-on licenses, activate those on both peers, too.

Guidelines: Automatic Activation (Cont.)

Your initial license activation requires contact information and acceptance of an End-User License Agreement (EULA). The CLI only requests this information if you have not already entered it through other means.

The contact information comes in two parts: general information and an E-mail address.

The CLI requests general information first, with prompts for the following items:

- First Name
- Last Name
- Company Name
- Phone
- Job Title
- Address
- City
- State/Province
- Postal Code
- Country

Then the CLI shows you a summary of these items; enter **yes** to accept the entries, or **no** to restart the list.

The CLI then prompts you for an E-mail address. If you see this, enter the E-mail address and then re-enter it as prompted.

The End-User License Agreement (EULA) appears next and prompts you for confirmation. The EULA also appears for a re-run of this command if the agreement text has changed in any way. If the EULA appears, review the agreement carefully and then enter **yes** to accept it.

An information message appears at the end indicating the final success or failure of your license activation.

Use the show active-license command to see all of the licensed limits after you install the base license.

An add-on license changes the limits expressed in show active-license; the CLI informs you of all the changes after you run this command with the **add-on-key** option.

Sample

```
stkbrgA# license activate base-reg-key CRJGV-QPDYW-SATNK-RGBYY-DMTMOBL
% INFO: The license has been successfully activated.
```

automatically activates the license for the “stkbrgA” switch.

Related Commands

ip route
show ip route
ip name-server
ip domain-list
show ip domain
ping license-server
clear active-license

license activate file

Purpose

Some installations do not permit a network connection from the ARX to the Internet, but an Internet connection is required to access a valid F5 license server. For these installations, a manual process is necessary to activate the license on the ARX. The manual-activation process involves downloading a valid license file from the F5 license server to a remote host, and then transferring the license file to the ARX. The final step in manual-license activation is activating that valid license file on the ARX.

Use the license activate file command to activate a valid license file, one that was downloaded from http://activate.f5.com/. You can also use it to activate a copy of the "active.license" file (which is created by the license activate command).

Mode

priv-exec

Security Role(s)

crypto-officer

Syntax

`license activate file` ***`file-name`***

file-name (1-255 characters) is the name of the license file to activate. For a full list of available license files, you can use the show license command.

Guidelines

You can use a faster method for activating your license if the ARX has a valid route to the license server at http://activate.f5.com/. The ping license-server command tests the connection between the ARX and the license server. If this ping command succeeds, you can use the license activate command instead of this one.

Guidelines: Manual Activation

The manual activation process starts at the ARX CLI, continues on a remote host with access to the F5 server, and then concludes at the ARX CLI:

- The license create license-dossier command creates a dossier file to identify the ARX and enumerate its capabilities. The file name is "arx.dossier;" you can use the show license-dossier command to confirm that it exists.
- Use your preferred form of the copy command (such as copy ftp, copy scp, or copy smtp) to upload the arx.dossier file to a host that can connect to the Internet.

  For sites where such an upload is not permitted, you can use an alternative that is described below.
- From an HTTP browser on the remote host, connect to the following site and fill out the forms for getting an ARX license:

  http://activate.f5.com/

  When prompted to enter the ARX dossier, you can either upload it from the location above, or you can copy and paste the encrypted output from show license-dossier.

  This creates a file with a ".license" extension (such as "arx.license") on the remote host.
- At the ARX CLI, use copy *source-url* license arx.license to download the ".license" file from the remote host to the "license" directory on the ARX.
- Use this command, license activate file arx.license, to activate this new license on the ARX. One or more information messages appear at the end indicating the final success or failure of your license activation.

Guidelines: Trial or Evaluation Licenses in a Redundant Pair

Trial licenses, obtained from https://www.f5.com/trial/, and evaluation licenses have an expiration date built into them. The timer for the license expiration begins when you run the license activate command or this command. If you are installing one of these temporary licenses on a redundant pair of ARX devices, the start dates should be no more than 24 hours apart. We recommend activating the license at both peers on the same day. This avoids a large gap in time where one peer is licensed and the other peer is unlicensed, and it also avoids an SNMP trap that warns you of the license mismatch.

Sample

```
stkbrgA# license activate file arx.license
% INFO: The license has been successfully activated.
```

activates the license in the "arx.license" file.

Related Commands

show license
license activate
license create license-dossier
copy ftp
copy {nfs|cifs}
copy scp
copy smtp
show license-dossier
clear active-license

license create license-dossier

Purpose

Some installations do not permit a network connection from the ARX to the Internet, but an Internet connection is required to access a valid F5 license server. For these installations, a manual process is necessary to activate a license on the ARX. The first step in manual-license activation is creating a *dossier* file that describes the current ARX. Use the license create license-dossier command to create the dossier file.

Mode

priv-exec

Security Role(s)

crypto-officer

Syntax

```
license create license-dossier base-reg-key base-key
license create license-dossier add-on-key add-key
```

base-reg-key *base-key* (1-33 characters) is the base registration key for the current ARX. If you purchased hardware with Release 6.00.000 or later, the key is installed before shipment; use the <tab> key to auto-complete your base registration key. To upgrade older hardware to Release 6.00.000 or later, contact your F5 Sales representative to purchase a license and receive the key in an E-mail message. For trial software, you can get a key for a trial license E-mailed to you from the same site where you downloaded the software, https://www.f5.com/trial/.

add-on-key *add-key* (1-15 characters) is a registration key for an add-on license. You can purchase an add-on license to add more capacity or additional features to your ARX. As above, contact your F5 Sales representative to purchase an add-on license and receive the key in an E-mail message.

Guidelines

You can use a faster method for activating your license if the ARX has a valid route to the license server at http://activate.f5.com/. The ping license-server command tests the connection between the ARX and the license server. If this ping command succeeds, you can use the license activate command instead of this one.

Guidelines: Manual Activation

This command creates a dossier file to identify the ARX and enumerate its capabilities. The file name is "arx.dossier;" you can use the show license-dossier command to confirm that it exists.

Then use your preferred form of the copy command (such as copy ftp, copy scp, or copy smtp) to upload the arx.dossier file to a host that can connect to the Internet.

For sites where such an upload is not permitted, you can use an alternative that is described below.

From an HTTP browser on the remote host, connect to the following site and fill out the forms for getting an ARX license:

http://activate.f5.com/

When prompted to enter the ARX dossier, you can either upload it from the location above, or you can copy and paste the encrypted output from show license-dossier.

This creates a file with a ".license" extension (such as "arx.license") on the remote host.

At the ARX CLI, use copy *source-url* license arx.license to download the ".license" file from the remote host to the "license" directory on the ARX.

Use license activate file arx.license to activate this new license on the ARX.

Sample

```
stkbrgA# license create license-dossier base-reg-key CRJGV-QPDYW-SATNK-RGBYY-DMTMOBL
```

creates a dossier file, “arx.dossier,” for the “stkbrgA” switch.

Related Commands

- license activate
- show license
- copy ftp
- copy {nfs|cifs}
- copy scp
- copy smtp
- show license-dossier
- license activate file
- clear active-license

ping license-server

Purpose Use the ping license-server command to send a simple request to your license server. This verifies that you can connect to the license server and use the automatic license activate command.

Mode (any)

Security Role(s) crypto-officer

Syntax **`ping license-server [base-reg-key`** ***`base-key`*****`]`**

> **base-reg-key** ***base-key*** (1-33 characters) is the base registration key for the current ARX. If you purchased hardware with Release 6.00.000 or later, the key is installed before shipment; use the <tab> key to auto-complete your base registration key. To upgrade older hardware to Release 6.00.000 or later, contact your F5 Sales representative to purchase a license and receive the key in an E-mail message. For trial software, you can get a key for a trial license E-mailed to you from the same site where you downloaded the software, https://www.f5.com/trial/. This option is unnecessary after activation.

Guidelines The output from a successful ping shows the current time at the F5 license server, http://activate.f5.com/.

If the output shows an error instead of a timestamp, the license server is unreachable. The ARX must have a static route (ip route) to the Internet for this command to succeed. Use the show ip route command for a list of static routes, and use show ip route monitor to check the status of each route.

The ARX must also have a DNS server available for translating server FQDNs into IP addresses; use show ip domain to see your DNS setup, and/or ip name-server to add a new DNS server to the ARX configuration.

If circumstances or network policy prevents you from connecting the ARX to the Internet, you can use a manual activation method, as described in the documentation for the license create license-dossier command.

Sample

```
stkbrgA# ping license-server

% INFO: Activation server response: 'Tue Nov 09 14:52:01 UTC 2010'
```

> indicates that you can reach the license server, and that you can therefore use license activate to automatically activate your software license.

Related Commands

license activate
ip route
show ip route
ip name-server
ip domain-list
show ip domain

show active-license

Purpose Use the show active license command to display a complete list of all features and limits that are currently licensed on this ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show active-license`

Guidelines This command shows a list of all currently-licensed options. This information is broken into several tables.

System License Information is information about the active license itself. Here are some key fields in the output:

- Auth Vers. is for internal use only.
- Usage indicates the intended usage for the license. This is also for internal use only.
- Registration Key is the base registration key used for license activate and similar commands.
- Licensed version identifies the software release that is licensed. This is a release number with an *x.y.*0 format. You are permitted to install any maintenance release off of the same base, such as *x.y.*1, *x.y.*2, and so on.
- License Date is the date when the license was first activated with license activate, license activate file, or a GUI equivalent.
- License Start is the earliest date when the license could be valid. This does not appear for a production license, which has no start or end date.
- License End is the expiration date for the license. This field also does not appear for a production license.
- Service Check Date is used for comparison to any software that is installed later. The date of a major release is included in its release file; if that date is earlier than the license's *service-check date* (in this field), you are permitted to install the new release. Use the show releases command to get a list of all release files on the ARX, and use boot system to prepare the ARX to install the new release.
- Platform ID identifies the ARX platform to which this license applies. The codes are Z100 for an ARX-VE, A106 for an ARX-500, D108 for an ARX-2000, D103 for an ARX-4000, C110 for an ARX-1500, or C111 for an ARX-2500.
- Service Status is a string to explain the implications of the Service Check Date above.

Module List contains a sub table for each add-on license. Add-on licenses will be supported in a future release.

Feature List shows all of the features and limits that the active license supports.

- Some of these features contain a suffix of "_4k." This indicates that the values are shown in blocks of 4,096. Multiply each such value by 4096 to get the actual limit.

Guidelines (Cont.) Use the license activate command to automatically activate the license for the current ARX. Automatic activation is only possible if the ARX has a network connection to the F5 web site, http://activate.f5.com/; if not, you can use the license create license-dossier command to start the process of manual activation.

Sample `bstnA# show active-license`

shows the active license and all the options that it enables. See Figure 11.1 on page 11-12 for sample output.

Related Commands license activate
license create license-dossier

Figure 11.1 *Sample Output: show active-license*

```
bstnA# show active-license
System License Information
--------------------------

Auth Vers:                      5b
Usage:                          F5 Internal Product Development
Registration Key:               W218274-067731-88692-6045-2946811
Licensed version:               6.0.0
License Date:                   Nov 23 2010
License Start:                  Nov 22 2010
License End:                    May  3 2011
Service Check Date:             Apr  3 2011
Platform ID:                    D103
Service Status:                 As of 2011-04-03 there is no active service contract.
 :                               This may inhibit your ability to upgrade your software.

License Load Date:              Apr  4 2011

Module List
-----------------
ARX 4000:
Reg Key:                        W428510-1430583

LIC-PKG-ARX4-ENT:


Feature List
-----------------
cifs_services_per_system:                64
direct_attach_points_per_system:         524288
direct_shares_per_system:                16384
direct_shares_per_volume:                255
direct_shares_per_volume_group:          2048
files_per_system_4k:                     500000
files_per_volume_4k:                     62500
files_per_volume_group_4k:               62500
global_servers_per_system:               64
namespaces_per_system:                   16
nfs_services_per_system:                 64
protocol_qty_allowed:                    3
redundancy:                              enabled
shares_per_system:                       1024
shares_per_volume:                       64
shares_per_volume_group:                 128
virtual_services_per_system:             64
```

```
volume_groups_per_system:                16
volumes_per_system:                      256
volumes_per_volume_group:                32
```

show license-dossier

Purpose A license *dossier* is a file that describes the ARX and its capabilities. The dossier file is used for license activation; the license server requires the ARX dossier to provide you with a license. The dossier is encrypted. When you activate your license manually, you can copy and paste the encrypted dossier into the form on http://activate.f5.com/. Use the show license-dossier command to display the dossier file as an encrypted string.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show license-dossier`**

Guidelines This command is only useful for a manual license activation. You can use a faster method for activating your license if the ARX has a valid route to the Internet. If the ARX can connect to the Internet, you can use the license activate command to automatically activate your license.

This command shows the encrypted string that represents your dossier file, if the file exists. You can copy this string from the CLI and paste it into the proper field at http://activate.f5.com/. As an alternative, you can use some form of the copy command (such as copy ftp, copy scp, or copy smtp) to copy the "arx.dossier" file to another host, and upload the dossier from that host to http://activate.f5.com/.

Sample `stkbrgA# show license-dossier`

shows the "arx.dossier" file in encrypted form. See Figure 11.2 on page 11-14 for sample output.

Related Commands license activate

Figure 11.2 *Sample Output: show license-dossier*

```
stkbrgA# show license-dossier
bb614ac893297bd563645a60e136d44952ad1990c831f65be13d88b5268ffec885a5ff45a13f8e25d1de0b60881e1c0e78
a11ec840e7010f02a4d313ec243ba2961935b11d186a924d74188a1f635024f5010716d91f156e1278f4de6a8bcced12d4
cbab70494df4b812e77737d25debb2c3f482c6f1b867cf44db73d316fe67ee09eef5fea1dc675e5caf665344ecfa9b36cd
4bb0b1f9303b848a8c2d602971c1e864c181ec0c0d04579ad47aaf6da42e1d0b80afffe28622e5a2433555c8ca9dde13ab
7fe7e6f9c81d4ceb583b7cc25e6f0798db29c41ef8a65eefa04854fbc12bff8f71f682f3fc94e09cb2bc24c70e47bf3391
238225b5be7478fa59118c53a0ddd347b14e8ad0f69b7826a540969f8d2f2d79290c239d759151e2691dc72d78e33d7fda
91a3271e9ca64a450f6359efd56267f658e1f2e36027b543d897bb61cebd4f129bec14040c58851eb837674469f0e10101
8c9ae035aa55d1682e9a1f30038da69142540cd7766c801d06ae24e8fe1e6456b7e492d724c68a9a8bb5909e8f3e2b22b8
f9daae0f5607fc12629bdc32822ff85cf2d62133349338e24c66710cc83a1df88a4f99f754366ccb639938552fbe43fd82
c3fc4f8515031c273297345017513837aaf4798950fef8a19ab001dd9543facd973bb02033ed0bf74ef919ac51db3b06af
9dd0ad4a8c0c7cae49b2b3070f09fce06ce77fb687c359814cf304b57b2117262b90704ab206b7f768bbe601768851d722
db6df07640fb4c4f8dda28f0c382d58a9e8f7b93d9f5cb18e16b12e31407e8e7b4b32d889aa6baea9f8447e9371c0ea586
0511cc410427c9ebb46ade4c844767d7389f3f385bb9dd4495a5073618c966d90501fc791e293a7febf744969a5893bbc7
bac7682fc00090d5ea954f4781d019ddf87d0e620d5583221b7e1564ac77e6c8c1a087abef548066c7a3ab7d96a3d02230
5f2f09c7fed78d0752384f60506a62f24fa9152523781aab5c5455db15334a4414825b89cf5e38d367471b85e5712acd40
9f56b587ddd4b223babee95de098ff6c1cc973e18521e405c664a0b4d3728fa5cbc050993bc25d5bcf30fd170ef1af8ab3
f5941b1c97e7d4774b5e0b7ff18bf81bc20ce01090db42951734a576e721b9381ecf472f3fac3e7aa89c717e67f0198d4d
7562d95d04f6d9d71c3804ddc0cbacaf3b2991cc5fcb0c3b3a86f61a771f8d57ceb8cec114be658a34b908b76549811aff
4714e955966ec7cef0808a21f82c928393fc805ac6537e021f7de59c02978ff73aeaa2ab6f4bf313841ef84c1c1634b340
6f27ce2d1d82133e0ccc2538fb7788b0c6410f32a944cc855f2345c4c683906d97e896c78a4efdca8ccf4ce7cae5dc94ba
fdf5ceff580732d6234da308abcd52c15b9a0d1766ae8dd6bd59dde0d5b5bac86b14a9ae8a2eb2a4019871a3695052b441
```

```
a36cf777c01354b1e0176de064ad250dffd256cc48490a32ce99bea3876fce8eb01bd2b7bbbe4635fb6c23fc0611b835b0
1effe9758f2e163fe324628ec386e4dd8d257b12a05f5e2d4da3e298248df0b2bfa85c8ab12f8843d1ca4d7ff88354d1e6
4daf9f2deae7842e1812adebdc4a71fbe8addc5e1d051ba51793442740756aa51bc2e34e4ab8add498ff1eccb74a3184db
faf9171ad6837ae2d2bc3e55b96e628179dc5146cc53a502279fb5d895b8634bdfdfe29a7332f0acf22688317d5c4a88cb
806a6a4c18d0f896a8fe854aff4f3bc33b18f20ff0393c40001ec53a435aa570d61faa393332326a4337b64856829a951a
3d2a53e3613b
```

12

SNMP

show snmp-server

Purpose Use the show snmp-server command to see all SNMP-configuration settings.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show snmp-server
```

Guidelines System Information contains configuration information:

- Name can be reset with the snmp-server name command.
- Contact can be set with snmp-server contact.
- Location can be changed with snmp-server location. Name, Contact, and Location are the configurable parts of the MIB-2 system table.
- Interface is the interface used for SNMP communication: this is either "Management," "VLAN," or "None," as set by the cfg-mgmt-access permit command.
- Traps shows the scope of SNMP traps that the ARX can send. The options are "Private Traps enabled" (the traps that are unique to the ARX, cataloged in the ARX SNMP Reference), "All Traps enabled" (ARX-only traps together with generic SNMP traps), or "Disabled." You can set this with the snmp-server traps command.

Access Mode/Communities is a table of access permissions (read-only, or read-write) and the community string associated with each. The community string is encrypted. You can use snmp-server community to set it.

Trap Targets shows all of the SNMP hosts configured to receive traps. Use snmp-server host to change this list.

Trusted Hosts are allowed to perform gets and/or sets. The snmp-server trusthost command controls this setting.

Sample

```
bstnA> show snmp-server
```

shows the SNMP configuration for the "bstnA" switch. See Figure 12.1 for sample output.

Related Commands

snmp-server contact
snmp-server name
snmp-server location
snmp-server community
snmp-server trusthost
snmp-server traps
snmp-server host

Figure 12.1 *Sample Output: show snmp-server*

```
bstnA> show snmp-server

System Information
----------------------------------------------------------
Name:           arx1
Contact:        jpublic, jpublic@mycompawells.me.org
Location:       2nd floor lab, row 3, bay 4, shelf 5
Interface:      Management/VLAN
Traps:          All traps enabled


Access Mode    Communities
----------------------------------------------------------
read-write     private
read-only      public


Trap Targets
Address          Community                      UDP Port
----------------------------------------------------------
10.1.1.68        public                          162
172.16.100.101   public                          162
172.16.100.183   public                          162


Trusted Hosts
Address
----------------------------------------------------------
10.1.1.68
172.16.100.183
```

snmp-server community

Purpose

External SNMP agents use a community string to access the local SNMP agent. Use the snmp-server community command to enter a valid community string for SNMP access.

Use the no form of snmp-server community to remove a community string from the list.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`snmp-server community` ***`community-string`*** **`{read-only | read-write}`**
`no snmp-server community` ***`community-string`*** **`{read-only | read-write}`**

community-string is up to 16 characters.

read-only | read-write is a required choice. This sets the permissions for agents that access the local SNMP agent.

Default(s)

None

Guidelines

You can use this command multiple times, once for each supported community string. The no form removes one of the strings from the list.

After you have configured a community string with this command, you can use the snmp-server trusthost command for each trusted SNMP host.

Samples

```
bstnA(cfg)# snmp-server community public read-only
```

establishes read-only access for trusted hosts that use the "public" community string in their SNMP queries.

```
bstnA(cfg)# no snmp-server community experimental read-write
```

removes a community string, "experimental," from the list of valid strings.

Related Commands

snmp-server trusthost

snmp-server contact

Purpose

Use the snmp-server contact command to enter contact information (sysContact, MIB-2.1.1.4) for the ARX.

Use no snmp-server contact to clear the contact information.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`snmp-server contact` ***`contact-info`***
`no snmp-server contact`

contact-info (1-255 characters) is the SNMP contact information. Use quotation marks around the string if it contains any spaces.

Default(s)

No contact information is set.

Guidelines

The sysContact MIB object is part of the MIB-2.system tree, MIB-2.1.1.4. Use the hostname command to set the sysName object, and use the snmp-server location command to set the sysLocation.

Samples

```
bstnA(cfg)# snmp-server contact "jpublic, jpublic@mycompany.com"
```

sets the sysContact MIB object for the ARX.

```
bstnA(cfg)# no snmp-server contact
```

erases the setting for the sysContact MIB object.

Related Commands

hostname
snmp-server location

snmp-server host

Purpose

Use the snmp-server host command to configure a target for SNMP notifications (such as SNMP traps).

Use the no form of the command to remove an SNMP server from the list.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`snmp-server host` ***`host-ip community-string`*** **`[`*****`udp-port`*****`]`**
`no snmp-server host` ***`host-ip`***

host-ip is the IP address of the SNMP host (for example, 192.168.70.65).

community-string (up to 16 characters) is sent in each SNMP packet.

udp-port (optional; 1-65535) is the UDP port where the remote host is listening for SNMP notifications. For most SNMP hosts, the default (162) is sufficient.

Default(s)

162, the well-known port for SNMP notifications, is the default *udp-port*.

Guidelines

You can use this command multiple times to define up to 6 trap destinations.

Use the snmp-server traps command to enable SNMP traps.

For communication with an SNMP server (host access or trap transmission), you must also permit access to the SNMP agent through the internal firewall. Use the management access snmp command to edit the rules for SNMP, then use the permit command to permit server access through the out-of-band management interface (labeled MGMT on the front panel), the in-band (VLAN) management interface(s), or both.

Samples

```
bstnA(cfg)# snmp-server host 10.1.1.68 public
bstnA(cfg)# snmp-server host 172.16.100.101 public
```

sets two destinations for SNMP notifications.

```
bstnA(cfg)# no snmp-server host 192.168.25.215
```

removes the SNMP host at 192.168.25.215 from the list of trap destinations.

Related Commands

snmp-server traps
management access -> permit

snmp-server location

Purpose Use the snmp-server location command to enter location information (sysLocation, MIB-2.1.1.6) for the ARX.

Use the no form of the command to clear the sysLocation setting.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
snmp-server location location-info
no snmp-server location
```

location-info (1-255 characters) is the SNMP-server location information. Use quotation marks around the string if it contains any spaces.

Default(s) No location information is set.

Guidelines The sysLocation MIB object is part of the MIB-2.system tree, MIB-2.1.1.6. Use the hostname command to set the sysName object, and use the snmp-server contact command to set the sysContact.

Samples

```
bstnA(cfg)# snmp-server location "2nd floor lab, row 3, bay 4, shelf 5"
```

sets the sysLocation MIB object for the ARX.

```
bstnA(cfg)# no snmp-server location
```

clears the sysLocation object for the ARX.

Related Commands

hostname
snmp-server contact

snmp-server name

Purpose Use the snmp-server name command to enter a new SNMP-advertised name for the ARX (sysName, MIB-2.1.1.5).

Use the no form of the command to erase the name.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
snmp-server name machine-name
no snmp-server name [machine-name]
```

machine-name (1 to 255 characters) is the administrative name of the server.

Default(s) None

Guidelines The sysName MIB object is part of the MIB-2.system tree, MIB-2.1.1.5. Use the snmp-server location command to set the sysLocation object, and use the snmp-server contact command to set the sysContact.

Sample

```
bstnA(cfg)# snmp-server name acopia1
```

sets the sysName MIB object for the "bstnA" ARX.

Related Commands snmp-server location
snmp-server contact

snmp-server traps

Purpose

Use the snmp-server traps command to enable SNMP traps from the ARX.

Use the no snmp-server traps command to disable SNMP traps.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
snmp-server traps [private]
no snmp-server traps
```

private (optional) enables only private traps, and disables the standard generic traps defined in RFC 1157. If you omit this option, the command enables all traps. Private traps are the enterprise traps defined specifically for the ARX. For a full list of F5's ARX-enterprise (private) traps, see the ARX SNMP Reference.

Default(s)

all traps enabled

Guidelines

This command enables all SNMP traps, with an option to omit the generic traps and only use private (ARX) traps. The *generic* traps are defined in RFC 1157, the specification for SNMP, as one of the following:

- coldStart
- warmStart
- linkDown
- linkUp
- authenticationFailure
- entityConfigChange (defined in the RFC as an enterpriseSpecific trap)

Use the snmp-server host command to provide a destination for SNMP traps.

For communication with an SNMP server (host access or trap transmission), you must also permit access to the SNMP agent through the internal firewall. Use the management access snmp command to edit the rules for SNMP, then use the permit command to permit server access through the out-of-band management interface (labeled MGMT on the front panel), the in-band (VLAN) management interface(s), or both.

If you are sending traps out through a channel on the client/server VLANs, use no trap shutdown to enable SNMP traps on the channel.

Samples

```
bstnA(cfg)# snmp-server traps
```

enables all trap types.

```
bstnA(cfg)# no snmp-server traps
```

disables all traps.

Related Commands

snmp-server host
management access -> permit
no trap shutdown

snmp-server trusthost

Purpose

Use the snmp-server trusthost command to permit SNMP gets and sets from a remote host.

Use the no form of the command to remove an SNMP host from the list of trusted hosts.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`snmp-server trusthost` ***`host-ip`***
`no snmp-server trusthost` ***`host-ip`***

host-ip is the IP address of the SNMP host (for example, 192.168.70.65).

Default(s)

All hosts are trusted (if they use the string(s) specified through snmp-server community) until you limit the number of trusted hosts with this command.

Guidelines

You can use this command multiple times to configure up to 8 trusted hosts. Before you use this command, all hosts are trusted (though they are denied access through the firewall; see below).

For communication with an SNMP server (host access or trap transmission), you must also permit access to the SNMP agent through the internal firewall. Use the management access snmp command to edit the rules for SNMP, then use the permit command to permit server access through the out-of-band management interface (labeled MGMT on the front panel), the in-band (VLAN) management interface(s), or both.

Samples

```
bstnA(cfg)# snmp-server trusthost 10.1.1.68
```

allows the SNMP host at 10.1.1.68 to access the local SNMP agent.

```
bstnA(cfg)# no snmp-server trusthost 172.16.100.101
```

removes the SNMP host at 172.16.100.101. The local SNMP agent will no longer accept SNMP gets or sets from that IP address.

Related Commands

snmp-server community
management access -> permit

13

Email Notifications (SMTP)

You can set up the Simple Mail Transfer Protocol (SMTP) to send email notifications for ARX events.

clear smtp queue

Purpose The SMTP *mail queue* holds email messages until they are successfully delivered. Use the clear smtp queue command to delete all such queued messages.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear smtp queue`**

Default None

Guidelines Messages are only in the mail queue for a very short time before the first attempt to deliver them. If you can see a message with show smtp queue, the message has not yet been successfully delivered. Use this command to delete all messages from the queue.

After deleting the messages, this shows information about each message: the "To" and "From" fields, the size, and the age (time in the mail queue).

Alternatively, you can retry all queued messages with smtp retry.

Use show smtp status to see details on the most-recent message delivery, and to view the current configuration for SMTP.

Sample

```
bstnA# clear smtp queue

  To: juser@wwmed.com From:admin@wwmed.com Size:393.00  Age:5h has been
deleted.
  To: juser@wwmed.com From:admin@wwmed.com Size:393.00  Age:5h has been
deleted.
  To: juser@wwmed.com From:admin@wwmed.com Size:393.00  Age:19m has
been deleted.
  To: juser@wwmed.com From:admin@wwmed.com Size:393.00  Age:18m has
been deleted.

bstnA#
```

clears four pending email messages.

Related Commands

smtp
show smtp queue
smtp retry
show smtp status

description (cfg-email-event)

Purpose

Use the optional description command to set a descriptive string for the current email-event role. An *email event role* is a configurable list of SNMP traps, along with a group of users who receive email notifications for those traps. This appears in the show command.

Use the no form of the command to delete the description.

Mode

cfg-email-event

Security Role(s)

network-engineer or crypto-officer

Syntax

`description` *`text`*
`no description`

text (1-255 characters) is your description. Surround the text with quotation marks (“ ”) if it contains any spaces.

Default(s)

“Built-In” for the default tech-support role

None for new email event roles.

Guidelines

The description appears in the output for show email-event.

Sample

```
bstnA(cfg-email-event[noc3])# description “support team at NOC3”
```

specifies a description for the “noc3” email event role.

Related Commands

email-event
show email-event

email-event

Purpose Use the email-event command to start defining a set of ARX events to be emailed, the conditions under which the switch sends the email(s), and any recipients who should receive them. This creates a profile, or *role*, for one or more users who require event notification.

Use no email-event to remove an email event role.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
email-event role
no email-event role
```

role (1-128 characters) is a name that you choose for this set of email events, trigger conditions, and mail recipients (for example, "nocGroup").

Default "tech-support" is a default email event role, shipped from the factory

Guidelines Before an email-event can notify any users of system events, smtp must be configured for this ARX.

This command places you in cfg-email-event mode, where you can access several commands that define the email event role. To enable groups of events at a time, use the group (cfg-email-event) command. You can also use this command to set thresholds or schedules for email notification (for example, send for every 3 occurrences of a group event, or once each week). You can enable individual events with the group ... event command, which can also override any default-threshold settings made for the event's group. The mail-to (cfg-email-event) command chooses an email recipient for this set of events; enter this command once for each desired recipient. You can use the description (cfg-email-event) command to add an optional description to the email event configuration; this appears in the output of show email-event. Use enable (cfg-email-event) to enable all of the rules in this email event configuration. Notification does not begin until you enable the email event role.

You can reuse the email-event command to define multiple email event roles.

Email messages have a field in the SMTP header where you can set a severity level. SMTP servers can use this field for sorting and flagging incoming messages; the server at F5 Support is configured to use this field for prioritizing customer issues. Use the email-severity command to set the email severity for a given event. The show email-severity command shows event severities.

To send a small test message to all configured email recipients, use the smtp test email-event command. After a successful test, you can use smtp welcome to send an introductory email to all of the recipients.

This email delivery mechanism can be used together with standard SNMP-trap delivery. Use snmp-server traps to start setting up SNMP traps.

Samples

```
bstnA(cfg)# email-event noc3
bstnA(cfg-email-event[noc3])# ...
```

creates an email event role named “noc3.”

```
bstnA(cfg)# no email-event testEmail
```

removes an email event configuration.

Related Commands

mail-to (cfg-email-event)
group (cfg-email-event)
group ... event
description (cfg-email-event)
enable (cfg-email-event)
smtp
snmp-server traps
smtp test email-event
smtp welcome

email-severity

Purpose

You can use the email-event feature to send emails to one or more users (possibly including F5 Support) for certain SNMP traps. A field in the SMTP-message header contains the severity for the message. This field is not shown in most email applications, but email servers can use it for prioritization. Use the email-severity command to modify the email severity for a particular event.

Use no email-severity to revert to a default email severity.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`email-severity` *`event-name`* `level {normal | warn | minor | major | critical}`
`no email-severity` *`event-name`*

event-name (1-64 characters) identifies an event. An event corresponds to one SNMP trap, cataloged in the ARX SNMP Reference. Use a **?** after the email-severity keyword for a full list of eligible email events (for example, **email-severity ?**).

normal | warn | minor | major | critical sets the severity for the event's email messages.

Default

Defaults vary for each event; use show email-severity all to see them.

Guidelines

To send SNMP traps in email, configure smtp and then set up at least one email-event role.

Use this command to set a severity level for an event's email messages. The severity level appears in a customized field in the SMTP header, "X-ARX-Severity." SMTP servers can use this field for sorting and flagging messages; the SMTP server at F5 Support is configured to sort email messages by this setting. You can use this to prioritize your events for Customer Support.

Each severity that you set with this command is system-wide; it applies to any email-event role that uses the event.

Use show email-severity to show the current severities for one or more events.

Samples

```
bstnA(cfg)# email-severity auto-reboot level critical
```

sets the email severity to "critical" for an "auto-reboot" event.

```
bstnA(cfg)# no email-severity disk-state
```

resets the "disk-state" event to a "normal" severity.

Related Commands

show email-severity
smtp
email-event

enable (cfg-email-event)

Purpose

You can define one or more SNMP-trap events to be emailed to one or more users, as well as the thresholds for sending each event. One collection of events, thresholds, and users is called an email event *role*. Use the enable command to activate the current role.

Use no enable to disable the current email event role.

Mode

cfg-email-event

Security Role(s)

network-engineer or crypto-officer

Syntax

```
enable
no enable
```

Default(s)

Disabled

Guidelines

You must enable an email event role to activate it.

Samples

```
bstnA(cfg-email-event[noc3])# enable
```

enables the current email event role, noc3.

```
bstnA(cfg-email-event[betaTest])# no enable
```

disables the “betaTest” role.

Related Commands

email-event

from (cfg-smtp)

Purpose

Use the from command to set the "From" field in all email messages from the ARX.

Use no from to revert to the default "From" string.

Mode

cfg-smtp

Security Role(s)

network-engineer or crypto-officer

Syntax

```
from hostname@domain
no from
```

hostname@domain is 1-132 characters (for example, "admin@arx1000-B.nemed.com"). Use only alphabetical characters (a-z and A-Z), numeric characters (0-9), - (dash), _ (underscore), and/or . (period).

Defaults

None

Guidelines

The string you set with this command appears in the header of all email messages, in the "From" field.

Enter an email address that uses acceptable characters (listed above in *Syntax*). If there are any invalid characters, the CLI rejects the string with an error.

To see the current setting for the "From" field, use the show smtp status command.

Samples

```
bstnA(cfg-smtp)# from aco_admin@acopia.wwmed.com
```

sets the "From" field for all outgoing messages.

```
bstnA(cfg-smtp)# no from
```

erases the "From" string and reverts to the default.

Related Commands

smtp

group (cfg-email-event)

Purpose You can define SNMP-trap events to be emailed to one or more users, as well as the thresholds for sending each event. One collection of events, thresholds, and users is called an email event *role*. Use the group command to add an entire group of events to an email event role.

Use the no form to remove a group of events from a role.

Mode cfg-email-event

Security Role(s) network-engineer or crypto-officer

Syntax

```
group group-name [threshold-counter counter |
threshold-interval time {minutes | hours | days}]

no group group-name [threshold-counter counter |
threshold-interval time {minutes | hours | days}]
```

group-name is **chassis**, **cifs**, **storage**, **metadata**, **nsck**, **policy**, **snapshot**, **stats-monitor**, **virtual-server**, **network**, **redundancy** (for any chassis except the ARX-VE), or **all**.

threshold-counter ***counter*** (optional, 1-128) sets a threshold based on the number of events that occur from this group. For example, **threshold-counter 4** means that each event from this group triggers an email if it occurs 4 times.

threshold-interval ***time*** **{minutes | hours | days}** (optional) sets a schedule for this group. For example, **threshold-interval 1 days** causes the system to accumulate traps throughout the day and, if there are any, send a single email describing all of them. All time calculations begin when you enter this configuration; if the threshold interval is 1 day, the switch sends the email at the current time of day.

Default no groups enabled

If you enable a group without setting any threshold, the default is **threshold-counter 1**.

Platforms You cannot choose the **redundancy** group on the ARX-VE.

Guidelines

This command enables one group of events for the current email event role. Whenever the system reaches a threshold you set with this command, it sends an email to any or all mail-to (cfg-email-event) recipients. You can optionally use the group ... event command to configure a single event in the group, or to set a different threshold for a particular event.

The system keeps a maximum of 128 events for any event group. This is a hard threshold, even if you use the threshold-interval option to create a time-based schedule. For example, suppose you configure the following:

```
group cifs threshold-interval 7 days
```

This sends an email every week. However, if 128 CIFS events occur before the week is over, the system delivers an interim email with those events. This resets the start time for the 7-day schedule: the ARX delivers the next set of CIFS events 7 days later, or after 128 more events, whichever comes first.

If you use group ... event to add a specific event, the no group command for its group does not delete that event. For example, the following command sequence does not delete the "warmstart" event:

```
group chassis event warmstart
no group chassis
```

Samples

```
bstnA(cfg-email-event[noc3])# group chassis
```

adds the group of "chassis" events to this email role. This command uses the default threshold; every chassis event triggers an email as soon as it happens.

```
bstnA(cfg-email-event[noc3])# group metadata threshold-counter 5
```

adds the group of "metadata" events to this email role. Every 5th metadata event triggers an email.

```
bstnA(cfg-email-event[noc3])# no group storage
```

removes the "storage" event group. This prevents any storage-group events from triggering emails.

Related Commands

email-event
group ... event

group ... event

Purpose

You can define one or more SNMP-trap events to be emailed to one or more users, as well as the thresholds for sending each event. One collection of events, thresholds, and users is called an email event *role*. Use the group ... event command to add one event to an email event role.

Use the no form to remove an event from a role, or to revert its threshold to the group default.

Mode

cfg-email-event

Security Role(s)

network-engineer or crypto-officer

Syntax

```
group group-name event event-name [threshold-counter counter |
threshold-interval time {minutes | hours | days}]
no group group-name event event-name [threshold-counter counter
| threshold-interval time {minutes | hours | days}]
```

group-name is **chassis**, **cifs**, **storage**, **metadata**, **nsck**, **policy**, **snapshot**, **stats-monitor**, **virtual-server**, **network**, **redundancy** (for any chassis except the ARX-VE), or **all**.

event-name (1-64 characters) is one of the events in the group. Each of these corresponds to one SNMP trap, cataloged in the ARX SNMP Reference. Use a **?** after the event keyword for a full list of events in this group (for example, **group chassis event ?**).

threshold-counter ***counter*** (optional, 1-128) sets a threshold based on the number of times this event occurs. For example, **threshold-counter 4** means that any four occurrences trigger an email.

threshold-interval ***time*** **{minutes | hours | days}** (optional) sets a schedule for this event. For example, **threshold-interval 1 days** causes the system to accumulate instances of this event throughout the day and, if there are any, send a single email describing all of them. All time calculations begin when you enter this configuration; if the threshold interval is 1 day, the switch sends the email at the current time of day.

Default

All defaults are defined by the event's group (see group (cfg-email-event)).

Platforms

You cannot choose the **redundancy** group on the ARX-VE.

Guidelines

This command enables one event for the current email event role. Whenever the system reaches its threshold for this event, it sends an email to any or all mail-to (cfg-email-event) recipients for this role.

If you do not set any threshold for an event, its group threshold is used (as set by the group (cfg-email-event) command).

The no form of the command disables the event's threshold (so that the event defaults to the threshold for its group), or removes the event altogether. See the *Samples*, below.

Samples

```
bstnA(cfg-email-event[noc3])# group metadata event online
threshold-counter 2
```

adds one “metadata” event, “online,” to this email role. The threshold is 2 events; the system sends an email for every second event that indicates that a metadata share is online.

```
bstnA(cfg-email-event[testteam])# group virtual-server event
server-offline
```

adds the “server-offline” event to email event role, “testteam.” The threshold defaults to whatever threshold is set for the “virtual-server” group.

```
bstnA(cfg-email-event[noc3])# no group metadata event online
threshold-counter 2
```

reverts the “online” event to the default threshold. This default is defined by the “metadata” group.

```
bstnA(cfg-email-event[noc3])# no group chassis event cpu-status
```

removes the “cpu-status” event from this email event role.

Related Commands

email-event
mail-to (cfg-email-event)
group (cfg-email-event)

mail-server

Purpose The mail-server command identifies the next-hop email server for the ARX.

Use no mail-server to delete the email server name.

Mode cfg-smtp

Security Role(s) network-engineer or crypto-officer

Syntax

```
mail-server ip-or-name
no mail-server
```

ip-or-name (1-132 characters) identifies the mail server, either by IP address or DNS name (for example, “192.168.25.44” or “mailServer3.myco.com”).

Default None

Guidelines This is required for SMTP to function. This is the local mail server to deliver the ARX’s email messages to the WAN.

You must have a DNS server configured to use a DNS name in this command. Refer to the ip name-server documentation to configure a DNS server.

To see the currently configured mail server, use the show smtp status command. To test the mail-server configuration, use the smtp test server command.

Important

The no mail server *command disables all email deliveries.*

Sample

```
bstnA(cfg-smtp)# mail-server email1.wwmed.com
```

identifies the mail server.

Related Commands

smtp
smtp test server
ip name-server

mail-to (cfg-email-event)

Purpose You can define one or more SNMP-trap events to be emailed to one or more users, as well as the thresholds for sending each event. One collection of events, thresholds, and users is called an email event *role*. Use the mail-to command to set one destination address for the current role's emails.

Use the no mail-to command to remove an email recipient.

Mode cfg-email-event

Security Role(s) network-engineer or crypto-officer

Syntax

```
mail-to recipient
no mail-to recipient
```

recipient (1-768 characters) is one email recipient (for example, "juser@nemed.com").

Default None

Guidelines This sets a destination for all emails from this email event role; this address appears in the "To" field of the emails. You can enter this command multiple times, once for each recipient. At least five separate destinations are supported.

A sample email appears in Figure 13.1 on page 13-15. You can use the ID (highlighted in the sample) to look up the trap in the ARX SNMP Reference.

Samples

```
bstnA(cfg-email-event[noc3])# mail-to juser@wwmed.com
```

sets up one email recipient for the "noc3" role. If a system event triggers a threshold, the system sends an email notification to juser@wwmed.com.

```
bstnA(cfg-email-event[noc3])# no mail-to ex@nemed.com
```

removes an email recipient from the "noc3" role.

Related Commands email-event

Figure 13.1 *Sample Email Event*

```
From:   admin@wwmed.com
Sent:   Friday, March 05, 2010 3:13 AM
To: Dan Owen; jqpublic@wwmed.com
Subject:    ::trap id:690(share-remove-complete)::cn:acopia1::switch:bstnA::loc:2nd
floor lab, row 3, bay 4, shelf 5::desc:ARX-4000

Software Version: Version 5.02.000.12543 (Mar  2 2010 20:13:33) [nbuilds]
Chassis Serial Number: BZDS72000182

Group:       storage
Event:       share-remove-complete (ID: 690)
Total Events:1

Time(UTC):   2010-03-05T08:12:49.972
Events:
 Share: medarcv:/test_results:2005_charts Name: [medarcv]
```

maximum age

Purpose

When the ARX fails to deliver an email message, it retries periodically. If the failures persist for too long, the ARX deletes the message. Use the maximum age command to determine the maximum amount of time to retry before deleting the message.

Use no maximum age to return to the default.

Mode

cfg-smtp

Security Role(s)

network-engineer or crypto-officer

Syntax

```
maximum age days
no maximum age
```

days (1-366) is the number of days to retry before deleting the message.

Default

4 (days)

Guidelines

The ARX retries a message every *n* minutes for the first two hours; you can use retry interval to set the value of *n*. The retry interval increases geometrically if the failures persist. After some number of days without success (set by this command, maximum age), the switch deletes the message from the mail queue.

Use show smtp queue to view all messages currently in the mail queue. These are the messages that have failed and are being saved for retries. You can use smtp retry to retry all queued messages now. The clear smtp queue command deletes all the messages without sending them.

Samples

```
bstnA(cfg-smtp)# maximum age 30
```

keeps all failed emails for up to 30 days before deleting them.

```
bstnA(cfg-smtp)# no maximum age
```

reverts to the default.

Related Commands

smtp
retry interval
smtp retry
clear smtp queue

retry interval

Purpose

When the ARX fails to deliver an email message, it waits for some number of minutes before retrying. Use the retry interval command to set the initial interval between retries.

Use no retry interval to return to the default interval.

Mode

cfg-smtp

Security Role(s)

network-engineer or crypto-officer

Syntax

```
retry interval minutes
no retry interval
```

minutes (1-119) is the number of minutes to wait between email retries.

Default

15 (minutes)

Guidelines

The ARX retries a message every *n* minutes for the first two hours; this command sets the value of *n*. The retry interval increases geometrically if the failures persist. After some number days without success (set by maximum age), the switch deletes the message from the mail queue.

Use show smtp queue to view all messages currently in the mail queue. These are the messages that have failed and are being saved for retries. You can use smtp retry to retry all queued messages now. The clear smtp queue command deletes all the messages without sending them.

Samples

```
bstnA(cfg-smtp)# retry interval 10
```

sets a 10-minute interval.

```
bstnA(cfg-smtp)# no retry interval
```

reverts to the default interval.

Related Commands

smtp
maximum age
show smtp queue
smtp retry
clear smtp queue

show email-event

Purpose You can define one or more SNMP-trap events to be emailed to one or more users, as well as the thresholds for sending each event. One collection of events, thresholds, and users can be applied to technicians with a certain *role*. Use the show email-event command to see the configuration for one or all email event roles.

Mode (all)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show email-event {role-name | all}`

role-name (1-128 characters) is one role (for example, "tech-support").

all selects all roles.

Guidelines The output contains up to four tables per email event role.

SMTP shows a summary of the SMTP settings for this ARX:

- From is used in the "From" field of all emails. You can set this with from (cfg-smtp).
- Local Mail Server is the SMTP server that is the next hop for all emails. You can set this with mail-server.

Email Event describes the configuration of one email-event role:

- Role is the name of the email-event role, set with the email-event command.
- Description shows the optional description for this role. You can use the description (cfg-email-event) command to set (or change) this description.
- To is a comma-separated list of email recipients. You can add one with the mail-to (cfg-email-event) command.
- Admin State shows whether or not this role is enabled. Use enable (cfg-email-event) to enable the role.

The next table only appears if at least one event group is defined. It contains one row per event group, with the following columns in each row:

- Group Name is set with the group (cfg-email-event) command.
- Threshold Counter is either a number or "n/a." If this many events from the group occur, the system sends an email. This is set with the threshold-counter option in the above group command.
- Threshold Interval (every) is either time value (such as "2 hours" or "7 days") or "n/a." If any events from the group occur during this time, the system sends an email at the end of the time interval. This is set with the threshold-interval option in the group command.

Guidelines (Cont.) The next table only appears if at least one individual event is defined. It contains one row per event, with the following columns in each row:

Group Name is set with the group ... event command.

Event Name is set with the same command.

Threshold Counter is either a number or "n/a." If this many instances of the event occur, the system sends an email. This is set with the threshold-counter option in the group ... event command.

Threshold Interval (every) is either time value (such as "10 minutes" or "12 hours") or "n/a." If any instances of this event occur during this time, the system sends an email at the end of the time interval. This is set with the threshold-interval option in the group ... event command.

Samples

`bstnA# show email-event all`

shows all email event roles. See Figure 13.2 on page 13-19 for sample output.

`bstnA# show email-event noc3`

shows the configuration for one email event role. See Figure 13.3 on page 13-20 for sample output.

Related Commands

email-event
mail-to (cfg-email-event)
description (cfg-email-event)
group (cfg-email-event)
group ... event
enable (cfg-email-event)
smtp
from (cfg-smtp)
mail-server

Figure 13.2 *Sample Output: show email-event all*

```
bstnA# show email-event all

SMTP
----
From             : admin@wwmed.com
Local Mail Server: email1.wwmed.com


Email Event
-----------
Role       : noc3
Description: support team at NOC3
To         : juser@wwmed.com,jqpublic@wwmed.com
Admin State: Enabled

Group Name                Threshold
                      Counter  Interval(every)
-------------------- --------- ---------------
chassis              1         n/a
metadata             5         n/a

Group Name                 Event Name                 Threshold
                                                   Counter Interval(every)
```

```
------------------- ------------------------------- ------- ---------------
metadata            online                          2       n/a
redundancy          ha-pair-qd-offline              1       n/a
storage             share-online                    3       n/a
storage             share-remove-complete           1       n/a


Email Event
-----------
Role       : tech-support
Description: Built-In
To         : e-support@acopiasupport.com
Admin State: Enabled

Group Name                Threshold
                      Counter  Interval(every)
-------------------- --------- ---------------
chassis              1         n/a
```

Figure 13.3 *Sample Output: show email-event noc3*

```
bstnA# show email-event noc3

SMTP
----
From             : admin@wwmed.com
Local Mail Server: email1.wwmed.com


Email Event
-----------
Role       : noc3
Description: support team at NOC3
To         : juser@wwmed.com,jqpublic@wwmed.com
Admin State: Enabled

Group Name                Threshold
                      Counter  Interval(every)
-------------------- --------- ---------------
chassis              1         n/a
metadata             5         n/a

Group Name                 Event Name                   Threshold
                                                      Counter Interval(every)
------------------- ------------------------------- ------- ---------------
metadata            online                          2       n/a
redundancy          ha-pair-qd-offline              1       n/a
storage             share-online                    3       n/a
storage             share-remove-complete           1       n/a
```

show email-severity

Purpose You can set up SNMP-trap events to be emailed to one or more users, and you can use a separate command to set the severity of each event's email message. The severity is expressed in a hidden field in the email header, which email servers can use for flagging and prioritizing email messages. Use the show email-severity command to see the severities for this switch's email events.

Mode (all)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show email-severity {`*`event-name`*` | all | non-default}`**

event-name (1-64 characters) is one event (for example, "module-failure").

all shows all events.

non-default shows any events that have a non-default severity.

Guidelines The output is a table with two columns: Event Name and Severity. You can use the email-severity command to reset the severity.

Samples `bstnA# show email-severity all`

shows all email events and their current severity settings. See Figure 13.4 on page 13-21 for sample output.

`bstnA# show email-severity non-default`

shows the email events whose severities have been reset. See Figure 13.5 on page 13-25 for sample output.

`bstnA# show email-severity nvram-battery-failure`

shows the severity for a particular email event. See Figure 13.6 on page 13-25 for sample output.

Related Commands email-severity
email-event

Figure 13.4 *Sample Output: show email-severity all*

```
bstnA# show email-severity all

   Event Name                         Severity
---------------------------------- ---------
System oom-error                   major
ad-policy_violation-clear          normal
ad-policy_violation-raise          major
archive-free-space-threshold-clear  normal
archive-free-space-threshold-raise  major
archive-offline                    major
archive-online                     normal
archive-remove                     normal
archive-write-access-clear         warning
archive-write-access-fail          warning
```

```
auto-diagnostics-failed             major
auto-diagnostics-failed-clear       normal
bounce-limit                        critical
bounce-limit-clear                  normal
clock-sync-fail-clear               normal
clock-sync-fail-raise               major
core                                critical
core-duplicate                      normal
cpu-halt                            major
cpu-temperature-failure             critical
cpu-temperature-status              normal
db-upgrade-fail-clear               normal
db-upgrade-fail-raise               critical
directory-attribute-inconsistency-clear normal
directory-attribute-inconsistency-clear-all normal
directory-attribute-inconsistency-raise major
directory-import-stalled-clear      normal
directory-import-stalled-clear-all  normal
directory-import-stalled-raise      major
disk-control-failure                critical
disk-control-status                 normal
disk-failure                        critical
disk-state                          normal
dnasInstanceStartupFailure          major
dnasforcerecovery                   major
dns-name-update-cancel              normal
dns-name-update-clear               normal
dns-name-update-raise               minor
dns-offline                         major
dns-online                          normal
dns-remove                          normal
down-rev-ad-forest-level-clear      normal
down-rev-ad-forest-level-raise      major
down-rev-ntlm-auth-srv-clear        normal
down-rev-ntlm-auth-srv-raise        major
dr-config_replication-clear         normal
dr-config_replication-raise         major
fan-failure                         critical
fan-status                          normal
filer-errors-clear                  major
filer-errors-raise                  major
filer-slow-clear                    major
filer-slow-raise                    major
firmware-mismatch                   normal
firmware-mismatch-raise             major
firmware-upgrade-initiated          critical
free-space-threshold                critical
free-space-threshold-clear          normal
gateway-offline                     critical
gateway-online                      normal
gateway-remove                      normal
ha-pair-cluster-offline             major
ha-pair-cluster-online              normal
ha-pair-qd-offline                  major
ha-pair-qd-online                   normal
ha-pair-qdisk-freespace-low         warning
ha-pair-qdisk-freespace-ok          normal
ha-pair-version-autosync-clear      normal
ha-pair-version-autosync-raise      major
ha-pair-version-mismatch-clear      normal
ha-pair-version-mismatch-raise      major
kerberos-cache-threshold-clear      major
```

```
kerberos-cache-threshold-cross      major
kerberos-dc-offline                 major
kerberos-dc-online                  normal
kerberos-dc-remove                  normal
kernel-nmi-error                    major
license-expired-clear               normal
license-expired-raise               major
license-ha-pair-different-clear     normal
license-ha-pair-different-raise     warning
license-ha-pairing-disabled-clear   normal
license-ha-pairing-disabled-raise   major
license-not-found-clear             normal
license-not-found-raise             warning
license-pending-expiration-clear    normal
license-pending-expiration-raise    warning
license-platform-limit-clear        normal
license-platform-limit-raise        warning
license-protol-qty-clear            normal
license-protol-qty-raise            warning
logging-failure-clear               normal
logging-failure-raise               major
metalog-latency-clear               normal
metalog-latency-raise               major
module-failure                      critical
module-status                       normal
no-ntlm-authdc-clear                normal
no-ntlm-authdc-raise                major
non-critical-resource-failure       major
non-critical-resource-failure-clear normal
nsckreimport                        normal
nsm-resource-threshold              major
nsm-resource-threshold-clear        major
nsm-standby                         minor
nsm-standby-clear                   normal
nsm-warm-restart                    major
ntlm-auth-srv-offline               major
ntlm-auth-srv-offline-clear         normal
ntp-reachable                       normal
ntp-unreachable                     normal
nvram-battery-degraded              major
nvram-battery-degraded-clear        major
nvram-battery-failure               critical
nvram-battery-status                normal
nvram-ecc-error                     critical
nvram-ecc-error-clear               normal
nvram-not-saved                     normal
offline                             critical
om-transactions-threshold           major
om-transactions-threshold-clear     normal
online                              normal
peer-critical-resources-failed      major
peer-critical-resources-healthy     normal
policyruleinlinequeueoverflow       warning
power-failure                       critical
power-status                        normal
prewin2k-mismatch                   major
prewin2k-mismatch-clear             normal
raid-verify-clear                   normal
raid-verify-raise                   critical
ram-ecc-correctable-error           major
ram-ecc-error                       major
ram-missing-clear                   normal
```

```
ram-missing-raise                   critical
res-file-near-full-cancel           warning
res-file-near-full-clear            warning
res-file-near-full-raise            warning
sam-reference-offline-clear         warning
sam-reference-offline-raise         warning
server-offline                      major
server-online                       normal
server-remove                       warning
service-acl-update-failure          major
service-acl-update-success          normal
service-errors-clear                major
service-errors-raise                major
service-offline                     major
service-online                      normal
service-rejoin-required-clear       normal
service-rejoin-required-raise       major
service-removed                     normal
service-slow-clear                  major
service-slow-raise                  major
shadowmetadatasharefreespaceerrorclear normal
shadowmetadatasharefreespaceerrorraise critical
shadowmetadatasharefreespacewarnclear normal
shadowmetadatasharefreespacewarnraise major
share-feature-mismatch-clear        warning
share-feature-mismatch-raise        warning
share-import-complete               normal
share-import-failure                major
share-logon-failure-clear           major
share-logon-failure-raise           major
share-offline                       major
share-online                        normal
share-probe-upgrade-clear           normal
share-probe-upgrade-raise           normal
share-remove-complete               normal
share-remove-failure                major
share-timeskew-clear                warning
share-timeskew-raise                warning
share-write-access-clear            warning
share-write-access-fail             warning
sharefreespacethresholdclear        normal
sharefreespacethresholdraise        major
snapshot-op-complete                normal
snapshot-op-fail                    warning
snapshot-op-start                   normal
spn-alias-update-clear              normal
spn-alias-update-raise              minor
subshare-export-degraded-clear      major
subshare-export-degraded-raise      major
suspend-failover-clear              normal
suspend-failover-raise              major
system-bus-error                    major
system-resource-threshold           major
system-resource-threshold-clear     major
temperature-failure                 critical
temperature-status                  normal
tenG-Phy-Unsupported                major
vcifssvcacctclear                   normal
vcifssvcacctraise                   major
vcifsworkjamclear                   normal
vcifsworkjamraise                   major
warm-start                          critical
```

```
x2-transceiver-fault               major
x2-transceiver-fault-clear         normal
xiplip-inconsistency-clear         normal
xiplip-inconsistency-raise         critical
```

Figure 13.5 *Sample Output: show email-severity non-default*

```
bstnA# show email-severity non-default

   Event Name                       Severity
---------------------------------- ---------
auto-reboot                        critical
cpu-failure                        critical
server-offline                     critical
```

Figure 13.6 *Sample Output: show email-severity nvram-battery-failure*

```
bstnA# show email-severity nvram-battery-failure

   Event Name                       Severity
---------------------------------- ---------
nvram-battery-failure              critical
```

show smtp queue

Purpose Use the show smtp queue command to see the email messages that are queued for delivery, if any.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show smtp queue`

Guidelines This command is used to see messages that failed delivery. email messages are in the queue for a very short time before they are sent out to the mail server; they stay in the queue longer if the delivery failed. Therefore, if all messages deliver on the first try, this command shows no messages in the queue.

The output is a table with two rows per message. The top row contains the following fields:

Status is "pending," to indicate that the message is waiting for final delivery.

From can be reset (for future emails) with from (cfg-smtp).

Size is in bytes if no unit (K, M, G, or T) appears after the number. K is for KiloBytes (1024 bytes), M is for MegaBytes (1024*1024 bytes), and so forth.

The bottom row contains two more fields:

Time is when the message was created.

To can be reset for future emails with to.

The ARX retries a message every few minutes (set by retry interval) for the first two hours. This retry interval increases geometrically if the failures persist. After some days without success (set by maximum age), the switch deletes the message from the mail queue. You can use smtp retry to retry all queued messages now. The clear smtp queue command deletes all the messages without sending them.

Use show smtp status to see details on the most-recent message delivery, and to view the current configuration for SMTP.

Sample `bstnA# show smtp queue`

shows the queue of pending email messages, if any. See Figure 13.7 for sample output.

Related Commands
smtp
to
from (cfg-smtp)
smtp retry
show smtp status

Figure 13.7 *Sample Output: show smtp queue*

```
bstnA# show smtp queue

show smtp queue
Status     From                            Size
    Time                         To
-----------------------------------------------
```

```
pending     admin@wwmed.com                461.00
    10/16/2007 13:33:51 GMT     juser@wwmed.com
pending     admin@wwmed.com                461.00
    10/16/2007 13:37:1 GMT      juser@wwmed.com
```

show smtp status

Purpose Use the show smtp status command to see the status of the most recent outbound email, as well as all SMTP-configuration settings.

Mode (all)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show smtp status`

Guidelines SMTP Last Transfer shows the last email message that was collected for delivery. This is not the most-recent retry; it is the most-recent message (or test) that was generated.

Last file is the file name and size for the attachment file.

Status shows the status of the most-recent delivery:

- Success
- Delivered to *mail-server*
- Deferred means that the message has been added to the mail queue.
- Configuration error means that the configuration is missing some information. (see the SMTP Current Configuration section, below)
- Could not communicate with mailserver indicates a network problem, a DNS-configuration error, or an SMTP configuration error (such as the wrong name for mail-server).
- Message size exceeds limit means that the local mailserver did not accept the large attachment with the file. You must configure the mail server to allow *very* large email attachments.
- Unable to relay
- Insufficient disk space indicates that there is not enough disk space to add the message to the mail queue.

Guidelines (Cont.)

– Error building mail header,

– SMTP error,

– Internal error mailing file, and

– Error uuencoding file are internal problems.

Message Created is the time when the message was first generated.

The remaining fields show the header for the email message.

SMTP Current Configuration shows the configuration that is in effect. This configuration will be applied to current and future email messages.

To can be set with to.

From can be set with from (cfg-smtp).

Via is the SMTP server that is the next hop for all emails. You can set this with mail-server.

Retry Interval is the time between retries for messages in the mail queue. The ARX uses this interval for the first two hours, then starts using larger intervals if the failures persist. You can set this initial interval with the retry interval command.

Maximum Age is maximum number of days to keep a message in the mail queue before deleting it. Use maximum age to set this value.

Use show smtp queue to see any pending email messages. To send a test message, use smtp test message. If messages are failing, you can run smtp test server to test SMTP communication with the local email server.

Sample

`bstnA# show smtp status`

shows the SMTP status for the current switch. See Figure 13.8 for sample output.

Related Commands

smtp
to
from (cfg-smtp)
mail-server
show smtp queue
smtp test message
smtp test server

Figure 13.8 *Sample Output: show smtp status*

```
bstnA# show smtp status
SMTP Last Transfer
------------------

Last file:       AcopiaTest (61 bytes)
Status:          Sent to email1.wwmed.com
Message Created: 9/15/2009 5:11:38 GMT
To:              juser@f5.com
From:            admin@wwmed.com
Via:             email1.wwmed.com
Subject:         Test diags for bstnA running 5.01.000.11891

SMTP Current Configuration
--------------------------

To:              juser@f5.com
From:            admin@wwmed.com
```

```
Via:             email1.wwmed.com
Retry Interval:  10  minutes
Maximum Age:     30  days
```

show smtp welcome

Purpose You can use the smtp welcome command to send an introductory email to all users of an email-event role. Use the show smtp welcome command to see the welcome message for this introductory email.

Mode (all)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show smtp welcome`**

Guidelines This command displays the welcome message that appears at the top of an smtp welcome email.

Sample

```
bstnA# show smtp welcome
Welcome to the "<EMAIL_EVENT_ROLE>" role.
The host <HOSTNAME> will send email to you based on the following
events:
```

shows the outline of the welcome message. The <EMAIL_EVENT_ROLE> variable is replaced by the role's name (set with email-event), and the <HOSTNAME> is replaced by the ARX's host name (set with hostname).

Related Commands

smtp welcome
smtp
email-event
hostname

smtp

Purpose

The ARX uses the Simple Mail Transfer Protocol (SMTP) to send email messages to other machines. Use the smtp command to begin SMTP configuration.

Use no smtp to erase all SMTP configuration parameters and disable email notifications.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
smtp
no smtp
```

Default

no smtp

Guidelines

You must configure DNS for the SMTP service to work. See the instructions for ip name-server. You must have layer-2 and layer-3 connectivity to your server network(s) before DNS or SMTP can function.

Use this command to enter cfg-smtp mode, where you set up the necessary parameters for email. You can use from (cfg-smtp) to set the "From" field in all out-bound emails (for example, "From: *hostname@domain*"). The mail-server command identifies the local mail server to be used as the next hop for mail messages. Use the to command to set one or more destinations for emails. If an email delivery fails, it retries at the frequency set by retry interval, up to a maximum number of days (maximum age). After the maximum age expires, the ARX deletes the message.

Guidelines: Testing and Debugging

To send a test message, use smtp test message reply-to with a local destination for the message. You can view any undelivered messages with the show smtp queue command. Use the smtp retry and clear smtp queue commands to manage this mail queue. If messages are consistently failing, you can run smtp test server to test SMTP communication with the local email server.

Guidelines: Applications for SMTP

After you finish configuring SMTP, you can configure certain system events to be delivered in email messages. Each event is analogous to an SNMP trap; see the ARX SNMP Reference for a full list of ARX Enterprise traps, including the traps that have email support. You can choose individual events or groups of them, one or more email recipients for the events, and thresholds (or a schedule) for sending the emails. All of these components can be assigned to technicians who perform a specific *role*. Use the email-event command to begin configuring one such role. You can create multiple roles by re-using this command.

SMTP is also useful for sending collected diagnostics and other maintenance files over email. After you configure SMTP, use the collect command to collect and send diagnostics. You can also use copy smtp to send log files, reports, or other files that could be useful for maintenance.

Samples

```
bstnA(cfg)# smtp
bstnA(cfg-smtp)# ...
```

starts SMTP configuration.

```
bstnA(cfg)# no smtp
```

deletes all SMTP configuration and disables email notifications.

Related Commands

ip name-server
email-event
collect
copy smtp

smtp retry

Purpose The SMTP *mail queue* holds email messages until they are successfully delivered. Use the smtp retry command to retry all queued messages at once.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `smtp retry`

Default None

Guidelines This is particularly useful if there was a network problem between the ARX and the next-hop mail server (mail-server). After you fix the network problem, you can use this command to re send all queued messages. You can use show smtp queue to see the mail queue before and after the retry.

As an alternative, you can use clear smtp queue to remove all messages from the queue.

For each message that is successfully delivered to the mail server, "Message sent" appears after the command. For each failed delivery, "Message deferred" appears.

When you change a parameter in cfg-smtp mode (see smtp) and then exit the mode, the ARX automatically retries all messages in the queue using the new parameters. There is no need to invoke this command under these circumstances.

Use show smtp status to see details on the most-recent message delivery, and to view the current configuration for SMTP.

Sample

```
bstnA# smtp retry
```

retries all pending email messages, if there are any.

Related Commands

smtp
show smtp queue
show smtp status

smtp test email-event

Purpose Use the smtp test email-event command to test an email-event configuration. This sends a test email to all configured recipients.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `smtp test email-event` ***role-name***

role-name (1-128 characters) is the role to test (for example, "tech-support"). Use show email-event all for a list of all configured email-event roles.

Default None

Guidelines Use this command to send an email message with a small test attachment. the message goes to all mail-to (cfg-email-event) recipients configured for the chosen email-event role. Use show email-event *role-name* to see all recipients for a given role. You can use this to verify the email-event configuration and SMTP setup.

If the message is successfully delivered to the next-hop mail-server, a message appears similar to the following:

```
Message sent to juser@wwmed.com,jqpublic@wwmed.com.
```

An error appears if there is a delivery problem. In this case, you can use show smtp queue to see the test message in the mail queue. To retry the message (perhaps after correcting the network configuration), use smtp retry. Use clear smtp queue to remove all messages from the queue. If the delivery problems persist, you can run an SMTP diagnostic test, smtp test server, and send the output back to F5 for analysis.

Note that the test message has a small attachment, so this does not test any errors that may come up with large attachments. An attachment from copy smtp or collect can be hundreds of MegaBytes.

Use show smtp status to see details on the most-recent message delivery, and to view the current configuration for SMTP.

After a successful test, you can use the smtp welcome command to send a welcome message to all email recipients. This message informs the recipients of the events that will trigger email messages in the future.

Sample

```
bstnA# smtp test email-event noc3

  Message sent to juser@wwmed.com,jqpublic@wwmed.com for role "noc3".
```

successfully sends a test message to the recipients of the "noc3" email event. See Figure 13.9 for a sample email.

Related Commands

email-event
mail-to (cfg-email-event)
show email-event
smtp
show smtp queue
smtp retry
clear smtp queue
smtp test server
show smtp status
smtp welcome

Figure 13.9 *Sample Email from 'smtp test email-event'*

```
From:    admin@wwmed.com
Sent:    Tuesday, September 15, 2009 12:40 AM
To:      Joe User; jqpublic@wwmed.com
Subject: Email Event Test for the event group noc3 for bstnA running 5.01.000.11891

This is a test of the email event mailer "noc3" on bstnA running
5.01.000.11891.


Email will be sent for events in the following groups:
Group Name                Threshold
                      Counter  Interval(every)
-------------------- --------- ---------------
chassis              1         n/a
metadata             5         n/a

Group Name                  Event Name                      Threshold
                                                         Counter Interval(every)
------------------- -------------------------------- ------- ---------------
metadata            online                           2         n/a
redundancy          ha-pair-qd-offline               1         n/a
storage             share-online                     3         n/a
storage             share-remove-complete            1         n/a
policy              shadowmetadatasharefreespaceerrorraise5         n/a
```

smtp test message

Purpose Use the smtp test message command to send a test email message.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `smtp test message [reply-to user@domain]`

user@domain (optional, 1-256 characters) is an email destination for the test-message reply.

Default **reply-to *user@domain*** defaults to the setting of the to command if you omit the option.

Guidelines Use this command to send an email message with a small test attachment to an SMTP server at F5. This tests the SMTP configuration settings and the network connection to the local email server. You can use this to verify the SMTP setup.

The message is delivered to the next-hop mail-server. An error appears if there is a delivery problem. In this case, or if no reply email arrives for 15 minutes or more, you can use show smtp queue to see the test message in the mail queue. Also, check the email filter at the destination mailbox. To retry the message (perhaps after correcting the network or filter configuration), use smtp retry. Use clear smtp queue to remove all messages from the queue. If the delivery problems persist, you can run an SMTP diagnostic test, smtp test server, and send the output back to F5 for analysis.

Note that the test message has a small attachment, so this does not test any errors that may come up with large attachments. An attachment from copy smtp or collect can be hundreds of MegaBytes.

Use show smtp status to see details on the most-recent message delivery, and to view the current configuration for SMTP.

Sample

```
bstnA# smtp test message reply-to jsmith@wwmed.com
```

sends a test message to an SMTP server at F5. If the delivery succeeds, the F5 server will send a reply to "jsmith@wwmed.com."

Related Commands

smtp
to
show smtp queue
smtp retry
clear smtp queue
smtp test server
show smtp status

smtp test server

Purpose At the advice of F5 personnel, use the smtp test server command to test the SMTP-layer connection to the email server.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `smtp test server`

Default None

Guidelines This command runs a series of STMP queries against the next-hop mail-server. After you enter this command, the CLI shows the detailed results of this test.

Run the test when you see email delivery problems. If the SMTP queue has undelivered messages in it (show smtp queue), or if smtp test message fails, you should run this test. Send the results back to F5 for further diagnosis.

Sample `bstnA# smtp test server`

tests the connection to the next-hop email server. See Figure 13.10 for sample output.

Related Commands smtp
show smtp queue
smtp test message

Figure 13.10 *Sample Output: smtp test server*

```
bstnA# smtp test server
  spawn telnet email1.wwmed.com 25
  Trying 192.168.25.209...
  Connected to smtp.wwmed.com.
  Escape character is '^]'.
  EHLO foo
  220 frontend.wwmed.com Internal SMTP Server Thu, 23 Mar 2006 16:02:44 -0500
  250-frontend.wwmed.com Hello [192.168.25.5]
  250-TURN
  250-ATRN
  250-SIZE
  250-ETRN
  250-PIPELINING
  250-DSN
  250-ENHANCEDSTATUSCODES
  250-8bitmime
  250-BINARYMIME
  250-CHUNKING
  250-VRFY
  250-X-EXPS GSSAPI NTLM LOGIN
  250-X-EXPS=LOGIN
  250-AUTH GSSAPI NTLM LOGIN
  250-AUTH=LOGIN
  250-X-LINK2STATE
  250-XEXCH50
  250 OK
```

smtp welcome

Purpose Use the smtp welcome command to send an introductory email message to all users in an email-event configuration. The introductory message informs the recipients of the types of system events they will be receiving through email.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `smtp welcome [role-name]`

role-name (optional, 1-128 characters) is the role to receive the email message (for example, "noc_7"). Use show email-event all for a list of all configured email-event roles.

Default ***role-name*** defaults to all roles if you omit it. That is, the command sends a welcome message to all users in all email-event roles.

Guidelines Use this command to send an introductory email message to the recipients of email events. the message goes to all mail-to (cfg-email-event) recipients configured for the chosen email-event role (or all roles, if no specific role was selected). Use show email-event *role-name* to see all recipients for a given role.

If the message is successfully delivered to the next-hop mail-server, a message appears similar to the following:

```
Message sent to juser@wwmed.com,jqpublic@wwmed.com.
```

Use show smtp status to see details on the most-recent message delivery, and to view the current configuration for SMTP. If the message fails, you can use smtp test email-event to test the email-event configuration and SMTP setup.

The email contains a welcome message and two tables to outline the types of events that the recipient can expect in email messages. The top table shows event groups, and the bottom table contains individual events that will trigger an email. The show smtp welcome command shows the format of the welcome message.

Sample

```
bstnA# smtp welcome noc3

  Message sent to to jqpublic@wwmed.com.
```

successfully sends a welcome message to the recipients of the "noc3" email event. See Figure 13.11 for a sample email.

Related Commands

email-event
mail-to (cfg-email-event)
show email-event
smtp test email-event
clear smtp queue
smtp test server
show smtp status
show smtp welcome

Figure 13.11 *Sample Email from 'smtp welcome'*

```
From:    admin@wwmed.com
Sent:    Tuesday, September 15, 2009 12:40 AM
To:      Joe User; jqpublic@wwmed.com
Subject: Welcome to the noc3 role  for bstnA running 5.01.000.11891

Welcome to the "noc3" role.
The host bstnA will send email to you based on the following events:

Group Name                Threshold
                       Counter  Interval(every)
-------------------- --------- ---------------
chassis              1         n/a
metadata             5         n/a

Group Name                  Event Name                  Threshold
                                                      Counter Interval(every)
------------------- -------------------------------- ------- ---------------
metadata            online                           2         n/a
redundancy          ha-pair-qd-offline               1         n/a
storage             share-online                     3         n/a
storage             share-remove-complete            1         n/a
policy              shadowmetadatasharefreespaceerrorraise5         n/a
```

to

Purpose

Use the to command to set destination address(es) for emails from the ARX.

Use the no to command to revert to the default destination.

Mode

cfg-smtp

Security Role(s)

network-engineer or crypto-officer

Syntax

```
to destination-list
no to
```

destination-list (1-1024 characters) is a comma-separated list of email recipients (for example, “juser@nemed.com” or “jqpublic@wwmed.com,juser@nemed.com”). Surround this field with quotation marks if it contains any spaces.

Default

e-support@f5.com

Guidelines

This sets the destination(s) for emails from the ARX; these addresses appear in the “To” field of the emails. These addresses only apply to emails from collect, copy smtp, and smtp test message. Each email-event role has its own set of email recipients, configured with the mail-to (cfg-email-event) command.

The default address sends messages back to Technical Support at F5. For most running systems, this default is appropriate. You can use this command to test email notifications, or to send email to one or more different support organizations.

To change any address in the list, re-run this command with an entirely new list of addresses.

To see the current email destination(s), use the show smtp status command.

Samples

```
bstnA(cfg-smtp)# to “juser@das1.wwmed.com,e-support@f5.com”
```

sets two destinations for all outgoing messages.

```
bstnA(cfg-smtp)# no to
```

resets the destination back to the default.

Related Commands

smtp
show smtp status
copy smtp
collect

14

RON

heartbeat failure

Purpose The ARX sends periodic heartbeats to check the health of the RON-tunnel connection. Use this command to determine the number of consecutive failures before declaring the tunnel "OFFLINE."

Mode cfg-if-vlan-ron-tnl

Security Role(s) network-engineer or crypto-officer

Syntax **`heartbeat failure`** ***`max-failures`***

max-failures (2-10) is the number of consecutive dropped heartbeats to tolerate *before* declaring a failure.

Default(s) 4

Guidelines Use the heartbeat interval command to set the number of seconds between heartbeats. Use the show ron command to see the current connection state for all RON tunnels.

Sample

```
bstnA(cfg-if-vlan-ron-tnl[25~toPortland])# heartbeat failure 6
```

sets the threshold to 6 consecutive heartbeat failures. This interface declares the tunnel "OFFLINE" if 7 heartbeats fail in a row.

Related Commands

ron tunnel
heartbeat interval
show ron

heartbeat interval

Purpose The ARX sends periodic heartbeats to check the RON-tunnel connection. Use this command to determine the number of seconds between heartbeats.

Mode cfg-if-vlan-ron-tnl

Security Role(s) network-engineer or crypto-officer

Syntax **`heartbeat interval`** ***`seconds`***

seconds (1-30) is the number of seconds between heartbeats.

Default(s) 3

Guidelines Use the heartbeat failure command to set the number of dropped heartbeats to tolerate before declaring the tunnel "OFFLINE." Use the show ron command to see the current connection state for all RON tunnels.

Sample

```
bstnA(cfg-if-vlan-ron-tnl[25~toPortland])# heartbeat interval 15
```

sets a 15-second interval between RON heartbeats.

Related Commands ron tunnel
heartbeat failure
show ron

interface ron

Purpose

** Deprecated **

Use this command to edit a legacy RON interface, one that was created before Software Release 2.0.

Use the no form of the command to remove a legacy RON interface. See the *Guidelines*, below.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
interface ron if-name
no interface ron if-name
```

if-name (1-32 characters) identifies the legacy RON interface.

Default(s)

None

Guidelines

This command is deprecated, and exists only to edit the configuration of legacy RON interfaces and tunnels. Use ron tunnel to create a new RON tunnel, which terminates at an in-band (VLAN) management IP instead of an additional RON IP address. If possible, use ron tunnel to duplicate all legacy RON tunnels, go to the remote switch(es) and run the peer address command with the new RON-tunnel address, then use no interface ron to remove the legacy tunnels.

This command places you in cfg-ron mode. From there, use the ip address (cfg-ron) command to edit the local IP address, and the peer address command to change the peer's address. You can optionally tune the tunnel's health-check parameters by setting the heartbeat interval and the threshold for consecutive heartbeat failure events before declaring the link OFFLINE. Use shutdown to disable the interface.

To view the current state and configuration of all tunnels, use show ron. For a full view of Link-State Advertisements from all switches connected via RON, use show ron database.

If the tunnel is connected and you use no interface ron, the CLI prompts you before disconnecting it. Enter **yes** to continue.

Samples

```
bstnA(cfg)# interface ron haPeer
bstnA(cfg-ron[haPeer])#
```

edits the RON interface, "haPeer."

```
bstnA(cfg)# no interface ron toPhilidelphia
Tunnel ''toPhilidelphia'' is currently connected.

Delete tunnel ''toPhilidelphia''? [yes/no] yes
bstnA(cfg)#
```

removes the legacy "toPhilidelphia" interface.

Related Commands

ron tunnel
show ron tunnel

ip address (cfg-ron)

Purpose

** Deprecated **

Use this command to provide a local IP address for the current (legacy) RON interface.

Mode

cfg-ron

Security Role(s)

network-engineer or crypto-officer

Syntax

`ip address` ***`address mask`*** **`[vlan`** ***`vlan-id`*** **`]`**

address is the IP address you choose for the RON interface (for example, 10.1.99.78).

mask defines the network part of the address (for example, 255.255.255.0).

vlan ***vlan-id*** (optional; 1-4096) specifies a VLAN to carry the tunnel.

Default(s)

vlan-id: 1

Guidelines

Note

This command is unique to a legacy RON interface, which is deprecated in favor of the new ron tunnel*. Whenever possible, duplicate all legacy RON interfaces as RON tunnels, go to the remote switch(es) and run the* peer address *command with the new RON-tunnel address, then return to the local switch to delete the RON interfaces. (RON tunnels re-use an in-band-management IP, whereas RON interfaces require an additional IP address.)*

This address must also be configured at the other end of the RON tunnel as the peer address; the tunnel is not functional until the configurations match at both switches.

This address must be unique from all proxy IPs, VIPs and management IPs.

Samples

`bstnA(cfg-ron[toProv])# ip address 192.168.25.50 255.255.255.0`
sets an IP address of 192.168.25.50 for the “toProv” interface.

`prtlndB(cfg-ron[test])# ip address 192.168.74.73 255.255.255.0 vlan 96`
sets the IP address for the “test” interface. This address is on VLAN 96.

Related Commands

interface ron
ron tunnel
peer address

ip private subnet reassign

Purpose Two switches in a RON are said to have a *conflict* if their private IP subnets are the same. Use this command to reassign a private subnet to the current switch and reboot it.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax **`ip private subnet reassign`**

Default(s) None

Guidelines Use the show ron or show ron conflicts command to confirm that this switch has a private subnet that conflicts with another.

Note that this command reboots the current chassis. The CLI prompts for confirmation before rebooting; enter **yes** to continue.

Sample

```
bstnA(cfg)# ip private subnet reassign
Reassign a new, unused, private subnet and reboot the chassis? [yes/no]
yes
...
```

Related Commands ron tunnel
show ron
show ron conflicts

peer address

Purpose Use this command to provide the IP address for the remote end of the RON tunnel.

Mode cfg-if-vlan-ron-tnl

Security Role(s) network-engineer or crypto-officer

Syntax **`peer address`** ***`remote-address`***

remote-address is the remote IP address (for example, 10.1.33.8) for the in-band (VLAN) interface at the other end of the tunnel.

Default(s) None

Guidelines Each end of the RON tunnel terminates at an in-band (VLAN) management interface, created with the interface vlan command. This command identifies the management interface's ip address (cfg-if-vlan) at the other end of the tunnel. At the other end of the tunnel, the peer address points back to the local in-band-management IP.

Sample

```
prtlndA(cfg-if-vlan-ron-tnl[74~toBoston])# peer address 192.168.25.5
```

sets an IP address of 192.168.25.5 for the remote end of the "toBoston" tunnel.

Related Commands ron tunnel
ip address (cfg-if-vlan)

rconsole

Purpose Use the rconsole command to start a new CLI session on an ARX at the other end of a RON tunnel.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`rconsole`** ***`hostname`*** **`[`*`username`*`]`**

hostname (1-128 characters) is the remote switch name.

username (optional, 1-32 characters) is a valid administrative user account on the remote ARX. If you omit this, the command uses the name you used to log into the current CLI session. Note that this administrative account may not exist at the remote ARX, or may have a different password.

Default(s) ***username*** - the administrative account you used to log into the local CLI.

Guidelines The rconsole command starts a new CLI session on a remote ARX. This occurs through a Secure Shell (SSH). The other switch must be reachable using a RON tunnel, and the name of the switch must be known to RON. To see the current switch names available through RON, use the show ron command. Those switches showing the connection status 'ONLINE' are available through the rconsole command.

The remote CLI prompts you for a password. This is the password for the ***username*** account on the remote switch; administrative accounts are configured independently at every switch in the RON (see user). Passwords are not guaranteed to be consistent throughout a RON.

Sample

```
bstnA# rconsole prtlndA admin
Password: myP@55w0RD
prtlndA>
```

starts a CLI session on the switch, 'prtlndA' connected through the RON. This login uses "admin" as a user name.

Related Commands

ron tunnel
show ron
user

ron evict

Purpose Use this command to remove an offline ARX from the RON.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`ron evict`** ***`host-name`***

host-name (1-128 characters) is the name of the switch to remove from the RON.

Default(s) None

Guidelines If the switch is to be replaced, we recommend that you avoid this command. The show ron command shows the UUID for the former RON member, and this number is very important for a smooth switch replacement.

The switch must be offline to be evicted. This means that there can be no working RON tunnels to the switch. The show ron command shows whether or not the switch is offline.

Sample

```
bstnA# ron evict pawtucket
```

removes the switch, "pawtucket," from the RON.

Related Commands ron tunnel
show ron

ron tunnel

Purpose

Use this command to create one end of a resilient overlay network (RON) tunnel. RON tunnels are used for communication between two or more ARXes.

Use the no form of the command to remove a RON-tunnel interface (see *Guidelines* below).

Mode

cfg-if-vlan

Security Role(s)

network-engineer or crypto-officer

Syntax

ron tunnel *name*
no ron tunnel *name*

name (1-32 characters) is a name you choose for the RON tunnel.

Default(s)

None

Guidelines

Use this command to begin the configuration at one end of a RON tunnel. A RON tunnel terminates at an in-band (VLAN) management interface; use interface vlan to create such an interface. The in-band-management IP is the IP address of the RON tunnel's local end.

This command places you in cfg-if-vlan-ron-tnl mode. From there, use the peer address command to identify the peer's address (that is, the IP address of the peer's in-band management interface). You can optionally tune the tunnel's health-check parameters by setting the heartbeat interval and the threshold for consecutive heartbeat failure events before declaring the tunnel "OFFLINE." Then use no shutdown (cfg-if-vlan-ron-tnl) to enable the tunnel interface. To start traffic on the tunnel, repeat this process (reversing the IP addresses) at the other end.

Multiple RON tunnels can terminate in a single in-band management interface.

After you connect two switches with a RON tunnel, you can access the peer switch's CLI through the rconsole command.

The show ron command shows a high-level status for the entire RON. To view the current state and configuration of a tunnel, use show ron tunnel. For a full view of Link-State Advertisements from all switches connected via RON, use show ron database. Use show ron conflicts to see which switches (if any) have a private-subnet conflict; to resolve a conflict, use ip private subnet reassign.

Samples

```
bstnA(cfg)# interface vlan 89
bstnA(cfg-if-vlan[89])# ron tunnel toEllesworth
bstnA(cfg-if-vlan-ron-tnl[toEllesworth])#
```

creates the tunnel, "toEllesworth."

```
bstnA(cfg-if-vlan[89])# no ron tunnel toPhilidelphia
Tunnel ''toPhilidelphia'' is currently connected.
Delete tunnel ''toPhilidelphia''? [yes/no] yes
bstnA(cfg-if-vlan[89])#
```

removes the tunnel, "toPhilidelphia."

Related Commands

interface vlan
show ron tunnel

show ron

Purpose Use this command to display the current RON configuration and status.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax `show ron [member]`

member (optional, 1-128 characters) specifies the hostname for one member of the RON. If you omit this, the command shows all RON members.

Default(s) None

Guidelines This command displays a two-line entry for each member of the RON. The first line identifies the switch and its high-level status.

Switch Name identifies the member switch.

HA Peer Switch for the switch, if there is one. This is the switch's redundant peer.

Uptime is the time since the switch's last reboot.

The second line has information about the RON (and possibly the redundant pair).

Status is ONLINE, OFFLINE, SUBNET CONFLICT, or unknown. The SUBNET CONFLICT status indicates that the private subnet is the same as some other switch in the RON; use show ron conflicts to find which switches conflict with which.

UUID is the Universally-Unique ID for the switch. All of the shares owned by this switch are marked with this UUID. In a redundant pair, a share is owned by a switch when its volume's volume-group is set at the switch. You set a switch's UUID during installation.

Management Addr is the Out-of-band MGMT interface, if configured. Otherwise, it is the address of the in-band (VLAN) management interface of the lowest-numbered VLAN.

If you choose one ***member***, the output focuses on that host only. The fields are the same, presented in a different format.

Use ron tunnel to create one end of a RON tunnel.

Samples `prtlndA# show ron`

displays all hosts in the RON. Figure 14.1 shows a sample.

`prtlndA# show ron prtlndB`

displays the RON status for one redundant peer. Figure 14.2 shows a sample.

Related Commands

ron tunnel
rconsole
show ron conflicts
show ron database
show ron route
volume-group

Figure 14.1 *Sample Output: show ron*

```
prtlndA# show ron

Switch Name             HA Peer Switch                              Uptime
Status                  UUID                                 Management Addr
-------------------------------------------------------------------------------
bstnA                   (None)                              0 days, 02:00:57
ONLINE                  d9bdece8-9866-11d8-91e3-f48e42637d58          10.1.1.7

gffstnA                 (None)                              0 days, 02:30:46
ONLINE                  e5d870ae-571e-1352-916b-ef324fbc05a2        10.1.49.60

minturnA                (None)                              0 days, 02:30:10
ONLINE                  3d17e8ce-571e-11dc-9852-ef323fbb290f        10.1.27.69

provA                   (None)                              0 days, 02:27:55
ONLINE                  db922942-876f-11d8-9110-8dtu78fc8329        10.1.38.19

prtlndA                 prtlndB                             0 days, 02:24:30
ONLINE                  876616f6-79ac-11d8-946f-958fcb4e6e35        10.1.23.11

prtlndB                 prtlndA                             0 days, 02:22:16
ONLINE                  64dcab94-a2b6-11d8-9d25-bf2c991c83f9        10.1.23.12
```

Figure 14.2 *Sample Output: show ron prtlndB*

```
prtlndA# show ron prtlndB

  Switch Name:           prtlndB
  HA Peer Switch:        prtlndA
  Status:                ONLINE
  Uptime:                0 days, 00:14:18
  UUID:                  64dcab94-a2b6-11d8-9d25-bf2c991c83f9
```

show ron conflicts

Purpose Two switches in a RON are said to have a *conflict* if their private IP subnets are the same. The show ron conflicts command shows all conflicting switches in the current RON.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax **`show ron conflicts`**

Default(s) None

Guidelines This shows a table of RON conflicts, one row per conflicting switch. The Accessible Switch is available from the local switch, and the Conflicting Switch is not; some other switch in the RON might show the same conflict with the switch roles reversed. Each of these switches can only communicate with a limited number of peers in the RON, if any. To correct the problem and bring conflicting switches fully into the RON, go to one conflicting switch and use ip private subnet reassign.

Sample

```
prtlndA> show ron conflicts

Accessible Switch        Conflicting Switch
-------------------------------------------------
BIG6000-WEST             ARX-500-SJC
```

displays all private-subnet conflicts in the RON.

Related Commands ip private subnet reassign
ron tunnel

show ron database

Purpose Use this command to display the RON-routing database.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax `show ron database`

Default(s) None

Guidelines You can connect one ARX to another with a RON tunnel: use this command to show routing information and statistics for each ARX connected to the current switch. The output displays one table per host switch. Each table contains the following information:

Hostname

Status is "Current," "Stale," or "ERROR." This represents the status of the peer's *Link-State Advertisement* (LSA). Peers exchange periodic LSAs with information about the tunnel state(s); this table expresses the latest LSA received from the Hostname.

"Stale" indicates that the peer is unreachable, so at least one LSA was missed since the one shown.

"ERROR" indicates that this peer has the same private subnet as another peer in the RON. This is an unsupportable configuration. An ERROR message at the bottom of the output shows which peers have the conflict. Go to either switch's CLI and use ip private subnet reassign to change its subnet.

Serial # is the serial number for the peer's latest LSA. This number increments whenever the switch receives a new LSA from this peer.

Age is the number of seconds since the last LSA was received from the remote peer.

Private Subnet(s) is the host's private subnet. If the host is part of a redundant pair, this also shows the private subnet for the host's peer.

A sub table shows all tunnels configured for the host. Each tunnel appears in one row with the following information:

Tunnel is the name of the tunnel, set with the ron tunnel command.

Peer is the host switch at the other end of the tunnel. You can change this with the cfg-if-vlan-ron-tnl peer address command.

State is Down, Connecting, Connected, Shutdown, or Unknown. A new tunnel transitions from "Shutdown" to "Connecting" to "Connected." "Down" indicates link failure: too many consecutive heartbeats were dropped (see heartbeat failure and heartbeat interval). "Shutdown" indicates that the tunnel was disabled with the shutdown (cfg-if-vlan-ron-tnl) command.

RTT(ms) is the average Round-Trip Time (RTT) through the tunnel, in milliseconds.

Loss(%) is the percentage of packets lost in the tunnel.

TCP (Kb/s) is the estimated TCP throughput (in Kilobits per second) on the tunnel.

Loss*RTT is reserved for future use.

Sample `prtlndA# show ron database`

displays the full RON database at the switch named "prtlndA." See Figure 14.3 on page 14-16 for sample output.

Related Commands ron tunnel
show ron tunnel
show ron
rconsole

Figure 14.3 *Sample Output: show ron database*

```
prtlndA# show ron database

Hostname          Status   Serial #    Age   Private Subnet(s)
------------------------------------------------------------------------------
bstnA             Current  4AEFC347    535   169.254.11.0/24

  Tunnel          Peer            State      RTT(ms) Loss(%) TCP(Kb/s) Loss*RTT
  ----------------------------------------------------------------------------
  toGoffstown     gffstnA         Connected     0.1     0.6   122070        1
  toMinturn       minturnA        Connected     0.1     0.6   122070        1
  toPortland      prtlndA         Connected     0.1     0.6   122070        1
  toPortlandB     prtlndB         Connected     0.1     0.6   122070        1
  toProvidence    provA           Connected     0.1     0.6   122070        0


Hostname          Status   Serial #    Age   Private Subnet(s)
------------------------------------------------------------------------------
gffstnA           Current  4AEFBDAC    340   169.254.104.0/24

  Tunnel          Peer            State      RTT(ms) Loss(%) TCP(Kb/s) Loss*RTT
  ----------------------------------------------------------------------------
  toBoston        bstnA           Connected     0.1     0.6   122070        0
  toPortland      prtlndA         Connected     0.1     0.6   122070        1
  toProvidence    provA           Connected     0.1     0.6   122070        1


Hostname          Status   Serial #    Age   Private Subnet(s)
------------------------------------------------------------------------------
minturnA          Current  4AEFBD94    230   169.254.80.0/24

  Tunnel          Peer            State      RTT(ms) Loss(%) TCP(Kb/s) Loss*RTT
  ----------------------------------------------------------------------------
  toBoston        bstnA           Connected     0.1     0.6   122070        1
  toPortland      prtlndA         Connected     0.1     0.6   122070        1


Hostname          Status   Serial #    Age   Private Subnet(s)
------------------------------------------------------------------------------
provA             Current  4AEFBD63    685   169.254.127.0/26

  Tunnel          Peer            State      RTT(ms) Loss(%) TCP(Kb/s) Loss*RTT
  ----------------------------------------------------------------------------
  toBoston        bstnA           Connected     0.1     0.6   122070        1
  toGoffstown     gffstnA         Connected     0.1     0.6   122070        0


Hostname          Status   Serial #    Age   Private Subnet(s)
------------------------------------------------------------------------------
```

```
prtlndA           Current  4AEFBD8C    695   169.254.66.0/24   169.254.96.0/24

  Tunnel          Peer            State      RTT(ms) Loss(%) TCP(Kb/s) Loss*RTT
  ------------------------------------------------------------------------------
  Redundancy      prtlndB         Connected        0     0.6   122070        0
  toBoston        bstnA           Connected      0.1     0.6   122070        1
  toGoffstown     gffstnA         Connected      0.1     0.6   122070        1
  toMinturn       minturnA        Connected      0.2     0.6   122070        1


Hostname          Status   Serial #    Age   Private Subnet(s)
------------------------------------------------------------------------------
prtlndB           Current  4AEFBE05    620   169.254.96.0/24   169.254.66.0/24

  Tunnel          Peer            State      RTT(ms) Loss(%) TCP(Kb/s) Loss*RTT
  ------------------------------------------------------------------------------
  Redundancy      prtlndA         Connected        0     0.6   122070        0
  toBoston        bstnA           Connected      0.1     0.6   122070        1
```

show ron route

Purpose This command shows the IP-routing table used by the RON process.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax **`show ron route`**

Default(s) None

Guidelines This shows a table of RON routes, one row per tunnel. Each row contains four fields:

Destination is the host name of the peer switch at the other end of the tunnel.

Subnet is the private-IP subnet of the peer switch.

Via Tunnel is tunnel for sending packets to the Subnet.

Milliseconds is average round-trip time, through the tunnel and back.

Sample `prtlndA>` **`show ron route`**
displays the RON routing table. See Figure 14.4 for sample output.

Related Commands ron tunnel

Figure 14.4 *Sample Output: show ron route*

```
prtlndA> show ron route

Default Policy
--------------

  Destination            Subnet           via Tunnel        Milliseconds
  ------------------------------------------------------------------------
  bstnA               169.254.11.0/24     toBoston               0.1
  gffstnA            169.254.104.0/24     toGoffstown            0.1
  minturnA            169.254.80.0/24     toMinturn              0.2
  provA              169.254.127.0/26     toGoffstown            0.2
  prtlndB             169.254.96.0/24     Redundancy               0
```

show ron tunnel

Purpose Use this command to display RON-tunnel configuration.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax `show ron tunnel [name | redundancy | all]`

name (optional, 1-32 characters) identifies a particular RON tunnel to display. If you omit this, the command shows a summary of all RON tunnels.

redundancy (optional) is the name of an automatically-generated tunnel. This only appears on an ARX-1500 or ARX-2500 with a redundant peer. This tunnel carries redundancy-related traffic over the redundancy link that connects the redundant pair. The ARX creates this tunnel when you enter the redundancy command to join a redundant pair.

all shows details for all RON tunnels on the ARX.

Default(s) None.

Guidelines A Resilient Overlay Network (RON) connects multiple ARX devices. One RON tunnel connects two of them; use the ron tunnel command to create a RON tunnel. This command shows the configuration and state of RON tunnels.

Guidelines: Summary Output The show ron tunnel command (without a specific tunnel or the all keyword) shows a table of all RON tunnels. Each row summarizes the tunnel configuration and state:

Name is the tunnel name, set by the ron tunnel command.

State is Shutdown, Connecting, Connected, Unreachable, No Response, Mismatch, Error, or Unknown. A new tunnel transitions from "Shutdown" to "Connecting" to "Connected." The error states are listed here.

- Unreachable means there is no local route to the tunnel's remote endpoint. Use ip route to create a static route.
- No response indicates that the remote peer is not responding to standard RON packets (such as RON heartbeats), nor is it responding to lower-level ICMP pings. Either the network is down or the remote switch is down.
- Mismatch means that ICMP pings worked but the standard RON heartbeats did not. This implies the local configuration of the remote IP address is wrong (see peer address), or that the remote switch does not have a tunnel coming back to the local switch.

  When the tunnel is first coming up, it is in this state between the time that the lower layers are connected and RON processes start. This is normal.
- Shutdown indicates that the tunnel was disabled with the shutdown (cfg-if-vlan-ron-tnl) command.
- Error should never appear. Contact F5 if you see this.

Guidelines: Summary Output (Cont.)

Interface shows the in-band management interface that serves as the local end of the tunnel (created with interface vlan).

Remote Addr is the IP address of the peer's end of the tunnel. Use the peer address command to change this. (The local address is the address of the tunnel's in-band management interface, set with ip address (cfg-if-vlan).)

Up Time is the amount of time that the RON tunnel has been "Connected."

Guidelines: Detailed Output

Include the interface name (or all) to display details for the tunnel(s):

Name is the tunnel name.

Peer is the host name of the remote peer.

Tunnel State is the same as State in the summary version.

Uptime shows how long the tunnel's state has been "Connected."

Interface shows the in-band management interface that serves as the local end of the tunnel (created with interface vlan).

Remote Address is the IP addresses at the other end of the tunnel. Use peer address to change the remote address.

Security Policy is reserved for future use.

Ping Fail Limit is the number of consecutive heartbeat failures to tolerate before declaring the Tunnel State "No response." Use heartbeat failure to change this threshold.

Ping Interval is the number of seconds between heartbeats. Use heartbeat interval to change it.

Round Trip Time is the average number of milliseconds for a packet to go through the tunnel and back.

Packet Lost Rate is the percentage of packets lost in the tunnel.

TCP Throughput is the tunnel's estimated throughput in bytes per second.

Loss RTT Product is reserved for future use.

RON Packets In and RON Packets Out counts the packets that are directly-related to RON, such as heartbeats and Link-State Advertisements (LSAs).

Data Packets In and Data Packets Out counts all other packets, including shadow-copy data.

Data Bytes In and Data Bytes Out is the total bytes (of Data Packets) exchanged.

Local Processor is the CPU where the tunnel terminates, in *slot.processor* format. Use show processors to see a list of all processors and their roles. This does not appear for the ARX-1500 or ARX-2500; all RON tunnels terminate at processor 1.1 on those platforms.

Last Error Code is 0 (zero), an error number, "Network Unreachable," or "Host Unreachable." If an error number appears here, contact F5 for interpretation.

Control Errors is the number of errors from RON control packets, such as RON heartbeats and LSAs.

Data Errors is the number of failed data packets. These are packets unrelated to RON control, such as replicated files.

Samples

`prtlndA> show ron tunnel`

displays a summary of all configured RON tunnels. See Figure 14.5 for sample output.

`prtlndA> show ron tunnel toBoston`

displays details for a RON tunnel named, "toBoston." See Figure 14.6 on page 14-21.

`prtlndA> show ron tunnel all`

displays details for all RON tunnels.

`stoweA> show ron tunnel redundancy`

displays details for the redundancy tunnel on an ARX-2500, "stoweA." See Figure 14.7 on page 14-22.

Related Commands

ron tunnel
rconsole
show ron
show ron database

Figure 14.5 *Sample Output: show ron tunnel*

```
prtlndA> show ron tunnel

Name            State       Interface       Remote Addr     Up Time
-------------------------------------------------------------------------
toBoston        Connected   VLAN/74         192.168.25.5    0d, 00:16:17
toGoffstown     Connected   VLAN/74         192.168.158.147 0d, 00:16:17
toMinturn       Connected   VLAN/74         192.168.81.22   0d, 00:16:27
```

Figure 14.6 *Sample Output: show ron tunnel toBoston*

```
prtlndA(cfg)# show ron tunnel toBoston

                Name: toBoston
                Peer: bstnA
        Tunnel State: Connected
              Uptime: 0d, 00:16:17
           Interface: VLAN/74
      Remote Address: 192.168.25.5
     Security Policy:
     Ping Fail Limit: 3
       Ping Interval: 10 (seconds)
     Round Trip Time: 0.12 (ms)
    Packet Lost Rate: 0.59 (%)
      TCP Throughput: 125000000 (Bytes/sec)
    Loss-RTT Product: 1 (us)
      RON Packets In: 235
     RON Packets Out: 138
     Data Packets In: 434741
    Data Packets Out: 97751
       Data Bytes In: 546980 (KB)
      Data Bytes Out: 6264 (KB)
     Local Processor: 2.2
     Last Error Code: 0
      Control Errors: 0
```

```
     Data Errors: 0
```

Figure 14.7 *Sample Output: show ron tunnel redundancy*

```
stoweA(cfg)# show ron tunnel redundancy

            Name: Redundancy
            Peer: stoweB
    Tunnel State: Connected
          Uptime: 0d, 01:01:55
       Interface: VLAN/1010
  Remote Address: 10.50.250.9
 Security Policy:
 Ping Fail Limit: 12
   Ping Interval: 1 (seconds)
 Round Trip Time: 0.08 (ms)
Packet Lost Rate: 0.59 (%)
  TCP Throughput: 125000000 (Bytes/sec)
Loss-RTT Product: 0 (us)
  RON Packets In: 9893
 RON Packets Out: 10052
 Data Packets In: 19151
Data Packets Out: 19291
   Data Bytes In: 2762 (KB)
  Data Bytes Out: 2902 (KB)
 Last Error Code: 0
  Control Errors: 0
     Data Errors: 0
```

shutdown (cfg-if-vlan-ron-tnl)

Purpose Use the no form of the command to open the local end of the RON tunnel.
Use this command to shut down a RON tunnel at the local end.

Mode cfg-if-vlan-ron-tnl

Security Role(s) network-engineer or crypto-officer

Syntax
```
no shutdown
shutdown
```

Default(s) shutdown

Guidelines Use the show ron tunnel command to view RON-tunnel status.

Samples
```
bstnA(cfg-if-vlan-ron-tnl[testTunnel])# shutdown
```
shuts down the current RON tunnel, “testTunnel.”

```
bstnA(cfg-if-vlan-ron-tnl[toLA])# no shutdown
```
activates the current RON tunnel.

Related Commands
ron tunnel
show ron tunnel

15

Redundant Pairs (HA)

clear counters redundancy

Purpose Use this command to clear the counters in the various show redundancy ... commands.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `clear counters redundancy [heartbeat | transition | network | critical-services]`

heartbeat (optional) clears only the heartbeat-related counters. This applies to heartbeats over the redundant-pair link (from show redundancy peer) as well as quorum-disk heartbeats (from show redundancy quorum-disk).

transition (optional) clears the counters associated with state/status transitions. This applies to all transition counters in all show redundancy ... commands.

network (optional) clears only the network counters, from the show redundancy network output. See clear counters redundancy network for details.

critical-services (optional) clears the counters associated with critical-services (from show redundancy critical-services).

Default(s) If you omit all of the optional flags, this clears all redundancy counters from all show redundancy ... commands.

Platforms any *except* ARX-VE

Guidelines These counters appear in the output of four show commands:

- show redundancy peer,
- show redundancy quorum-disk,
- show redundancy network, and
- show redundancy critical-services.

The first two of these commands have heartbeat counters that track the number of exchanged heartbeats. All of them have transition counters that track the number of times that the connection state has changed.

Samples

```
bstnA# clear counters redundancy
bstnA#
```

clears all counters.

```
bstnA# clear counters redundancy heartbeat
bstnA#
```

clears only the heartbeat-related counters. This does *not* clear any transition-related counters.

```
bstnA# clear counters redundancy transition
bstnA#
```

clears all transition-related counters from all show redundancy ... output.

Related Commands

show redundancy peer
show redundancy quorum-disk
show redundancy network
show redundancy critical-services
clear counters redundancy network

critical route

Purpose

Use this command to designate an external subnet as "critical," so that a failover may occur if there is no route to the subnet.

Use the no form of the command to make a subnet non-critical.

Mode

cfg-redundancy

Security Role(s)

network-engineer or crypto-officer

Syntax

```
critical route subnet mask
no critical route subnet mask
```

subnet (0.0.0.0-255.255.255.255) is the IP address of the subnet. This must be reachable through at least one static route; use show ip route to view all static routes.

mask (0.0.0.0-255.255.255.255) is the netmask, which identifies the network part of the subnet address.

Default(s)

No critical routes

Platforms

any *except* ARX-VE

Guidelines

The CLI prompts for confirmation if the route is down, the peer is down, or there is some other issue. Enter **yes** if you want to proceed anyway.

You can re-issue this command multiple times to establish multiple critical routes.

If a critical route fails on the current peer *and* the other peer has no failures, control fails over to the other peer. If the other peer has any failures that would ordinarily cause a failover (such as a major hardware fault), no failover occurs. This prevents unnecessary failovers.

The ARX tests for failure with regular ARP requests. Every 20 seconds, the ARX sends an ARP to the route's gateway. (The gateway is configured with the ip route command.) If the gateway fails to respond, the ARX waits an additional 20 seconds before asking the peer if it is possible to fail over. The ARPs continue indefinitely at 20-second intervals. If the gateway responds before the failover is initiated, the failover does not occur.

The ARX uses one of its in-band (VLAN) management addresses to send those ARP requests. The subnet you identify with this command must have a management address for its VLAN. Use the interface vlan command to create an in-band management address for a the subnet's VLAN.

From gbl-ns-vol-shr mode, you can declare a namespace share as a critical resource with the critical command.

Critical routes (unlike critical shares) are *not* shared between redundant switches. This is because the switches may have different visibility into the network. To duplicate a critical route, issue this command on both peers.

To show all critical routes and shares on the current peer, use the show redundancy critical-services command.

Samples

```
bstnA(cfg-redundancy)# critical route 172.16.54.0 255.255.255.0
```

designates a class-C subnet as a critical route. If the ARX loses all routes to that subnet *and* its redundant peer has no serious issues, a failover occurs.

```
bstnA(cfg-redundancy)# no critical route 172.16.0.0 255.255.0.0
```

removes a class-B subnet from the list of critical subnets.

Related Commands

redundancy
show redundancy critical-services
critical
interface vlan

enable (cfg-redundancy)

Purpose Use this command to enable redundant pairing on the current switch.

Mode cfg-redundancy

Security Role(s) network-engineer or crypto-officer

Syntax

```
enable
no enable
```

Default(s) None

Platforms any *except* ARX-VE

Guidelines Before you enable redundant pairing through enable (cfg-redundancy), you must complete the following steps:

1. For ARX-2000 or ARX-4000: Prepare interfaces or a channel to connect the switches.

   At both switches, define an interface that can be used in a redundant-pair link (redundancy protocol), or define a larger channel between the two switches (redundancy protocol (cfg-channel)).

   This is unnecessary for the ARX-500, which has a dedicated port for this link.

2. For ARX-500 only: configure a ron tunnel between the switches to carry the redundancy heartbeats.

3. At each switch, identify the redundant peer using the peer command.

4. Define the *quorum disk* by using the quorum-disk command on both switches. Each peer writes its heartbeats to the same external-filer share, called a quorum disk, and reads the heartbeats from its peer. The quorum disk configuration must be identical on both switches.

The enable command then starts the initial rendezvous of the switches. Once this command is invoked on the peer switch, the rendezvous can proceed. The pair is joined after the rendezvous is complete.

Guidelines: Resilvering The enable command invokes the initial-rendezvous process, which then starts the metalog (namespace log) *resilvering* process. Metalog resilvering is duplicating the metalog data from the active peer to the backup peer. The metalog data must be mirrored on both peers to ensure that the namespace software can fail over. You may want to monitor the resilvering process if the peers are separated by a long distance: you can use the show redundancy metalog command for this. If resilvering times out due to excessive latency, you can use the resilver-timeout command to reset the timeout value.

Guidelines: Senior-Switch Election The senior switch is elected based on its global configuration; if switch A has any namespaces or global servers and switch B does not, switch A becomes the senior switch. If neither has any namespaces or global servers, the one with the lower rendezvous address (established with the peer command) is elected senior. If both have namespaces and/or global servers, the rendezvous fails; see below.

Guidelines: Failure Recovery

One of the switches (if not both) should be devoid of any namespaces or global servers before you enable redundancy. If both have examples of these global-configuration objects, the rendezvous fails. Use the show redundancy history command to detect this failure. To recover, use the clear global-config command to remove all global-config from one of the peers and reboot it. Then enable redundancy again.

The rendezvous also fails if the private subnets of the peers match, or if one peer's private subnet matches any private subnet in the other's RON. The show redundancy history command indicates this problem if it occurs. To recover, use no enable to undo the redundancy configuration, use ip private subnet reassign on one of the switches, then try enabling redundancy again.

The switches continue to retry the rendezvous indefinitely. If you decide to abandon the pairing, you can use no enable to stop the switches from retrying. The no enable feature is disabled after the pair successfully forms.

Sample

```
prtlndA(cfg-redundancy)# enable
```

enables redundant pairing on the current switch.

Related Commands

redundancy protocol
redundancy protocol (cfg-channel)
redundancy
peer
quorum-disk
show redundancy
show redundancy history

nsm binary-core-files

Purpose By default, an NSM processor creates a detailed binary-core file if it has a failure. On the advice of F5 Support, you can use no nsm binary-core-files to return to the smaller ASCII form of NSM-core files.

The nsm binary-core-files command re-instates detailed core files from NSM processors. The core file is useful for diagnosing processor failures.

Mode cfg

Security Role(s) crypto-officer

Syntax
```
nsm binary-core-files
no nsm binary-core-files
```

Default(s)
```
nsm binary-core-files
```

Platforms ARX-500, ARX-2000, and ARX-4000

Guidelines Only use this command on the advice of F5 Support.

An NSM processor requires very-little time to generate the smaller ASCII form of NSM-core files, but they offer less information for later diagnosis. An NSM processor requires up to three minutes to generate the detailed binary-core file, but it fails over to its peer processor at the beginning of the process. If two peer processors fail at the same time, the ARX fails over to its redundant peer before they write their core files. The extra time to write the binary-core file(s) has no effect on clients.

Important

A standalone ARX, without a redundant peer, loses service for up to three minutes if two peer processors fail at the same time. Without an ARX peer, clients must wait for the peer processors to write their binary-core files. For this reason and others, we recommend redundant ARXes at all live customer deployments.

An ARX-500 requires nsm recovery before you can enable this feature.

Use the show nsm command to see the current setting for binary-core files.

Sample
```
bstnA(cfg)# nsm binary-core-files
```
sets the current chassis to produce detailed, binary-core files in the event of any NSM-processor failure.

Related Commands nsm recovery
show nsm

nsm recovery

Purpose

Each NSM-processor core has a redundant peer that takes over in case of a failure. A *recovered* processor goes into a "Standby" state while its peer processor manages all network traffic for both. If the peer processor fails, all network traffic fails back to the peer that failed first. On the advice of F5 Support, you can use no nsm recovery to prevent NSM-processor recoveries on this ARX.

The nsm recovery command re-enables this ARX's NSM processors to recover from failures.

Mode

cfg

Security Role(s)

crypto-officer

Syntax

```
nsm recovery
no nsm recovery
```

Default(s)

```
nsm recovery
```

Platforms

ARX-500, ARX-2000, and ARX-4000

Guidelines

Only use this command on the advice of F5 Support.

The show processors command shows all NSM processors that are in the standby state.

This command cannot run on an ARX-2000 chassis. An ARX-2000 contains four cores on a single chip, so it cannot benefit from the nsm recovery command. If a core fails on the ARX-2000, it remains in the "Failed" state while its peer core processes all of its traffic.

On an ARX-500, this must be enabled before you can enable nsm binary-core-files. The nsm binary-core files command enhances the diagnostic files produced by an NSM processor when it fails. It is enabled by default.

You can use the show nsm command to see the current setting for NSM-processor recovery and/or binary-core files.

Guidelines: Enabling Warm Restarts

You also have the option to make individual cores restart independently, without necessarily putting the entire NSM processor into the "Standby" state. Use the nsm warm-restart command to enable these restarts.

Samples

```
prtlndA(cfg)# no nsm recovery
```

stops NSM-processor recovery on the current chassis. If a processor fails, it fails over to its peer once; if the peer fails later, the ARX reboots.

```
bstnA(cfg)# nsm recovery
```

sets up the NSM-processors to recover after a failure and assume a hot-standby state.

Related Commands

nsm binary-core-files
nsm warm-restart
show nsm

nsm warm-restart

Purpose

An NSM-processor core can attempt to restart after a failure, without causing its entire NSM processor to reboot and go into a "Standby" state. Other cores on the same NSM processor are unaffected by this warm-restart of the failed core. Use the nsm warm-restart command to make NSM cores attempt these warm restarts.

The no nsm warm-restart command returns to the default: if any core fails on an NSM processor, the entire processor reboots.

Mode

cfg

Security Role(s)

crypto-officer

Syntax

```
nsm warm-restart
no nsm warm-restart
```

Default(s)

```
nsm warm-restart
```

Platforms

ARX-2000 and ARX-4000

Guidelines

This command localizes the effects of an NSM-core failure. Only the core that failed restarts, without affecting any *sister cores* on the same processor.

A warm restart produces a core-memory file, visible with the show cores command. The core-memory files produced by this failure are smaller than those produced by a full NSM restart, but they contain enough information for F5 Engineering to analyze the failure. The size of the core-memory file is unaffected by the nsm binary-core-files command.

There are two internal counters, *restart count* and *restart limit*, that prevent repetitive warm restarts. The restart limit is 3, and the restart count applies to an entire NSM processor (or CPU). If core 1 fails on a four-core CPU, and then core 3 fails on the same CPU, the NSM CPU has a restart count of 2. If the restart count reaches the restart limit (3), the entire NSM CPU reboots and behaves as dictated by the nsm recovery setting.

Each restart event times out after 24 hours, decrementing the restart count by one. For example,

- At noon on Thursday, core 1 restarts - restart count = 1.
- At 5PM on Thursday, core 3 restarts - restart count = 2.
- At noon on Friday, core 1's restart times out - restart count = 1.
- At 5PM on Friday, core 3's restart times out - restart count = 0.

You can use the show nsm command to see the current settings for NSM warm restarts, NSM-processor recovery, NSM binary-core files. To see the history of warm restarts on this system, use the show nsm warm-restart history command.

Hardware faults are too severe to allow for a warm restart. For hardware faults, the NSM processor fully reboots and restarts as set by the nsm recovery command.

Sample

```
bstnA(cfg)# nsm warm-restart
```

sets up the NSM-processor cores to independently restart on failure, without affecting any other cores on the same processor.

Related Commands

nsm recovery
nsm binary-core-files
show nsm
show nsm warm-restart history

peer

Purpose

At each redundant peer, use this command to identify the other switch.

Use the no form of the command to remove the redundant-peer configuration.

Mode

cfg-redundancy

Security Role(s)

network-engineer or crypto-officer

Syntax

peer *peer-address*
no peer

peer-address (0.0.0.0-255.255.255.255) is one of the peer switch's management-IP addresses. This can be the out-of-band (MGMT) interface or one of the switch's inband (VLAN) management interfaces. For the ARX-1500 and ARX-2500, you must use the inband management interface defined for the redundancy link: define a special VLAN for the link (with vlan and members (cfg-vlan) for a standalone link, or with vlan (cfg-channel) for a channel), and use the interface vlan and ip address (cfg-if-vlan) commands to create the inband-management IP.

Default(s)

Port default is 49800.

Platforms

any *except* ARX-VE

Guidelines

A *rendezvous* occurs after you issue enable (cfg-redundancy) at the second switch. Each switch uses its ***peer-address*** to contact the other switch and exchange information for the join operation. It also uses this address for regular heartbeat exchanges.

Define the redundant peer *before* issuing enable (cfg-redundancy) on the switches.

Use the show redundancy peer command to see the current configuration for the peer, as well as heartbeat counters.

Sample

```
bstnA(cfg-redundancy)# peer 10.1.1.12
```

identifies the switch at 10.1.1.12 as the current switch's redundant peer.

Related Commands

redundancy
enable (cfg-redundancy)
show redundancy peer

quorum-disk

Purpose Use this command to set up a quorum disk for a redundant pair.

Mode cfg-redundancy

Security Role(s) network-engineer or crypto-officer

Syntax: NFS `quorum-disk nfs-server:/export[path] {nfs2 | nfs3 | nfs3tcp}`

***nfs-server*:/*export*[*path*]** (1-1024 characters) selects an NFS export:

nfs-server is the IP address for the filer (for example, 192.168.70.65). This address must be on a server (proxy-IP) subnet (see ip proxy-address) or reachable through a gateway on that subnet (via static route: see ip route to create a static route).

export is the path to an NFS export on the server.

path (optional) is the specific directory to use.

nfs2 | nfs3 | nfs3tcp is a required choice; this is the NFS protocol to use for accessing the quorum-disk share.

Syntax: CIFS `quorum-disk \\cifs-server\share[\path] cifs [DOMAIN/]username spn spn`

***cifs-server**share*[*path*]** (1-1024 characters) is the syntax for a CIFS share.

cifs-server is the IP address for the filer (for example, 192.168.23.23). This address must be reachable, as described above for an NFS filer.

share is the specific share to use.

path (optional) is a path within the share.

cifs is a required keyword.

[*DOMAIN*/]*username* (1-1024 characters) is the username that the redundancy software can use to write to the CIFS share. If you use a short *DOMAIN* name, like "medarch," you authenticate with NTLM or NTLMv2. If you use an FQDN for the domain, like "medarch.org," you use Kerberos authentication.

spn *spn* (required for a Windows 2008 cluster, optional for other servers; 1-255 characters) is the Service-Principal Name (SPN) for the back-end server. You require a SPN to connect to a CIFS service on any Windows 2008 cluster.

Default(s) ***path*** defaults to the top of the NFS export or CIFS share.

Platforms any *except* ARX-VE

Guidelines The quorum disk is an NFS or CIFS share on any reliable, external filer. Each peer records its heartbeat messages on the quorum disk, and reads the heartbeat messages from the other. This is another path for reading the peer's heartbeats, in addition to the path through the redundant-pair link.

The quorum disk also records ballot information for senior-switch election (see show redundancy ballots).

Guidelines: Filer/Share Requirements

Use a highly-available, high-performance filer for the quorum disk. Round trip data write times for 1 block of quorum-disk data must be below 1 second. Slow-performing filers may cause the redundancy state to fluctuate and could lead to unnecessary switch down time.

The quorum disk's volume must be capable of storing 1 MB of file data. Use a share that blocks until the data is written to disk; no caching should be enabled for the share.

If the quorum disk is an NFS export, it must be configured (at the filer) for *synchronous* writes. Use the 'sync' option. We also recommend that you specify the no_wdelay option. CIFS shares do not have this configuration issue; they perform synchronous writes on request.

Guidelines: Filer Connectivity

Each peer uses one of its in-band (VLAN) management addresses to communicate with the quorum disk's filer. Use the interface vlan command to create an in-band management address for a particular VLAN. The switch can reach the quorum-disk filer through this interface if

- the filer is on the same VLAN as the management interface, or
- the filer is reachable through a static ip route that goes through a gateway on the same VLAN as the management interface.

Guidelines: Command Usage

If you connect to a CIFS share, the CLI prompts for a password to use for accessing the filer. It prompts a second time to confirm the password. For both prompts, enter the password for the ***username***.

Use the show redundancy quorum-disk command to view the current configuration for the quorum disk, as well as some counters.

Samples

```
bstnA(cfg-redundancy)# quorum-disk 172.16.4.98:/lhome/qdisk1 nfs3
```

configures an NFS quorum disk.

```
provB(cfg)# quorum-disk \\10.10.201.8\qd cifs BOSTONCIFS/juser spn
svcA@BOSTONCIFS
Password: jpasswd
Confirm: jpasswd
```

configures a CIFS quorum disk.

Related Commands

redundancy
show redundancy quorum-disk

redundancy

Purpose Use this command to start configuring redundancy between two switches. If an ARX with a redundant peer experiences a catastrophic failure, all of its services fail over to its peer ARX.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax `redundancy`

Default(s) None.

Platforms any *except* ARX-VE

Guidelines: Before You Begin Before you configure redundancy, you must connect the two peers together with a redundancy link. You can use a single port for this, but we recommend two or more ports in a channel.

The ARX-1500 and ARX-2500 use layer-3 (IP) networking software for exchanging heartbeats and important metalog data between the peers, so they require layer-3 configuration for their redundancy link. Use these commands to set up a layer-3 connection:

- Establish a new VLAN for this link. This requires different CLI commands for a channel than it does for a stand-alone port:
  - For a channel, use the vlan (cfg-channel) command to assign the channel to the VLAN.
  - For a single port, use the vlan command to create a new VLAN, then use members (cfg-vlan) to assign the single port to that VLAN.
- Use the interface vlan to create a management-IP interface on the VLAN; this puts you into cfg-if-vlan mode.
  - From cfg-if-vlan mode, use the ip address (cfg-if-vlan) command to establish an in-band (VLAN) IP address. You later use this VLAN-management IP address to identify this ARX to its peer, as described below.
  - From the same mode, use redundancy (cfg-if-vlan) to designate the interface for exchanging metalog data and heartbeats.
  - From the same mode, use no shutdown (cfg-if-vlan) to enable the management interface.

Other platforms use a layer-2 connection for their redundancy link. After cabling the peers together, you use the redundancy protocol command on the link's interface to designate it for use as this link. If you use multiple links in a channel, as recommended, you use the redundancy protocol (cfg-channel) command instead.

Guidelines

The cfg redundancy command brings you to cfg-redundancy mode, where you configure the parameters for creating a redundant pair. From this mode, you use the peer command to identify the redundant peer; for the ARX-1500 and ARX-2500, you must use the other peer's in-band (VLAN) IP address at the other end of the redundancy link. You also use quorum-disk to identify an external-filer share to be used as a quorum disk. Repeat these steps at the peer switch, which must have redundancy parameters that agree. Once the parameters match on both switches, you enable the redundant pair with enable (cfg-redundancy) at each peer.

The enable command invokes the initial-rendezvous process, which then starts the metalog (namespace log) *resilvering* process. Metalog resilvering is duplicating the metalog data from the active peer to the backup peer. The metalog data must be mirrored on both peers to ensure that the namespace software can fail over. You may want to monitor the resilvering process if the peers are separated by a long distance: you can use the show redundancy metalog command for this. If resilvering times out due to excessive latency, you can use the resilver-timeout command to reset the timeout value.

For rare situations where network maintenance may cause unwanted failovers, you can use the suspend-failover command to suspend failovers for a short time. Use the no form of the command to lift the suspension when the maintenance is finished.

Guidelines: ARX-1500 and ARX-2500 Redundancy

The ARX-1500 and ARX-2500 store their metalog data on their internal disks, along with logs, software-release files, and other management data. Managed volumes write their metalog data as clients change the volume state; the metalog is used to restore the volume configuration in the event of a failure. The metalog is also copied to the redundant peer. The speed of many volume operations depends on fast metalog writes.

Some other system operations create a large number of writes to the internal disk, potentially slowing metalog writes. This can slow volume performance, even if it occurs on the backup peer. For example, the process of upgrading the software release is extremely disk intensive, and may cause a noticeable performance degradation. During an upgrade, you use

- the copy command (copy ftp, copy {nfs|cifs}, copy scp, or copy tftp) to copy a full release file to the disk, and
- the boot system command to unpack the release file on the disk.

You should perform such disk-intensive operations during off-peak hours on the ARX-1500 and ARX-2500. This is true whether you run the operations on the active peer or the backup.

Guidelines: ARX-VE Redundancy

The ARX-VE is a Virtual Appliance (VA, similar to a VM) that runs on a standard hypervisor, and uses the same redundancy mechanisms that are used in a VM cluster. If the ARX-VE's hypervisor host fails, an identical ARX-VE on a peer hypervisor resumes processing. This does not require any additional configuration on the ARX, so the redundancy command, cfg-redundancy mode, and other ARX-redundancy operations are excluded from the ARX-VE CLI.

Sample

```
bstnA(cfg)# redundancy
bstnA(cfg-redundancy)# ...
```

starts redundancy configuration.

Related Commands

enable (cfg-redundancy)
quorum-disk
show redundancy
show redundancy metalog
resilver-timeout
suspend-failover

redundancy force-active

Purpose Use this command to resolve a rare scenario (described below) where one peer has failed unrecoverably and the other will not take over as the active peer.

Mode priv-exec

Security Role(s) network-technician, network-engineer, or crypto-officer

Syntax `redundancy force-active`

Default(s) None.

Platforms any *except* ARX-VE

Guidelines This command forcibly promotes the backup peer to active status.

Important

This command discards any forward progress made at the other peer. It also re-imports all namespaces, forcing all clients to re-mount all front-end services. Never use this command unless under the circumstances below, and with the advice of F5 personnel.

Consider a redundant pair where Peer A is active and Peer B is either active or backup. Peer B fails, then (later) Peer A fails unrecoverably. When Peer B recovers, it refuses to take an active role because it knows that it was down while Peer A was active: Peer A may therefore have changed its configuration or managed some transactions with back-end filers. If Peer B were to take control, all such configuration changes and transactions would be lost. In this scenario, Peer B waits indefinitely for Peer A and the redundant pair never comes back online.

You need to force Peer B to take an active role, thus abandoning any configuration changes or transactions that occurred on Peer A. The peer switch must be offline, and the current switch must have been up and waiting for its peer to come online for at least 5 minutes. Also, the quorum disk must be online and available. Use the show redundancy command to verify that the peer is down and the quorum disk is connected.

The CLI warns you before discarding transactions and rebooting both peers. Enter **yes** to proceed.

Sample

```
bstnA(cfg)# redundancy force-active
CAUTION: To avoid data corruption, the peer switch MUST BE OFFLINE and
remain offline while this command executes. Any global configuration
changes made on the Active peer switch while this switch was
unavailable will be lost. All namespaces will be re-imported
automatically to ensure consistency in the metadata. All in-flight
transactions not synchronized will be lost. NFS-based clients will
need to remount. Please contact Technical support for further advice.

Proceed? [yes/no] yes
bstnA(cfg)# ...
```

forces the “bstnA” switch to take active status in its redundant pair.

Related Commands show redundancy

resilver-timeout

Purpose

The redundant peers exchange their metalog (namespace-transaction) data during initial rendezvous, or after a failover. This is called *resilvering* the metalog data. At some sites, the latency between peers is high enough that the resilvering process times out before the pair can form. On the advice of F5 Support, you can use this command to increase the time allowed for resilvering.

Use the no form of the command to reset the resilvering timeout to its default.

Mode

cfg-redundancy

Security Role(s)

network-engineer or crypto-officer

Syntax

```
resilver-timeout minutes
no resilver-timeout
```

minutes (6-60) is the maximum number of minutes for resilvering. If this time expires before resilvering is complete, the redundant pair cannot form.

Default(s)

6 minutes

Platforms

any *except* ARX-VE

Guidelines

A *rendezvous* occurs after you issue enable (cfg-redundancy) at the second switch. A *failover* occurs whenever the active switch fails and the backup switch takes control. You can use this command to increase the time allotted for resilvering metalog data during a rendezvous or a failover.

The default is sufficient for most sites. Use this command only on the advice of F5 Support. It is unnecessary if you use a direct layer-2 connection for the redundancy link. The packet latency on a direct connection is typically very short.

If resilvering times out, the redundancy software retries until it succeeds. This severely impacts system performance. You can use the show redundancy history command to see the results of resilvering (referenced as "synchronization" in that output), and determine whether or not it is repetitively timing out. You can use show redundancy metalog to monitor the resilvering process as it occurs.

Use the show redundancy resilver-timeout command to see current timeout value for resilvering.

Sample

```
provA(cfg-redundancy)# resilver-timeout 10
```

sets the resilvering timeout to 10 minutes on the "provA" switch. The resilvering process has 10 minutes to complete after rendezvous or failover.

Related Commands

redundancy
show redundancy resilver-timeout
show redundancy metalog

show nsm

Purpose The show nsm command shows the current state of the NSM-maintenance features: processor recovery and binary-core files. For each of these features, this command shows the administrative setting as well as whether or not the setting is operational.

Mode (any)

Security Role(s) crypto-officer

Syntax `show nsm [recovery | binary-core-files | warm-restart]`

recovery | binary-core-files | warm-restart (optional) focuses the output on the state of a single NSM-maintenance feature. If you omit these, the output shows the state of all NSM features.

Platforms ARX-500, ARX-2000, and ARX-4000

Guidelines The output contains a separate table for each NSM-maintenance feature, **Recovery**, **Binary Core Files**, and **Warm-Restart**. Each table has one row per NSM, where each row has the following fields:

Proc identifies an NSM processor, in *slot.processor* format.

Status - Admin is the status that was last set for this feature.

Status - Operational is the feature's actual status.

You can use [no] nsm recovery to disable or re-enable NSM-processor recovery. You can also use [no] nsm binary-core-files to change NSM-processor core files to an ASCII format or a larger binary format. NSM recovery must be enabled before you can enable NSM binary-core files. Only change these settings on the advice of F5 Support. The final table concerns the [no] nsm warm-restart command; this allows a processor core to fail and recover independently, without rebooting the entire processor.

Sample

```
bstnA# show nsm binary-core-files

Binary Core Files:
                  Status
  Proc     Admin        Operational
  ------   ----------   ------------
  2.1      Enabled      Enabled
  2.2      Enabled      Enabled
  2.3      Enabled      Enabled
  2.4      Enabled      Enabled
  2.5      Enabled      Enabled
  2.6      Enabled      Enabled
  2.7      Enabled      Enabled
  2.8      Enabled      Enabled
  2.9      Enabled      Enabled
  2.10     Enabled      Enabled
  2.11     Enabled      Enabled
  2.12     Enabled      Enabled
```

shows that the current chassis is configured to produce detailed, binary core files in the event of an NSM-processor failure.

Related Commands

nsm recovery
nsm binary-core-files
nsm warm-restart

show nsm warm-restart history

Purpose To show the history of warm restarts by the NSM-processor cores, use the show nsm warm-restart history command.

Mode (any)

Security Role(s) crypto-officer

Syntax `show nsm warm-restart history [processor slot.processor]`

processor *slot.processor* (optional) specifies one NSM processor for which to view this history. If you omit this option, the output includes warm-restart history for all NSM processors.

slot (2 for ARX-4000; 1 for any other) is the slot number.

processor is the processor number. Use the show processors command to show all processors and their associated *slot.processor* IDs.

Platforms ARX-500, ARX-2000, and ARX-4000

Guidelines An NSM *warm restart* occurs when an NSM-processor core encounters a catastrophic software failure and reboots, without causing any other cores to reboot with it. A warm restart is only possible on a system where nsm warm-restart is enabled.

The output contains two tables, one that shows the most-recent recent warm restart and another that shows the current values for *restart count* and *restart limit*. The restart count is the sum of all restarts for a given CPU's cores, and the restart limit is the maximum number of restarts for each CPU. If a core fails after its CPU's restart count reaches the restart limit, the entire CPU reboots. The full-CPU reboot follows the rules set by the nsm recovery command.

The first table has one row per NSM processor, with the following fields:

Proc identifies the NSM processor, in *slot.processor* format. This identifies a specific core on the CPU.

CPU identifies the NSM-processor CPU, where each CPU typically contains multiple cores. This field identifies the CPUs with a letter, such as A, B, or C.

Restart Number is the most-recent restart number for this core. The full restart count for the CPU is the sum of all its cores' Restart Numbers. As mentioned above, you use the nsm warm-restart command to make warm restarts possible for NSM cores.

Date/Time (UTC) is time stamp (if any) for the most-recent warm restart.

The second table shows the number of restarts remaining for each hardware CPU from the first table:

Slot is the slot number for the NSM processor.

CPU identifies the NSM-processor CPU, where each CPU typically contains multiple cores. This maps to the CPU field in the table above.

Restart Remaining is the number of warm restarts remaining for this NSM CPU. This count decreases by one every time one of the CPU's cores has a warm restart. It increases by one every time a warm restart ages by 24 hours.

Restart Limit is the total number of restarts allowed for each CPU.

Samples

`bstnA# show nsm warm-restart history`

shows the warm-restart history for all NSM processors in the "bstnA" chassis. See Figure 15.1 for sample output., without any warm restarts.

`prtlndA# show nsm warm-restart history`

shows the warm-restart history for all NSM processors in the "prtlndA" chassis. See 15.2, on page 15-26 for sample output., where two warm restarts occurred in the last 24 hours.

Related Commands

nsm recovery
nsm binary-core-files
nsm warm-restart

Figure 15.1 *Sample Output: show nsm warm-restart history*

```
bstnA# show nsm warm-restart history

Proc          CPU         Restart           Date/Time(UTC)
                          Number
-----------------------------------------------------------
2.1            A             0

2.2            A             0

2.3            A             0

2.4            A             0

2.5            B             0

2.6            B             0

2.7            B             0

2.8            B             0

2.9            C             0

2.10           C             0

2.11           C             0

2.12           C             0


Slot          CPU         Restart
                          Remaining
-----------------------------------
2              A             3
2              B             3
2              C             3

Restart Limit:      3
```

Figure 15.2 *Sample Output: show nsm warm-restart history (with restarts)*

```
prtlndA# show nsm warm-restart history

Proc          CPU         Restart           Date/Time(UTC)
                          Number
-----------------------------------------------------------
1.2            A             1               06/22/2011 02:51:36 -0400

1.3            A             1               06/22/2011 02:52:00 -0400

1.4            B             0

1.5            B             0


Slot          CPU         Restart
                          Remaining
-----------------------------------
1              A             1
1              B             3

Restart Limit:      3
```

show redundancy

Purpose Use this command to show high-level status for the redundant pair.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show redundancy`

Default(s) None

Platforms any *except* ARX-VE

Guidelines Node is 1 for the initial-senior switch, 2 for the initial-junior switch, or QD for the quorum disk. The asterisk (*) indicates the local node.

Switch/Quorum Disk identifies each node with a hostname or IP address.

Status is Up, "Up,NoFovr," Down, Suspended, or Unknown for a redundant peer. The "UpNoFovr" status indicates that someone used suspend-failover to temporarily freeze the Active/Backup status of the peers and suspend failovers.
The status is Up, "Up,NoHb," Pending, or Down for the quorum disk. "Up,NoHb" means that the quorum disk is up but not showing any heartbeats from the peer yet.

Role is Active, Backup, or Quorum. The senior peer is always Active, meaning it can run namespace software and virtual servers. The junior peer is "Backup;" it is a hot standby for the active peer.

Transitions: Total is the number of times that the Status has changed. This increments for each failover. If there has never been a failover, the Total is "Never."

Transitions: Last (UTC) is the timestamp for the last status change, in Universal Coordinated Time (UTC).

Sample See the sample output in Figure 15.3.

Related Commands redundancy

Figure 15.3 *Sample Output: show redundancy*

```
prtlndA> show redundancy

                                          Transitions
Node Switch/Quorum Disk   Status     Role    Total Last (UTC)
---- -------------------- ---------- ------- ----- -------------------
*1   prtlndA              Up         Active  Never -
 2   prtlndB              Up         Backup  1     05:33:19 09/14/2009
 QD  192.168.74.83        Up         Quorum  1     05:33:07 09/14/2009
```

show redundancy all

Purpose Use this command to show all redundant-pair information with a single command.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show redundancy all`

Default(s) None

Platforms any *except* ARX-VE

Guidelines This command shows all flavors of show redundancy ... output, in the following order:

1. show redundancy
2. show redundancy history
3. show redundancy peer
4. show redundancy ballots
5. show redundancy network
6. show redundancy critical-services
7. show redundancy resilver-timeout
8. show redundancy quorum-disk

For details about the show redundancy all output, refer to the command descriptions for these individual commands.

Sample See the sample output in Figure 15.4.

Related Commands redundancy

Figure 15.4 *Sample Output: show redundancy all*

```
prtlndA> show redundancy all

                          Transitions
Node Switch/Quorum Disk   Status     Role    Total Last (UTC)
---- -------------------- ---------- ------- ----- -------------------
*1   prtlndA              Up         Active  Never -
 2   prtlndB              Up         Backup  2     06:21:10 08/25/2010
 QD  192.168.74.83        Up         Quorum  1     06:18:55 08/25/2010

Date/Time(UTC) Recent History
--------------------------------------------------------------------------
08-25 06:21:15 Quorum disk is online, system is ready for failover
08-25 06:21:10 Peer switch 'prtlndB' is now online
08-25 06:19:20 Local switch wins seniority election (sr=1, epoch=3)
08-25 06:19:10 Quorum disk is online, system is ready for failover
08-25 06:19:05 Peer switch 'prtlndB' is now offline
08-25 06:19:05 Peer is down, local switch will take over services (sr=1, epoch=3)
```

```
08-25 06:18:55 Quorum disk 192.168.74.83:/exports/quorum-disk/portland1 is now online.
08-25 06:18:45 Neither peer nor quorum disk is reachable, networking may not yet be online.
08-25 06:18:24 Quorum disk 192.168.74.83:/exports/quorum-disk/portland1 not responding after 4
seconds (possibly offline).
08-25 06:16:18 Site quorum manager daemon started.


Peer
  Name:           prtlndB
  IP Address:     10.1.23.12
  Port:           49800

Status:           Backup

Heartbeats
  Sent:           838
  Received:       792

Transitions
  Count:          2
  Last:           06:21:10 08/25/2010
  Reason:         Peer switch 'prtlndB' is now online

                      Ballot Cast
Node Switch/Quorum Disk   Senior Switch        Epoch
---- -------------------- -------------------- -----
*1   prtlndA               prtlndA              3
 2   prtlndB               prtlndA              2
 QD  192.168.74.83         prtlndA              3

Last vote occurred at:  06:19:20 08/25/2010


Network  VLAN  Port(s)   Admin    Link   Spanning-Tree
                         State    Status Status
-------- ----- --------- -------- ------ -----------------
External 1     1/1       Enabled  Down   Disabled
External 1     1/2       Enabled  Down   Disabled
External 1     1/3       Enabled  Down   Disabled
External 1     1/4       Enabled  Down   Disabled
External 1     1/5       Enabled  Up     Manual Forwarding
External 1     1/6       Enabled  Up     Manual Forwarding
External 1     1/7       Disabled Down   Disabled
External 1     1/8       Disabled Down   Disabled
External 1     1/9       Disabled Down   Disabled
External 1     1/10      Disabled Down   Disabled
External 1     1/11      Disabled Down   Disabled
External 74    1/5       Enabled  Up     Manual Forwarding
External 74    1/6       Enabled  Up     Manual Forwarding
Private  1008  1/12      Enabled  Up     Manual Forwarding
Metalog  1009  1/12      Enabled  Up     Manual Forwarding


Link Transitions:
  Count:          3
  Last:           06:18:15 08/25/2010
  Reason:         Port 1/12 link up
  Last Cleared:   Never

                                                        Transitions
Type        Service                          Status  Count Last (UTC)
--------------------------------------------------------------------------------
```

```
meta-only   192.168.74.89:/vol/vol1/mdata_A     Up       1  2010-08-25 06:23:02
meta-only   192.168.74.89:/vol/vol1/mdata_B     Up       1  2010-08-25 06:23:12
quorum      192.168.74.83:/exports/quorum-disk/portland1 Up       1  2010-08-25 06:18:55
route       0.0.0.0/0                           Up       0  Never
share       nemed~/acctShdw~back2               Up       0  Never

Counters last cleared: Never


Resilver Timeout:  6 minutes

Path:             192.168.74.83:/exports/quorum-disk/portland1
Protocol:         nfs2
Status:           Up

Heartbeats
  Sent:           1054
  Received:       1006

Transitions
  Count:          1
  Last:           06:18:55 08/25/2010
  Reason:         Quorum disk 192.168.74.83:/exports/quorum-disk/portland1 is now online.

Heartbeat Latency:
                        Heartbeat Latency Intervals (msec)
  Time Interval    [0-499]   [500-999]   [1000-3999]   [No Response]
  -------------    ------------------------------------------------
  02:00 - 02:36      1056          0             0              0
  01:00 - 02:00         0          0             0              0
  00:00 - 01:00         0          0             0              0
  23:00 - 24:00         0          0             0              0
  22:00 - 23:00         0          0             0              0
  21:00 - 22:00         0          0             0              0
  20:00 - 21:00         0          0             0              0
  19:00 - 20:00         0          0             0              0
  18:00 - 19:00         0          0             0              0
  17:00 - 18:00         0          0             0              0
  16:00 - 17:00         0          0             0              0
  15:00 - 16:00         0          0             0              0
  14:00 - 15:00         0          0             0              0
  13:00 - 14:00         0          0             0              0
  12:00 - 13:00         0          0             0              0
  11:00 - 12:00         0          0             0              0
  10:00 - 11:00         0          0             0              0
  09:00 - 10:00         0          0             0              0
  08:00 - 09:00         0          0             0              0
  07:00 - 08:00         0          0             0              0
  06:00 - 07:00         0          0             0              0
  05:00 - 06:00         0          0             0              0
  04:00 - 05:00         0          0             0              0
  03:00 - 04:00         0          0             0              0

Heartbeat latency summary:
  0-499msec  : 100.00%
  500-999 msec :   0.00%
  1000-3999 msec :   0.00%
  No response  :   0.00%
```

show redundancy ballots

Purpose Each failover is followed by an election process to choose the senior switch. Use this command to show the ballots from the most-recent redundancy election.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show redundancy ballots`

Default(s) None

Platforms any *except* ARX-VE

Guidelines Node is 1 for the initial-senior switch, 2 for the initial-junior switch, or QD for the quorum disk. The asterisk (*) indicates the local node.

Switch/Quorum Disk identifies each peer with a hostname or IP address.

Ballot Cast shows the most-recent seniority vote from each node:

Senior Switch is the switch that the node believes should be senior, and

Epoch is the epoch number that the node last recorded. All three quorum members keep a common epoch number that increments with each failover. If one node has a lower epoch number than the others, its vote for Senior Switch is discounted.

Guidelines: Redundancy Elections Each switch, as well as the quorum disk, stores an epoch number and the identity of the senior switch. After a switch failure, this information is exchanged as election *ballots* to determine which peer should have seniority. Ballots with higher epoch numbers carry more weight in the election. If two nodes agree but disagree with the third, the majority rules. The results determine which peer is senior.

None of the ballots are authoritative; use the show redundancy command to see which peer(s) is/are currently active. (The senior switch is always active, the junior switch is active after the initial rendezvous but backup after a failover.)

For example, the following table shows that the Quorum Disk and the second peer, “prtlndB,” agree on the Epoch number (9) and the senior peer (node 2, which is prtlndB). The “prtlndA” peer cast the dissenting vote (1, or itself, as senior) with a lower Epoch number (8). From this output, we can infer that “prtlndA” was senior before it failed, then became junior when it rejoined the pair.

```
...
                                  Ballot Cast
Node Switch/Quorum Disk   Senior Switch        Epoch
---- -------------------- -------------------- -----
*1   prtlndA              1                    8
 2   prtlndB              2                    9
 QD  192.168.74.83        2                    9
...
```

Ballots are only cast when a node fails, so all ballots are 0 (zero) when the pair first forms. In this case, both nodes are always active and Node 1 is always senior.

Sample See the sample output in Figure 15.5.

Related Commands redundancy
show redundancy

Figure 15.5 *Sample Output: show redundancy ballots*

```
prtlndA> show redundancy ballots

                      Ballot Cast
Node Switch/Quorum Disk   Senior Switch        Epoch
---- -------------------- -------------------- -----
*1   prtlndA              prtlndA              3
 2   prtlndB              prtlndA              2
 QD  192.168.74.83        prtlndA              3

Last vote occurred at:  06:19:20 08/25/2010
```

show redundancy critical-services

Purpose If the current switch loses contact with a critical service (such as a critical namespace share or a critical subnet), a failover may occur. Use this command to show all of the critical services configured for the current peer.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show redundancy critical-services`

Default(s) None

Platforms any *except* ARX-VE

Guidelines Each critical service has one row in the table:

- Type is "quorum" "share," "route," or "meta-only." The quorum disk is a critical resource by default, and you cannot remove it from this list. To declare a namespace share as "critical," use critical in gbl-ns-vol-shr mode. To create a critical route, use critical route in cfg-redundancy mode. A "meta-only" resource is a dedicated metadata share that is critical: use metadata share to configure a dedicated metadata share, and use metadata critical to make it a critical resource.
- Service specifies the exact route or share. For a quorum disk, this shows the machine and path to the share. For shares, this shows the share in *namespace~volume~share-name* format. For a critical route, this shows the critical subnet in *ip-address*/*subnet-length* (CIDR) format.
- Status is "Up," "Down," or "Config." A "Config" status indicates that the critical service is configured but the Up or Down status has not been determined yet.
- Transitions: Count is the number of times that the Status has changed.
- Transitions: Last (UTC) is the timestamp for the last status change, in Universal Coordinated Time (UTC). If there has never been a failover, this shows "Never."

Counters Last Cleared is the timestamp for the last time someone ran clear counters redundancy [critical-services].

Guidelines: Failover Algorithm Only an active switch can initiate a critical-services failover. By default, both switches are active. After a failover, only the senior switch is active.

If an active switch loses a critical share or route, it requests a failover. If the peer has all critical routes, shares, and access to the quorum disk, the failover occurs. Conversely, the peer rejects the failover if it has lost any of its own critical resources: a failover in this case could cause more service outages than it resolves.

Note the difference in the way that the quorum disk is treated. Peer A does not initiate a failover if it loses access to the quorum disk. If peer A loses some other critical service (such as a critical share), it does initiate a failover. In this case, peer B *rejects* the failover if peer B cannot access the quorum disk. Loss of the quorum disk can prevent a failover, but never causes a failover.

Sample See the sample output in Figure 15.6.

Related Commands

- critical
- critical route
- metadata critical
- clear counters redundancy

Figure 15.6 *Sample Output: show redundancy critical-services*

```
prtlndA> show redundancy critical-services

                                                          Transitions
Type        Service                            Status  Count Last (UTC)
-------------------------------------------------------------------------------
meta-only   192.168.74.89:/vol/vol1/mdata_A     Up       1  2010-03-02 07:41:21
meta-only   192.168.74.89:/vol/vol1/mdata_B     Up       1  2010-03-02 07:41:28
quorum      192.168.74.83:/exports/quorum-disk/portland1 Up       1  2010-03-02 07:39:42
route       0.0.0.0/0                          Up       0  Never
share       nemed~/acctShdw~back2              Up       0  Never

Counters last cleared: Never
```

show redundancy history

Purpose Use this command to show some high-level redundancy logs kept at the local switch.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show redundancy history`

Default(s) None

Platforms any *except* ARX-VE

Guidelines Recent History messages include:

- Remote peer has gone offline
- Pair status changed from
  - Inactive to Rendezvous,
  - Rendezvous to Joining, or
  - Joining to Formed.
- Incumbent switch wins seniority election (sr=1, epoch=4)
- Site quorum manager daemon started

Errors with forming the pair appear here when they occur. Many of them are self-explanatory, but some of them (such as "DbSync" errors) may require intervention from F5 Support.

Sample See the sample output in Figure 15.7.

Related Commands redundancy

Figure 15.7 *Sample Output: show redundancy history*

```
prtlndB# show redundancy history

Date/Time(UTC) Recent History
--------------------------------------------------------------------------
09-14 05:33:24 Quorum disk is online, system is ready for failover
09-14 05:33:19 Synchronization complete, ready for failover
09-14 05:33:19 Pair status changed from Joining to Formed
09-14 05:33:19 Peer switch 'prtlndA' is now online
09-14 05:33:07 Quorum disk 192.168.74.83:/exports/quorum-disk/portland1/ is online, proceeding to
verify rendezvous.
09-14 05:33:07 Quorum disk 192.168.74.83:/exports/quorum-disk/portland1 is now online.
09-14 05:33:01 Pair status changed from Rendezvous to Joining
09-14 05:33:01 Switch 'prtlndA' has a lower rendezvous IP address than 'prtlndB'. [10.1.23.11,
10.1.23.12]
09-14 05:32:56 Cannot rendezvous because quorum disk 192.168.74.83:/exports/quorum-disk/portland1
is offline.
09-14 05:32:56 Pair status changed from Inactive to Rendezvous
09-14 04:21:50 Site quorum manager daemon started.
```

show redundancy license

Purpose Both peers in a redundant pair require the same license, to prevent any loss of service after a failover. A failover to a peer with a lesser license would result in fewer supported features and fewer available system resources. To confirm that the licenses are the same between the current peers, use the show redundancy license command.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show redundancy license`**

Default(s) None

Platforms any *except* ARX-VE

Guidelines The output contains a success message if the licenses are the same for both peers. Otherwise, it returns a failure message along with all mis-matched features.

In the case of a failure, you need to determine which peer has the desired license, and then activate or re-activate the license at the other peer. Run show active-license at each peer to see details about the licenses there. If the peer can connect to the Internet, you can use the license activate command to automatically activate the license there. Otherwise, you can use a manual activation method, as described in the documentation for the license create license-dossier command.

Sample

```
prtlndA> show redundancy license
% INFO: The HA peer license is equivalent to the license on this ARX.
```

shows that the licenses match between the peers.

Related Commands
show active-license
license activate
license create license-dossier
redundancy
show redundancy

show redundancy metalog

Purpose Metalog (namespace transaction) data is mirrored between redundant peers during their initial rendezvous, and it is duplicated between the peers during normal operation. The duplication process is called *resilvering*. You can use the show redundancy metalog command to see the current state of the metalog-resilvering process between peers. This command is especially useful for monitoring the connection between redundant peers that are separated by long distances.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show redundancy metalog`

Default(s) None

Platforms any *except* ARX-VE

Guidelines The active peer's metalog data must be mirrored at the backup peer so that namespace software can recover after a failover. Use this command to monitor the resilvering process.

This command only functions on the active peer.

Use this output to gauge the latency and performance of the redundant-pair link. An overly-long latency between peers can cause a large number of metalog-packet retransmits (see below), and may lead to slow responses that are noticeable to namespace/volume/front-end service clients.

On all platforms, the namespace software keeps metalog read/write statistics. You can use the show statistics metalog command to see these metalog-usage statistics from a namespace-software perspective.

The ARX-1500 and ARX-2500 store their metalog data on their internal disks. On those platforms, you can use the show metalog usage command to see the metalog-usage statistics.

The ARX-VE stores its metalog data on an external data store. On the ARX-VE, you can use the probe metalog latency command to test the latency between the ARX and this external storage.

Guidelines: Output

The output shows the current state of the resilvering process:

State is

- Resilvering - the active peer is copying all of its metalog data to the backup peer during a rendezvous.
- Peer Online - the metalog data was 100% duplicated after rendezvous, and now the active peer is sending metalog updates as they occur.
- Peer Offline

The next three fields only appear for the Resilvering state, which occurs during rendezvous:

Started is the start time for the rendezvous process.

Timeout indicates the total time allowed for resilvering. The resilvering process times out, sends an SNMP trap, and restarts if it exceeds this time limit. You can use the resilver-timeout command to reset this time limit.

Time Remaining is the estimated time left for resilvering. This estimate is based on the current data-transfer rate and the amount of data left to transfer.

Byte Count is the number of bytes of metalog data transferred to the backup peer. If the initial Resilvering process is still underway, this also shows the total size of the metalog data to be copied to the backup peer.

Retransmits counts the retransmissions of individual metalog packets. A retransmit occurs if an internal timeout passes before the packet is acknowledged by the backup peer. This field counts the total retransmits that occurred since the most recent start of the resilvering operation. A long latency between peers may increase the number of retransmits. This always displays 0 (zero) for the ARX-1500 or ARX-2500 when they are resilvering. These platforms use a different transmission mechanism for their metalog packet.

Latency shows the minimum, maximum, and average latency for sending packets of metalog data to the backup peer. These are measured in micro-seconds (us).

Data Rate shows the average megabits per second for the transfer of metalog data. This field only appears while resilvering is occurring.

Samples

```
nyc15> show redundancy metalog
State:              Resilvering

Started:            01/31/2012 22:48:03 +0000
Timeout:            00:06:00 sec
Time Remaining:     00:00:58 sec

Byte Count:    992 M of 3.7 G (26%)

Retransmits:   0

Latency:
      Min:     18538 usec
      Max:     77860 usec
      Avg:     22715 usec


Data Rate:     47.2 Mb/s
```

shows the state of metalog resilvering on the “nyc15” chassis.

```
prtlndA> show redundancy metalog
State:         Peer Online

Byte Count:    55 M
Retransmits:   0

Latency:
      Min:     80 usec
      Max:     127 usec
      Avg:     83 usec
```

shows the state of metalog resilvering on the “prtlndA” chassis. In this case, the initial resilvering is complete and the active peer is sending metalog updates as it generates them.

Related Commands

resilver-timeout
redundancy
show redundancy
show metalog usage
show statistics metalog
probe metalog latency

show redundancy peer

Purpose Use this command to show the configuration and counters associated with the current switch's redundant peer.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show redundancy peer`

Default(s) None

Platforms any *except* ARX-VE

Guidelines Peer is the heading for peer-identification parameters:

- Name is set by the hostname command at the peer's CLI.
- IP Address is the peer's management IP that is chosen for the rendezvous. You can change this with the peer command.
- Port is the *peer* port used for rendezvous. You can use an option in the peer command to change the port number.

Status is Active, Backup, Down, or Unknown..

Heartbeats are counters for the number of redundancy heartbeats sent and received. You can use the clear counters redundancy command to clear this counter.

Transitions are the number of changes in redundancy Status for this peer. You can use the clear counters redundancy command to clear this counter, too.

Sample See the sample output in Figure 15.8.

Related Commands redundancy
peer
clear counters redundancy

Figure 15.8 *Sample Output: show redundancy peer*

```
prtlndA> show redundancy peer

Peer
  Name:           prtlndB
  IP Address:     10.1.23.11
  Port:           49800

Status:           Backup

Heartbeats
  Sent:           573
  Received:       572

Transitions
  Count:          1
  Last:           05:33:19 09/14/2009
  Reason:         Peer switch 'prtlndB' is now online
```

show redundancy quorum-disk

Purpose A quorum disk is a filer share that each switch uses to write its own heartbeats and read the heartbeats from its peer. Use this command to show the configuration and counters associated with the quorum disk.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show redundancy quorum-disk`

Default(s) None

Platforms any *except* ARX-VE

Guidelines Path is in *machine*:/*path* format (for an NFS export) or *machine**path* format (for a CIFS share). This is the external-filer share used as the quorum disk. From cfg-redundancy mode, use quorum-disk to reset this.

Protocol is the file-access protocol (nfs2, nfs3, nfs3tcp, or cifs) used to access the quorum-disk share. This flag is also set with the quorum-disk command.

If the protocol is cifs, the following fields appear to describe the CIFS options used to access the quorum disk. These are all options from the quorum-disk command:

- QD User is the Windows username that the ARX uses as its identity when accessing the quorum disk.
- User Domain is the Windows domain for the above username.
- QD SPN is the Service Principle Name (SPN) for the quorum disk's host server.

Status is Up, Pending, or Down. Pending indicates that the quorum disk is functional but the redundant pair is in the process of forming.

Heartbeats are counters for the number of redundancy heartbeats sent to the quorum disk by the current node, and received from the quorum disk by the current node. You can use the clear counters redundancy command to clear this counter.

Transitions shows the changes in Status for the quorum disk:

- Count is the number of changes. This should be a low number; the only valid reason for a quorum-disk transition is a planned outage of the quorum-disk filer (such as a hardware upgrade).
- Last is the time stamp for the most-recent transition.
- Reason is the log message associated with the Status change.

You can use the clear counters redundancy command to clear the transitions counter, too.

Guidelines (Cont.) Heartbeat Latency is a chart of latency measures over the past 24 hours. This is the latency (round-trip time) for heartbeat packets between the current peer and the quorum disk. Each row shows the latency measures in one hour:

Time Interval shows the start and end time for the hour, in local time.

[0-499] is a count of heartbeats with a latency of 0-499 milliseconds. This column should have by far the highest count of all of them.

[500-999] is the number of heartbeats that took 500-999 milliseconds to make a round trip. This is a long latency and should be uncommon.

[1-3999] shows how many heartbeats took 1 second to 3,999 milliseconds (almost 4 seconds). This latency should be extremely rare, if it ever occurs.

[No Response] is the number of heartbeats that were lost. This number should be 0, unless there is a planned outage for the quorum disk. If either peer reboots while one of them is disconnected from the quorum disk, they both may reboot simultaneously.

Heartbeat Latency Summary shows the percentage of heartbeats in each of the above time intervals.

If the Heartbeat Latency and Heartbeat Latency Summary tables indicate long latencies, choose a faster quorum disk. You can run the quorum-disk command on both peers to change the quorum disk.

Sample See the sample output in Figure 15.9.

Related Commands redundancy
quorum-disk
clear counters redundancy

Figure 15.9 *Sample Output: show redundancy quorum-disk*

```
prtlndA> show redundancy quorum-disk

Path:             192.168.74.83:/exports/quorum-disk/portland1
Protocol:         nfs2
Status:           Up

Heartbeats
  Sent:           886
  Received:       884

Transitions
  Count:          1
  Last:           07:39:42 03/02/2010
  Reason:         Quorum disk 192.168.74.83:/exports/quorum-disk/portland1 is now online.

Heartbeat Latency:
                        Heartbeat Latency Intervals (msec)
  Time Interval    [0-499]   [500-999]   [1000-3999]   [No Response]
  -------------    ------------------------------------------------
  02:00 - 02:54       887           0             0              0
  01:00 - 02:00         0           0             0              0
  00:00 - 01:00         0           0             0              0
  23:00 - 24:00         0           0             0              0
  22:00 - 23:00         0           0             0              0
  21:00 - 22:00         0           0             0              0
  20:00 - 21:00         0           0             0              0
  19:00 - 20:00         0           0             0              0
```

```
  18:00 - 19:00         0           0              0              0
  17:00 - 18:00         0           0              0              0
  16:00 - 17:00         0           0              0              0
  15:00 - 16:00         0           0              0              0
  14:00 - 15:00         0           0              0              0
  13:00 - 14:00         0           0              0              0
  12:00 - 13:00         0           0              0              0
  11:00 - 12:00         0           0              0              0
  10:00 - 11:00         0           0              0              0
  09:00 - 10:00         0           0              0              0
  08:00 - 09:00         0           0              0              0
  07:00 - 08:00         0           0              0              0
  06:00 - 07:00         0           0              0              0
  05:00 - 06:00         0           0              0              0
  04:00 - 05:00         0           0              0              0
  03:00 - 04:00         0           0              0              0

Heartbeat latency summary:
  0-499msec  : 100.00%
  500-999 msec :   0.00%
  1000-3999 msec :   0.00%
  No response  :   0.00%
```

show redundancy reboot-history

Purpose Use this command to show the history for all redundancy-related reboots.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show redundancy reboot-history`**

Default(s) None

Platforms any *except* ARX-VE

Guidelines Version is the software version running at the time the reboot was issued.
Time of reboot is a timestamp from the beginning of the reboot.
Message is the message that appeared on the Console to announce the reboot.

Sample See the sample output in Figure 15.10.

Related Commands redundancy

Figure 15.10 *Sample Output: show redundancy reboot-history*

```
prtlndA# show redundancy reboot-history

Version 5.02.000.12539 (Feb 18 2010 18:03:22) [nbuilds]
Time of reboot: Mon Mar  1 12:57:53 2010

Message: user initiated reboot: admin at console

-------------------

Version 5.02.000.12542 (Mar  1 2010 10:54:49) [nbuilds]
Time of reboot: Mon Mar  1 14:19:44 2010

Message: user initiated reboot: admin at console

-------------------

Version 5.02.000.12542 (Mar  1 2010 10:54:49) [nbuilds]
Time of reboot: Tue Mar  2 00:05:03 2010

Message: user initiated reboot: admin at console

-------------------

Version 5.02.000.12542 (Feb 26 2010 18:21:06) [nbuilds]
Time of reboot: Tue Mar  2 07:31:39 2010

Message: software reboot: Dual reboot request was issued
```

show redundancy resilver-timeout

Purpose Initial rendezvous and standard managed-volume processing involves duplicating the metalog (namespace transaction) data from the active peer to the backup. This process is called *resilvering*. Use this command to show the maximum time allowed for resilvering. If this time expires before resilvering is complete, the redundant pair cannot form.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show redundancy resilver-timeout`**

Default(s) None

Platforms any *except* ARX-VE

Guidelines The output is a single value, the timeout allowed for resilvering. You can use the resilver-timeout command to reset this. The show redundancy metalog command shows the current state of the resilvering process.

Sample

```
prtlndA# show redundancy resilver-timeout

Resilver Timeout:  6 minutes
```

Related Commands
resilver-timeout
redundancy
show redundancy
show redundancy metalog

suspend-failover

Purpose

There are occasions when site maintenance may inadvertently cause redundancy-inspired failovers and reboots. These can complicate the maintenance procedure. You can use the suspend-failover command to prevent any failovers until further notice.

You can later lift the suspension with no suspend-failover.

Mode

cfg-redundancy

Security Role(s)

network-engineer or crypto-officer

Syntax

```
suspend-failover
no suspend-failover [force]
```

force (optional) is only necessary when the peer switch is disabled or otherwise unreachable. This causes the current switch to assume the active role in the pair, whether or not it previously had the active role. This is similar to the behavior from the redundancy force-active command.

Important

If you use the force *option on the backup peer when the other peer is still active and connected, this command creates a "split-brain" condition.*

Default(s)

```
no suspend-failover
```

Platforms

any *except* ARX-VE

Guidelines

Use this command only under the guidance of F5 personnel.

This command locks the redundancy state until the suspension is lifted. The peers retain their active and backup roles during the suspension. The CLI prompts for confirmation before doing this; enter **yes** to proceed.

You can only use this command when redundancy is enabled, the peer is reachable and ready to take control, the quorum disk is functioning normally, and all critical services are healthy on the active peer. Use the show redundancy command to verify that the peer and quorum disk are functioning, and use show redundancy critical-services to check on critical services.

During a period of failover suspension, both peers continue to record conditions that would typically cause failovers and/or failover-related reboots. These appear in the syslog (use show logs syslog to view the syslog), and some result in SNMP traps (see the ARX SNMP Reference for a full list of traps). Traps and logs also appear for any suppressed failovers or reboots.

Guidelines: Force Option

As stated above, you should only use the force option if the peer is disabled or otherwise unreachable. If the peer is reachable and active, the force option may cause both peers to take the active role and work at cross purposes. This is called a "split-brain" situation, and it can create serious issues that are very difficult to repair. The CLI prompts for special confirmation before lifting suspension in this way.

A switch that is properly forced can successfully re-join its peer later, but there is another possibility of "split brain" during that rendezvous. The rendezvous happens automatically when you first boot the repaired peer. A "split-brain" may occur if the running switch has any connectivity issues during the rendezvous. Verify that the running switch has full network connectivity and all of its critical resources, especially the quorum-disk link, before booting its repaired peer.

Samples

```
prtlndA(cfg-redundancy)# suspend-failover

Are you sure you want to suspend redundancy? [yes/no] yes
```

suspends failovers and failover-related reboots on "prtlndA" and its redundant peer.

```
bstnA(cfg-redundancy)# no suspend-failover
```

lifts the failover suspension from "bstnA" and its redundant peer.

```
bstnA(cfg-redundancy)# no suspend-failover force

The force option should only be used when the peer switch is disabled.
The peer should be off or isolated from the network. Otherwise,
services could be severely disrupted.

Are you sure you want to proceed? [yes/no] yes
```

forces a resumption of standard-redundancy. This should only be used if the redundant peer is offline.

Related Commands

show redundancy
show redundancy critical-services

16

Active Directory Discovery

active-directory update seed-domain

Purpose

The ARX must understand the Active-Directory (AD) hierarchy in the Windows network. That is, each ARX-CIFS service (and, in some sites, the authentication software for the CLI/GUI) must know where to find the correct DC for each domain in the AD forest. The ARX database requires a full representation of the AD forest, with the IP addresses of at least one DC per domain.

Use the active-directory update seed-domain command to discover an AD forest in your network and add its representation to the ARX.

Mode

priv-exec

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
active-directory update seed-domain seed proxy-user proxy
  [domain-controllers max-dcs] [site-name site]
  [verbose] [tentative]
```

seed (1-255 characters) is the name of one domain in the forest. The ARX uses this domain name to begin its forest discovery. This becomes the name of the AD-forest in the ARX configuration.

proxy (1-32 characters) is a proxy-user with credentials for accessing the seed domain's DC(s). These credentials can belong to the seed domain itself, or any domain that is trusted by the seed domain. The ARX queries the DC for the names of other domains in the same AD forest.

max-dcs (optional, 1-100) sets a maximum number of DCs used in each domain. The ARX queries its DNS server to discover all the DCs in each domain; if the DNS server returns more DCs than ***max-dcs***, the ARX takes the top DCs from the DNS list. The ARX uses the order returned from DNS.

site (optional, 1-64 characters) identifies the AD site for the ARX. If it knows of multiple DCs that can answer the same query, the ARX prefers DCs in its own site (if there are any) over DCs in any other site. The site name is defined on a DC with the Active Directory Sites and Services plugin. The site name is case insensitive, so "boston" and "BOSTON" are equivalent. If you omit this, the ARX software uses the AD site configured for the ip proxy-address subnet. Use this option if the AD's site configuration does not include the proxy-IP subnet.

verbose (optional) causes the command to show the results of the forest discovery as it progresses.

tentative (optional) makes the ARX perform the AD-forest discovery without creating the actual active-directory-forest configuration.

Default(s)

max-dcs - all DCs returned from each DNS query.

site - the AD site configured on the external Active Directory for the proxy-IP subnet. This default requires that the proxy-IP subnet is defined in the AD; you can add the subnet on a DC with the Active Directory Sites and Services plugin.

Guidelines

The ARX uses a DNS query to find the DC(s) in this domain, then queries one of the DCs for the other domains in the forest. The DNS server that you use must be able to translate for all DCs in the target forest. Use the ip name-server command to add a DNS server to the ARX configuration.

The active-directory update seed-domain command automatically discovers an Active Directory (AD) forest and adds it to the ARX configuration. The switch uses this information to support CIFS authentications in single- and multi-domain environments. After a successful discovery, the ARX configuration contains an AD-forest object with the name of the ***seed*** domain that you provided.

This creates a report named "active-directory-*seed_domain*.rpt," where *seed_domain* is the seed domain that you chose in the command. The CLI displays the name of the report after you issue the command. Use show reports type AdUp to list all AD-discovery reports. To follow the progress of the AD-discovery operation, you can use tail reports *report-name* follow. Use show reports *report-name* to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands. If you want to truncate the report before it finishes, use the truncate-report command. See Figure 16.1 on page 16-6 for a sample report.

Use the show active-directory command to show the configuration of the AD forest, as recorded on the ARX. To see the status of all DCs in the forest, use the show active-directory status command.

Guidelines: Manual AD-Forest Configuration

As an alternative to this automatic-discovery process, you can use the active-directory-forest command to manually create or edit the AD-forest configuration. The active-directory-forest command brings you to gbl-forest mode, where you can add all types of domains and DCs in the forest. Typically, this is only used to add dynamic-DNS servers, described later.

Guidelines: Domain Types

The report and the show commands all identify a domain type for each of the forest's domains. The discovery process categorizes all of the domains as it finds them.

Domains in an AD forest have parent-child relationships based on their names. For example, the domain, "myco.com," can have child domains named "usa.myco.com" and "britain.myco.com." In this relationship, "myco.com" is called the *forest-root* domain and the other two domains are *child domains*.

An AD forest can also support trees that are outside the forest root's tree. These are domains with two-way-transitive-trust relationships with one or more domains in the root tree, but their domain names are entirely different. They are called *tree domains*. To continue the above example, a tree domain named "nonprofit.org" may also be in the AD forest, and it may have a child domain named "euro.nonprofit.org."

Guidelines: Dynamic DNS

Kerberos authentication requires that the forest's DNS servers have up-to-date mappings of the ARX's CIFS services and their IP addresses. To facilitate DNS updates as the ARX's CIFS services are added and removed, you can use the gbl-forest name-server command to identify a dynamic-DNS server for each domain in the forest. The discovery process does not discover dynamic-DNS servers.

In many cases, the DC itself can also double as the name server for its domain.

Guidelines: Trusts Between Forests

For AD forests with two-way-trust relationships, where clients from one forest are allowed to access CIFS services from the other forest, you can use the kerberos auto-realm-traversal command to automatically discover all such trust relationships. Alternatively, you can use the active-directory forest-trust command to declare a forest-to-forest trust relationship in the ARX configuration.

Guidelines: Offline DCs Slow the Discovery

If any of the discovered DCs are offline or unreachable, they slow the time required to discover the forest. Testing has shown that the forest discovery requires an additional 2.5 minutes for each offline DC; the delay may vary at your site.

Samples

```
bstnA# active-directory update seed-domain ny.com proxy-user ny_admin

Report File : active-directory-NY.COM.rpt
```

discovers an AD forest named "NY.COM" and adds the forest to the ARX configuration. This also creates a report, active-directory-NY.COM.rpt. See Figure 16.1 on page 16-6 for a sample report.

```
bstnA# active-directory update seed-domain vt.com proxy-user ny_admin verbose

Report File : active-directory-vt.com.rpt

DNS lookup vt.com from 172.16.108.139
    10.51.100.8
    172.16.213.8
    172.16.213.9

Enumerate Forest vt.com from 10.51.100.8
    Domain Name                                Pre-Win2k Name        Domain Type
    -----------------------------------------  --------------------  ------------
    BSH.ATLANTIC.ME.ORG                        BASSHARBOR            child-domain
    MCNIELS.VT.COM                             MCNIELS               child-domain
    VT.COM                                     VT                    tree-domain
    ATLANTIC.ME.ORG                            ATLANTIC              tree-domain

Domain Controller Discovered:

    Domain Type  Domain Name                              Pre-Win2k Name    IP Address
    -----------  ------------------------------------  ----------------  -----------
    tree-domain  ATLANTIC.ME.ORG                       ATLANTIC          172.16.210.7
    tree-domain  ATLANTIC.ME.ORG                       ATLANTIC          172.16.210.14
    forest-root  VT.COM                                VT                10.51.100.8
...
    child-domain BSH.ATLANTIC.ME.ORG                   BASSHARBOR        172.16.210.9
    child-domain BSH.ATLANTIC.ME.ORG                   BASSHARBOR        172.16.110.11

Change Summary:

    Site:  (null)
        No Domain Controllers discovered for this site.

    tree-domain ATLANTIC.ME.ORG    ATLANTIC  172.16.210.7 added.
    tree-domain ATLANTIC.ME.ORG    ATLANTIC  172.16.210.14 added.
    forest-root VT.COM    VT        10.51.100.8 added.
    forest-root VT.COM    VT        172.16.213.8 added.
    forest-root VT.COM    VT        172.16.213.9 added.
    child-domain MCNIELS.VT.COM MCNIELS   172.16.240.88 added.
    child-domain MCNIELS.VT.COM MCNIELS   172.16.240.70 added.
    child-domain BSH.ATLANTIC.ME.ORG BASSHARBOR   172.16.210.9 added.
    child-domain BSH.ATLANTIC.ME.ORG BASSHARBOR   172.16.110.11 added.
```

discovers the "vt.com" forest and shows the discovery process as it occurs.

Related Commands

ip name-server
proxy-user
active-directory-forest
name-server
kerberos auto-realm-traversal
active-directory forest-trust
show active-directory
show active-directory status

Figure 16.1 *Sample Report: Active-Directory Discovery*

```
bstnA# show reports active-directory-vt.com.rpt
**** Active-Directory Update Report: Started at 10/10/2011 12:05:48 -0400 ****
**** Software Version: 6.02.000.14267 (Oct  3 2011 20:08:59) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

**** Command : active-directory update seed-domain vt.com proxy-user ny_admin verbose

DNS lookup vt.com from 172.16.108.139
    172.16.213.79
    172.16.213.8
    172.16.213.9

Enumerate Forest vt.com from 172.16.213.79
    Domain Name                                Pre-Win2k Name        Domain Type
    -----------------------------------------  --------------------  -------------
    BSH.ATLANTIC.ME.ORG                        BASSHARBOR            child-domain
    MCNIELS.VT.COM                             MCNIELS               child-domain
    VT.COM                                     VT                    tree-domain
    ATLANTIC.ME.ORG                            ATLANTIC              tree-domain

Current Configuration:


  active-directory-forest VT.COM

    exit


Change Summary:

    Site:  (null)
        No Domain Controllers discovered for this site.

    tree-domain ATLANTIC.ME.ORG    ATLANTIC  172.16.210.7 added.
    tree-domain ATLANTIC.ME.ORG    ATLANTIC  172.16.210.14 added.
    forest-root VT.COM    VT        172.16.213.79 added.
    forest-root VT.COM    VT        172.16.213.8 added.
    forest-root VT.COM    VT        172.16.213.9 added.
    child-domain MCNIELS.VT.COM MCNIELS   172.16.240.88 added.
    child-domain MCNIELS.VT.COM MCNIELS   172.16.240.70 added.
    child-domain BSH.ATLANTIC.ME.ORG BASSHARBOR172.16.210.9 added.
    child-domain BSH.ATLANTIC.ME.ORG BASSHARBOR172.16.110.11 added.

New Installed Configuration:


  active-directory-forest VT.COM
```

```
    forest-root    VT.COM                                VT              172.16.213.79
    forest-root    VT.COM                                VT              172.16.213.8
    forest-root    VT.COM                                VT              172.16.213.9
    tree-domain    ATLANTIC.ME.ORG                       ATLANTIC        172.16.210.7
    tree-domain    ATLANTIC.ME.ORG                       ATLANTIC        172.16.210.14
    child-domain   MCNIELS.VT.COM                        MCNIELS         172.16.240.88
    child-domain   MCNIELS.VT.COM                        MCNIELS         172.16.240.70
    child-domain   BSH.ATLANTIC.ME.ORG                   BASSHARBOR      172.16.210.9
    child-domain   BSH.ATLANTIC.ME.ORG                   BASSHARBOR      172.16.110.11

    exit

**** Total processed:              9
**** Elapsed time:           00:00:00
**** AD Update Report: DONE at 10/10/2011 12:05:48 -0400 ****
```

description (gbl-proxy-user)

Purpose

Use the optional description command to set a descriptive string for the current proxy user. This appears in the show proxy-user command.

Use the no form of the command to delete the description.

Modes

gbl-proxy-user

Security Role(s)

crypto-officer

Syntax

`description` ***`text`***
`no description`

text (1-255 characters) is your description. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s)

`no description`

Guidelines

This description appears in the output for show global-config and show proxy-user.

Sample

```
stoweA(gbl-proxy-user[jckilley])# description "Jean-Claude's user identity"
```

specifies a description for the current proxy-user configuration.

Related Commands

show proxy-user
show global-config

proxy-user

Purpose

Use this command to provide the ARX software with Windows credentials, so that it can perform the following functions:

- Query the Windows Active Directory for its Windows Domains and domain-controller (DC) addresses (see active-directory update seed-domain);
- Validate inventory of back-end filers and import them into a managed volume;
- Maintain access control and permissions on back-end filers;
- Monitor free disk space on back-end filers;
- Enforce CIFS policy rules by migrating files between back-end shares.

You can also use a proxy user for other functions, such as logging into a filer's CLI and invoking snapshot operations. For these uses, the credentials may be a Unix username and password with permission to run snapshot commands.

Use the no form of the command to remove a proxy-user configuration.

Mode

gbl

Security Role(s)

crypto-officer

Syntax

```
proxy-user name
no proxy-user
```

name (1-32 characters) is a name you choose for the proxy user.

Default(s)

None

Guidelines

All CIFS namespaces require a proxy user, whether the network uses NTLM, NTLMv2, or Kerberos.

The switch uses a proxy-user configuration as its identity when it accesses back-end filers. This command places you in gbl-proxy-user mode, where you then use user (gbl-proxy-user) and windows-domain (gbl-proxy-user) to define a valid Windows username, password, and Windows domain for this proxy-user configuration.

The documentation for the user (gbl-proxy-user) command describes the exact privileges required for the Windows-user account. You can use the probe exports command with this proxy user to determine whether or not the proxy user has (at least) Backup Operator privileges. To prove that the proxy user has full Administrator privileges, use it with the show exports ... paths command; the paths are blank unless the proxy user belongs to the Administrators group on the filer.

After you create proxy-user configurations, you apply one or more to a namespace through the proxy-user (gbl-ns) command. Use the show proxy-user command to display all configured proxy users and their associated domains and usernames.

Guidelines: Changing the Password

For security reasons, you may choose to change the password for the proxy-user account from time to time. For a Windows proxy user, this change must occur in two places: the Active Directory, where the corresponding Windows user is defined, and in the ARX configuration. Use the following steps to make the change:

- Change the password on a Domain Controller (DC).
- Wait for all of the DCs in the proxy user's Windows Domain to synchronize their databases.
- After the DCs are synchronized with the new password, re-run the user (gbl-proxy-user) command to change the password on the ARX.

During the time that the proxy-user password on the ARX does not match the one in the external Active Directory, the volumes that use the proxy user cannot access any of their back-end-CIFS storage. This prevents the managed volumes from performing autonomous operations, such as policy-driven migrations, and causes the CIFS shares to go "offline." We recommend that you perform this procedure during non-busy hours.

Sample

```
bstnA(gbl)# proxy-user acoProxy2
bstnA(gbl-proxy-user[acoProxy2])#
```

starts the configuration of proxy user, "acoProxy2."

Related Commands

user (gbl-proxy-user)
windows-domain (gbl-proxy-user)
description (gbl-proxy-user)
proxy-user (gbl-ns)
show global-config **namespace**
show proxy-user
probe exports
show exports **... paths**

show active-directory

Purpose To support CIFS authentication, the ARX requires an accurate representation of your network's current Active Directory (AD) forest(s). Use the show active-directory command to review the AD-forest configuration as currently recorded on the ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show active-directory
show active-directory forest forest-name
show active-directory domain domain-name
```

forest-name (optional, 1-256 characters) identifies a particular forest to show.

domain-name (optional, 1-256 characters) identifies a particular domain.

Guidelines For the current status of the DCs in this output, use the show active-directory status command.

The show active-directory command displays the current AD configuration in one or two sections: Active Directory Domains and, if there are any forest-to-forest trusts, an additional Forest Trust table.

Guidelines: Active Directory Domains Output The Active Directory Domains section contains one table per AD forest. Each table has a Forest Name heading, showing the name configured with the active-directory-forest or active-directory update seed-domain command. The forest table contains one sub table per domain. Each sub table row contains the following fields:

forest-root, child-domain, or tree-domain is the type of the Windows domain in the AD forest, followed by the domain name and its pre-Windows 2000 name. This is the heading for the sub table of DCs for the domain. This domain name is discovered automatically, or it is manually set with the forest-root, child-domain, or tree-domain commands.

IP Address is the address of the Key-Distribution Center (KDC, a domain controller used for Kerberos) and/or dynamic-DNS server for this Windows domain. You establish this with the forest-root, child-domain, or tree-domain command, or the ARX automatically discovers it if you use active-directory update seed-domain. You can also use name-server to identify a dynamic-DNS server; in many cases, dynamic DNS runs on the same server as the KDC.

Services are KDC and/or DNS, as explained above.

Preferred shows whether or not this DC is on the "preferred" list for its domain. The choices are "YES" or "NO." The Kerberos software chooses its active DC(s) from the preferred list if there are any online. The preferred DCs are the DCs in the same AD site as the ARX (see active-directory update seed-domain or active-directory update forest), or you can manually set the DC preference with an optional flag in the forest-root, child-domain, or tree-domain commands.

Guidelines: Forest Trust Table

The Forest Trust table only appears if at least one forest-to-forest trust has been declared in the ARX configuration. You can declare a forest-to-forest trust with the active-directory forest-trust command. Each trust relationship appears as one row in the table, with the following fields:

Forest-1 and

Forest-2 are the two forests in the trust relationship. Clients from either of these forests can access CIFS services in the other.

Trust Type is always "bidirectional" in the current release. Clients in Forest-1 can access services in Forest-2, and clients in Forest-2 can access services in Forest-1.

Sample

`bstnA# show active-directory`

shows all active-directory forests on the ARX. See Figure 16.2 for sample output.

Related Commands

active-directory-forest
show active-directory status

Figure 16.2 *Sample Output: show active-directory*

```
bstnA# show active-directory
Active Directory Domains
------------------------

Forest Name:  MEDARCH.ORG
  forest-root    MEDARCH.ORG             MEDARCH
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    192.168.25.102        KDC DNS          YES
    192.168.25.104        KDC DNS          YES
    192.168.25.111        KDC              YES
    192.168.25.109        KDC              YES
    192.168.25.110        KDC              YES

  tree-domain    bostonmed.org         DNAS
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    172.16.74.89          KDC              YES
    172.16.74.88          KDC              YES

  tree-domain    fdtestnet.net         GUMMIE
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    172.16.168.22         KDC              YES
    172.16.168.21         KDC              YES

  child-domain   bostoncifs.fdtestnet.net       CUTE
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    10.19.230.88          KDC              YES
    10.19.230.94          KDC              YES

  child-domain   westcoast.medarch.org        EAST
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    192.168.202.9         KDC              YES
    192.168.202.10        KDC              YES
```

```
  child-domain   ma.ne.medarch.org                    HARVARD
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    192.168.25.105        KDC              YES
    192.168.25.103        KDC              YES

  child-domain   ne.medarch.org                NORTH
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    172.16.124.19         KDC              YES
    172.16.124.73         KDC              YES


Forest Name:  WELLS.ME.ORG
  forest-root    wells.me.org          NY
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    172.16.108.136        KDC              YES
    172.16.108.139        KDC              NO

  child-domain   adk.wells.me.org      ADK
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    172.16.110.8          KDC              NO
    172.16.110.5          KDC              YES

  child-domain   catskills.wells.me.orgCATSKILLS
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    172.16.120.22         KDC              YES
    172.16.120.5          KDC              YES


Forest Name:  VT.COM
  forest-root    VT.COM                VT
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    172.16.213.8          KDC              NO
    172.16.213.9          KDC              YES

  tree-domain    ATLANTIC.ME.ORG       NH
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    172.16.210.14         KDC              YES
    172.16.210.7          KDC              YES

  child-domain   MCNIELS.VT.COM        MCNIELS
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    172.16.240.70         KDC              YES
    172.16.240.88         KDC              YES

  child-domain   RDI.ATLANTIC.ME.ORG   NEWPORT
    IP Address            Services         Preferred
    --------------------  ---------------  ------------
    10.51.100.6           KDC              YES
    10.52.160.1           KDC              YES


Forest Trust
------------
```

```
Forest-1                          Forest-2                          Trust Type
--------------------------------  --------------------------------  ----------
MEDARCH.ORG                       wells.me.org                      bidirectional
wells.me.org                      vt.com                            bidirectional
```

show active-directory status

Purpose Use the show active-directory status command to see the status of all Domain Controllers (DCs) in the Active-Directory configuration.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show active-directory status [detailed]
show active-directory status forest forest-name [detailed]
show active-directory status domain domain-name [detailed]
```

detailed (optional) expands the output to include statistics about each DC's health.

forest-name (optional, 1-256 characters) identifies a particular forest to show.

domain-name (optional, 1-256 characters) identifies a particular domain.

Guidelines: Output The top line in the output shows the Offline timeout, which is the time that the ARX waits for a response from a DC before declaring it offline. This is a system-wide variable, set with the kerberos health-check threshold command.

The remaining output is divided into sections, one per control-plane processor. Each processor section has the following heading:

PROCESSOR *slot.proc*, where *slot.proc* identifies the processor.

Each processor section contains a subsection for each AD forest. The subsections have the following headings:

Forest is the name of the AD forest (set by the active-directory update seed-domain or the active-directory-forest command).

One table per DC appears in each forest subsection. Each DC table contains the following fields:

Domain Name is the Windows domain.

Domain Controller is the IP address of one of the domain's DCs.

Status is "Active," "Backup," "Offline," "Unusable," or "NoStatus."

- "Active" means that the processor is using this DC for authentication.
- "Backup" indicates that there are redundant DCs configured for this domain, where one or more other DCs are "preferred" over this one (see below), and the processor is currently using the DC's redundant peer(s).
- "Offline" indicates that the DC is unreachable. An unreachable DC is one that fails to respond to an LDAP query before the time set for the kerberos health-check threshold. The maximum time appears in the Offline timeout field at the top of the output.

Guidelines: Output (Cont.)

– "Unusable" means that the DC is not usable for forest-to-forest trusts (configured with the active-directory forest-trust command) because the ARX cannot verify that it is a Windows 2003 (or later) server. The DC must confirm that it is running Windows 2003 or a later release to be usable for these trusts.

  If such a DC is in the forest, no cifs service can support the recommended "constrained delegation" feature (described with the domain-join command).

– "NoStatus" appears very briefly before the processor first queries the DC. You are unlikely to ever see this status.

Preferred is "1" (one) if this DC is on the preferred list for its domain, or "0" (zero) otherwise. If any preferred DC is online for a given domain, a processor chooses the DC for Active status instead of its non-preferred peers. The preferred DCs are the DCs in the same AD site as the ARX (see active-directory update seed-domain or active-directory update forest), or you can manually set the DC preference with an optional flag in the forest-root, child-domain, or tree-domain commands.

Guidelines: Forest Trust Table

The Forest Trust table only appears if at least one forest-to-forest trust has been declared in the ARX configuration. You can declare a forest-to-forest trust with the active-directory forest-trust command. Each trust relationship appears as one row in the table, with the following fields:

Forest-1 and

Forest-2 are the two forests in the trust relationship. Clients from either of these forests can access CIFS services in the other.

Trust Type is always "bidirectional" in the current release. Clients in Forest-1 can access services in Forest-2, and clients in Forest-2 can access services in Forest-1.

Last Transition (UTC) is the timestamp showing the most-recent status change. (This time stamp is in UTC, not local time.) The ARX periodically checks each forest-root DC to confirm that it is online. If the forest has redundant root DCs, only the active DC has its status checked. The next field shows the results of the status check.

Status should be "Forest roots are online." This indicates that the trust is functional. If either or both are offline, check the connectivity to each DC, and the DC itself. You can use the show active-directory command to identify the forest root(s) for each forest.

Guidelines: Detailed Output

If you use the detailed flag in the command, the output changes to a separate table for each DC. The DC tables are organized by processor, then by forest. Each table contains the following fields:

Domain Name is the DC's Windows domain.

Domain Controller is the IP address of one of the domain's DCs.

Status is "Active," "Backup," "Offline," "Unusable," or "NoStatus," as described above.

Active Count is the number of times that the current DC has been Active for its domain.

Health Check Failure counts the number of times that the ARX failed to connect to the DC.

Health Check Timeout counts the number of times that the DC exceeded the kerberos health-check threshold before responding to the LDAP query.

Transition Total is the count of transitions to "Active" from some other status.

Last Transition (UTC) is the timestamp for the most-recent transition, in Coordinated Universal Time (known as UTC).

LDAP Health Check shows round-trip times for the LDAP query and response. These statistics only appear for a DC that is either "Active" or "Backup."

- Since last transition displays the average, minimum, and maximum round-trip times since the last time a new DC became reachable in this domain.
- Last 5/30/60 minutes shows the average round-trip times over the last 5 minutes, the last half hour, and the last hour.

Guidelines: Alternative Commands

Use show active-directory to see the switch's representation of the AD configuration.

Samples

```
bstnA# show active-directory status
```

shows the status of all AD forests configured on the current switch. See Figure 16.3 on page 16-18 for sample output.

```
bstnA# show active-directory status domain FDTESTNET.NET
```

shows the status of one AD domain. See Figure 16.4 on page 16-19 for sample output.

Related Commands

active-directory-forest
show active-directory

Figure 16.3 *Sample Output: show active-directory status*

```
bstnA# show active-directory status

Offline timeout is set to  3500 milliseconds.

PROCESSOR  1.1:

Forest :  MEDARCH.ORG
Domain Name                                        Domain Controller    Status    Preferred
-------------------------------------------------  -------------------  --------  ----
MEDARCH.ORG                                        192.168.25.102       Active    1
MEDARCH.ORG                                        192.168.25.104       Active    1
MEDARCH.ORG                                        192.168.25.111       Active    1
MEDARCH.ORG                                        192.168.25.109       Active    1
MEDARCH.ORG                                        192.168.25.110       Active    1
BOSTONMED.ORG                                      172.16.74.89         Active    1
BOSTONMED.ORG                                      172.16.74.88         Active    1
FDTESTNET.NET                                      172.16.168.22        Active    1
FDTESTNET.NET                                      172.16.168.21        Active    1
BOSTONCIFS.FDTESTNET.NET                           10.19.230.88         Active    1
BOSTONCIFS.FDTESTNET.NET                           10.19.230.94         Active    1
BOSTONCIFS.FDTESTNET.NET                           10.51.7.2            Active    1
WESTCOAST.MEDARCH.ORG                              192.168.202.9        Active    1
WESTCOAST.MEDARCH.ORG                              192.168.202.10       Active    1
WESTCOAST.MEDARCH.ORG                              10.51.2.4            Active    1
MA.NE.MEDARCH.ORG                                  192.168.25.105       Active    1
MA.NE.MEDARCH.ORG                                  192.168.25.103       Active    1
NE.MEDARCH.ORG                                     172.16.124.19        Active    1
NE.MEDARCH.ORG                                     172.16.124.73        Active    1

Forest :  WELLS.ME.ORG
Domain Name                                        Domain Controller    Status    Preferred
-------------------------------------------------  -------------------  --------  ----
WELLS.ME.ORG                                       172.16.108.136       Active    0
WELLS.ME.ORG                                       172.16.108.139       Active    0
ADK.WELLS.ME.ORG                                   172.16.110.8         Active    0
ADK.WELLS.ME.ORG                                   172.16.110.5         Active    0
YORK.WELLS.ME.ORG                                  172.16.120.22        Active    0
YORK.WELLS.ME.ORG                                  172.16.120.5         Active    0

Forest :  VT.COM
Domain Name                                        Domain Controller    Status    Preferred
-------------------------------------------------  -------------------  --------  ----
VT.COM                                             172.16.213.8         Active    0
VT.COM                                             172.16.213.79        Active    0
ATLANTIC.ME.ORG                                    172.16.210.14        Active    0
ATLANTIC.ME.ORG                                    172.16.210.7         Active    0
MCNIELS.VT.COM                                     172.16.240.70        Active    0
MCNIELS.VT.COM                                     172.16.240.88        Active    0
BSH.ATLANTIC.ME.ORG                                172.16.110.11        Active    0
BSH.ATLANTIC.ME.ORG                                172.16.210.9         Active    0


Forest Trust

    Forest-1             :  MEDARCH.ORG
    Forest-2             :  wells.me.org
    Trust Type           :  bidirectional
    Last Transition(UTC) :  05:37:17 01/27/2012
    Status               :  Forest roots are online
```

Figure 16.4 *Sample Output: show active-directory status domain FDTESTNET.NET*

```
bstnA# show active-directory status domain FDTESTNET.NET

Offline timeout is set to  3500 milliseconds.

PROCESSOR  1.1:

Forest :  MEDARCH.ORG
Domain Name                                      Domain Controller    Status    Preferred
-----------------------------------------------  -------------------  --------  ----
FDTESTNET.NET                                    172.16.168.22        Active    1
FDTESTNET.NET                                    172.16.168.21        Active    1


Forest Trust

    Forest-1             :  MEDARCH.ORG
    Forest-2             :  wells.me.org
    Trust Type           :  bidirectional
    Last Transition(UTC) :  05:37:17 01/27/2012
    Status               :  Forest roots are online
```

show proxy-user

Purpose The ARX needs user credentials (a username and password) for autonomous operations, such as migrating files from one filer to another. These credentials are called a *proxy user*. Use this command to view a list of configured proxy-user configurations on the switch.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show proxy-user [name]`

name (optional; 1-32 characters) is the name of a proxy user. If no name is specified, this command shows all configured proxy users.

Guidelines This command shows a table of proxy-user configurations, one per row. Each row contains the following columns:

Name identifies the proxy-user object. You choose this when you use the proxy-user command to create the proxy-user configuration.

Windows Domain identifies the domain of the proxy user, if there is one. This is typically an FQDN. Use the windows-domain (gbl-proxy-user) command to change this.

Pre-Win2k is the domain name that the ARX uses for its NTLM authentications; some filers do not accept FQDNs for NTLM. This is an old-style domain name, in the format used by Windows networks before Windows 2000. By default, this name is discovered automatically, through active-directory update seed-domain. If this domain's name was not discovered, the ARX uses the first name in the Windows Domain's FQDN, up to 15 characters before the first period. You can manually set the name with an option in the windows-domain (gbl-proxy-user) command, though we recommend setting it at the AD and discovering it from there.

User is the username used by this proxy-user object. If the ARX uses this proxy-user object, this is the username it presents to back-end filers. You can change this with the user (gbl-proxy-user) command.

Description shows a descriptive string for the proxy user, if one is defined. You can use the description (gbl-proxy-user) command to add a description to a proxy-user configuration.

Include a proxy-user "name" to show details for that configuration only.

Sample

```
bstnA> show proxy-user
```

shows all proxy-user configurations on the ARX. See Figure 16.5 on page 16-21 for sample output.

Related Commands proxy-user
windows-domain (gbl-proxy-user)
user (gbl-proxy-user)
show ntlm-auth-db

Figure 16.5 *Sample Output: show proxy-user*

```
bstnA> show proxy-user

Name             Windows Domain               Pre-Win2k        User Name
                 Description
-------------------------------------------------------------------------------
acoProxy1        WWMEDNET.COM                                  jqprivate
                 jq's admin account

acoProxy3        FDTESTNET.COM                BOSTONCIFS       jqtester

cifs_admin       MEDARCH.ORG                                   Administrator

nas_admin                                                      root

ny_admin         WELLS.ME.ORG                                  jqpublic

acoProxy2        MEDARCH.ORG                                   jqpublic
                 user with backup and admin creds on our servers
```

user (gbl-proxy-user)

Purpose

The ARX needs user credentials (a username and password) for autonomous operations, such as migrating files from one filer to another. These credentials are called a *proxy user*. Use this command to enter a username and password for the current proxy-user configuration.

Use the no user command to remove the username from the current configuration.

Mode

gbl-proxy-user

Security Role(s)

crypto-officer

Syntax

```
user username
no user
```

username (1-64 characters) is a valid username you choose for this proxy-user configuration. This is typically a Windows user, though it may be a Unix user for certain applications of the proxy user (for example, see the documentation for proxy-user (gbl-filer)). Any Windows proxy user must be a member of the Backup Operators or Administrators group on all filers that use these credentials; see below for details.

◆ Note

The password must be the same password set for this user in the external network. If you change the user's password on the network, you must also change it on the switch through this command.

Default(s)

None

Guidelines

The system prompts you for a password, then prompts you to validate the password by entering it a second time. Enter a password (it is masked with asterisk (*) characters as you enter it) and confirm when prompted.

Guidelines: CIFS Proxy User

A CIFS-supporting namespace can contain one or more managed volumes, and uses a proxy user as its identity for accessing CIFS shares behind those volumes. The proxy user requires specific privileges to access those CIFS shares. You must set these privileges on every CIFS filer or server behind every managed volume in the proxy user's namespace.

Minimally, the proxy user must have Backup Operator privileges at all of its filers. This means that the user you enter with this command must be a member of the Backup Operators group at every filer/server behind every managed volume. Backup Operator privileges are required for importing shares into the volume (with the enable (gbl-ns-vol-shr) command) and for migrating files onto the share. Also, verify that the user account has "full control" (read and change control) at the back-end shares to be used in managed volumes.

If a managed volume supports

- Access-Based Enumeration (ABE; see cifs access-based-enum)
- CIFS subshares (see the documentation for the filer-subshares command), or
- file-history archiving (a snapshot rule with a file-history archive),

the proxy user should belong to the Administrators group on that volume's filers/servers. The volume software needs this higher level of access for file-path queries, which are required for all of the above features.

If the user is a member of the local Backup Operators or Administrators group on a Windows cluster, add the Windows user to the group at every cluster node. The local groups are not shared in the cluster.

A Windows server with User Access Control (UAC) places further constraints on a CIFS proxy user. This requires a user account in the Domain Users group. UAC was introduced with Windows Server 2008.

Guidelines: Changing the Password

For security reasons, you may choose to change the password for the proxy-user account from time to time. For a Windows proxy user, this change must occur in two places: the Active Directory, where the corresponding Windows user is defined, and in the ARX configuration. Use the following steps to make the change:

- Change the password on a Domain Controller (DC).
- Wait for all of the DCs in the proxy user's Windows Domain to synchronize their databases.
- After the DCs are synchronized with the new password, re-run this command to change the password on the ARX.

During the time that the proxy-user password on the ARX does not match the one in the external Active Directory, the volumes that use the proxy user cannot access any of their back-end-CIFS storage. This prevents the managed volumes from performing autonomous operations, such as policy-driven migrations, and causes the CIFS shares to go "offline." We recommend that you perform this procedure during non-busy hours.

Sample

```
bstnA(gbl)# proxy-user acoProxy2
bstnA(gbl-proxy-user[acoProxy2])# user jqpublic
Password: ********
Validate Password: ********
bstnA(gbl-proxy-user[acoProxy2])#
```

Related Commands

proxy-user
show proxy-user

windows-domain (gbl-proxy-user)

Purpose

A *proxy user* is a set of user credentials (a username and password) that the ARX can use to access its back-end filers. The ARX uses the proxy-user credentials when it must perform autonomous operations, such as importing files into a managed volume. For a Windows proxy user, you can use this command to identify the Windows domain.

Use the no form of the command to remove the domain name.

Mode

gbl-proxy-user

Security Role(s)

crypto-officer

Syntax

```
windows-domain domain-name [pre-win2k-name old-style-domain]
no windows-domain
```

domain-name (1-64 characters) is the name of the Windows domain to identify with this Windows user. This must be a Fully-Qualified-Domain Name (FQDN) for the proxy user to authenticate through Kerberos. Choose a domain in the same Active-Directory forest as the filers.

old-style-domain (optional, 1-15 characters, no periods) is the legacy Windows domain for the proxy-user account. The proxy user may have to use NTLM when it authenticates with back-end filers, and not all back-end filers accept an FQDN for NTLM authentication. Use this option only if the default (below) is incorrect; this option is rarely necessary.

Default(s)

old-style-domain - The old-style name discovered with active-directory update seed-domain. If the old-style name was never discovered for this domain, the ARX uses the first part (before the first "."), up to 15 characters) of the FQDN in the ***domain-name***.

Guidelines

When assigned to a namespace through proxy-user (gbl-ns), a proxy-user provides a login ID for the ARX to access back-end filers. The ARX uses this account for imports, file migrations from one back-end filer to another, and other autonomous operations.

Ideally, the domain should be one that the ARX has already discovered with active-directory update seed-domain. If you are creating the proxy user before you run that command, the CLI warns that the domain is unknown. This warning is designed for situations where the domain is typed incorrectly, and is benign in this case.

Use the show proxy-user command to view all configured proxy users, and their associated Windows domains and usernames.

Samples

```
bstnA(gbl)# proxy-user acoProxy2
bstnA(gbl-proxy-user[acoProxy2])# windows-domain medarch.org

% INFO: The windows-domain name 'medarch.org' is not in
active-directory-forest configuration. The specified name is accepted, but
reconfigure the windows-domain if it is misconfigured.

bstnA(gbl-proxy-user[acoProxy2])#
```

sets the proxy-user's domain to medarch.org. The informational message is benign if it occurs before anyone has run the active-directory update seed-domain command.

```
stoweA(gbl)# proxy-user acoProxy3
stoweA(gbl-proxy-user[acoProxy3])# windows-domain FDTESTNET.COM
pre-win2k-name BOSTONCIFS
stoweA(gbl-proxy-user[acoProxy3])#
```

sets another proxy-user's domain to FDTESTNET.COM, but specifies a legacy-domain name of "BOSTONCIFS." The last option is only necessary if this pre-Win2K name was not discovered earlier with the active-directory update seed-domain command.

Related Commands

ntlm-auth-server
proxy-user (gbl-ns)
show global-config **namespace**
show proxy-user

17

Management Access

authentication

Purpose Use the authentication command to set a primary, secondary, or tertiary authentication service for the current management-access point.

Use the no form of the command to remove an authentication service.

Mode cfg-mgmt-access

Security Role(s) crypto-officer

Syntax

```
authentication {primary | secondary | tertiary}
  {active-directory | radius | local}
no authentication {primary | secondary | tertiary}
```

primary | secondary | tertiary is a required choice for setting the precedence of this authentication service:

primary is the preferred authentication service,

secondary is the authentication service to be used if the primary service fails, and

tertiary is used if the secondary fails, too.

active-directory | radius | local is another required choice, for setting the authentication type.

active-directory authenticates through the externally-configured Active Directory. (You can use the active-directory update seed-domain command to automatically discover the Active Directory (AD) and all of its Domain Controllers.) This option allows Windows clients with sufficient privileges (see below) to log in with their Windows credentials.

radius authenticates using a RADIUS (Remote Authentication Dial-In User Service) server on the network.

local authenticates using services local to the ARX.

Default(s) None

Guidelines The management-authentication software uses its primary authentication method first and tries the secondary method only if the primary method fails. The software then drops down to the tertiary method if the secondary method fails.

Use the show management access command to show the current configuration mapping of authentication services to management-access points.

To remove an authentication service with no authentication, you must remove the lowest-precedence service first. For example, if you have a management service configured with primary, secondary, and tertiary authentication, you must remove the tertiary authentication service before you can remove the secondary service.

Guidelines: Active-Directory Authentication

The **active-directory** option requires additional configuration: you must also identify Windows-user groups with sufficient privileges for administrative access. Use the group command to create a group with the same name as the desired Windows group, then use windows-domain (gbl-group) to declare a domain for that group, and finally use the role command to add one or more administrative roles for that group. For example, you could create a "Domain Admins" group in the "medarch.org" domain and give it highly-privileged roles, such as "storage-engineer" and "crypto-officer." As another example, you could create a "Domain Users" group with the lesser role of "operator."

The ARX only considers a user's first 1000 Active Directory groups when using AD authentication. If the AD group needed for authentication is not within the first 1000 groups returned by the DC, the user will not be able to authenticate.

To ensure that you can log in during a large-scale AD failure, we recommend that you configure a secondary or tertiary local service behind any primary AD service.

Guidelines: RADIUS Authentication

Before you configure RADIUS authentication, use the show radius-server command to verify that a RADIUS server is properly configured. To insure against RADIUS-server failures, we recommend that you configure a secondary local service behind any primary RADIUS service.

Guidelines: Keep Each User Account in Only One Authentication Service

If the same username is defined for two or more authentication types (such as RADIUS and AD) with different passwords, the user can log in with either password. The ARX software uses the given username/password combination for all authentication methods until authentication succeeds, or until all authentication methods are exhausted. For example, if the "jsmith" user has one password defined at the RADIUS server and a different password defined in the Active Directory, and both those authentication services are assigned to the HTTPS access point, jsmith can log into the GUI with either password.

This situation can create security issues, so we recommend defining the user account in only one authentication database (RADIUS, AD, or locally on the ARX).

Guidelines: Authentication for the ARX API

A client application can access the *ARX API* to make queries about the ARX configuration and state. These queries come through one of two management access points, http-api or https-api. Each query from the client application results in a separate authentication.

We recommend using local authentication only for the API access points; active-directory or radius authentication may create an excessive load on your domain controllers (for active-directory) or RADIUS server. If you must duplicate accounts in more than one database, we recommend using the same password for all of them.

Samples

```
bstnA(cfg)# management access https
bstnA(cfg-mgmt-access[HTTPS])# authentication primary active-directory
```

sets authentication configuration for HTTPS (GUI) access: the primary authentication service is Active Directory. This allows an administrator to log into the GUI with his or her Windows username and password, provided the username is a member of a properly-configured Windows group (see the *Guidelines* above).

```
bstnA(cfg-mgmt-access[Telnet])# no authentication secondary
```

removes the secondary authentication service from Telnet.

Related Commands

management access
show management access
radius-server
show radius-server

auth-port (gbl-radius)

Purpose

Most RADIUS servers listen at a well-known port. By default, the ARX sends all RADIUS traffic to that port. Use the auth-port command to change the destination-port number for outbound RADIUS traffic.

Use the no form of the command to reset the port number to the default.

Mode

gbl-radius

Security Role(s)

crypto-officer

Syntax

auth-port ***port-number***
no auth-port

port-number (1024-65535) is the destination-port number at the RADIUS server.

Default

1812

Sample

```
bstnA(gbl)# radius-server 192.168.25.207
bstnA(gbl-radius[192.168.25.207])# auth-port 5555
bstnA(gbl-radius[192.168.25.207])# ...
```

changes the port number to 5555 for the RADIUS server at 192.168.25.207:

Related Commands

radius-server
key (gbl-radius)
retries (gbl-radius)
timeout (gbl-radius)

clear session

Purpose Use the clear session command to end a current CLI or GUI user session on the ARX.

Mode priv-exec

Security Role(s) crypto-officer, network-engineer, or network-technician

Syntax **`clear session`** ***`login-id`***

login-id (1-128 characters) is the ID of the session to disconnect. The show sessions command shows a list of valid IDs.

Guidelines If you choose the login ID for the current CLI session, a prompt warns you that you will be disconnected if you proceed. Enter **yes** to clear the CLI session and disconnect.

Use the show sessions command to view the list of current user sessions.

Samples

```
bstnA# clear session 25250
```

ends session 25250, a CLI session.

```
bstnA# clear session F8DAB350CE624E50AD767A717E931826
```

ends a GUI session.

Related Commands show sessions

clear statistics authentication

Purpose Use this command to clear all login statistics, shown in the show statistics authentication output.

Mode priv-exec

Security Role(s) crypto-officer, network-engineer, or network-technician

Syntax **`clear statistics authentication`**

Default(s) None

Guidelines Use the show statistics authentication command to view the number of successful and failed login attempts. This command clears all of those counters.

Sample `bstnA# clear statistics authentication`

Related Commands show statistics authentication

key (gbl-radius)

Purpose

Use the key command to set the key (password or shared secret) on the ARX to match the key on the RADIUS server. The switch and its RADIUS server must have identical keys.

Use the no form of the command to erase the key. Erasing the key disables authentication at the current RADIUS server.

Mode

gbl-radius

Security Role(s)

crypto-officer

Syntax

```
key
no key
```

Default(s)

None

Guidelines

The CLI challenges you for the key, and you must re-enter it at the Validate Key prompt. See the example below.

Sample

```
bstnA(gbl)# radius-server 192.168.25.201
bstnA(gbl-radius[192.168.25.201])# key
Key: $3cretPa$$w0rd
Validate Key: $3cretPa$$w0rd
bstnA(gbl-radius[192.168.25.201])# ...
```

The following example uses the “$3cretPa$$w0rd” key:

Related Commands

radius-server

management access

Purpose You can manage the ARX through the following management-access points:

- Console (a serial connection),
- Telnet (on a management interface),
- SSH (on a management interface),
- HTTP (on a management interface),
- HTTPS (on a management interface),
- HTTP for the API (on a management interface),
- HTTPS for the API (on a management interface), and
- SNMP.

Each of these services is gated by an *authentication service*. Use the management access command to start configuring authentication service(s) for a management-access point.

Use the no form of the command to remove all authentication services from a management-access point, thus shutting down access.

Mode cfg

Security Role(s) crypto-officer

Syntax

```
management access {console | telnet | ssh | http | https |
  http-api | https-api | snmp | all}

no management access {console | telnet | ssh | http | https |
  http-api | https-api | snmp | all}
```

Default(s) console - access permitted, authentication through a local service (as opposed to a remote Active-Directory service or RADIUS server)

telnet - access denied

ssh - access permitted through any management interface, local authentication

http - access denied

https - access permitted through any management interface, local authentication

http-api - access denied

https-api - access denied

snmp - access denied

Guidelines

Administrators can access the CLI through telnet or SSH, they can access the GUI through HTTP or HTTPS, and they can access the API through HTTP or HTTPS. SSH and HTTPS, the secure protocols, are available by default for the CLI and the GUI.

You can use this command to start provisioning the management-IP interfaces that accept these communication protocols, and the authentication services (Active Directory, RADIUS, or local) to be used. This command puts you into cfg-mgmt-access mode, where you use the permit command to permit access through one or more management IPs. You can also use the authentication command to configure a primary, secondary, and/or tertiary authentication service.

If you disable all management access to the GUI or the API, you shut down those services. Conversely, you can restart those services by permitting management access through any interface.

Use the show management access command to show the current configuration mapping of access points to their authentication services and available interfaces.

Guidelines: API

You can use a SOAP client to access the API. The HTTP-API port is 83, and the HTTPS-API port is 843. Use the following URL syntax to access the API documentation through HTTP:

http://*arx-management-ip*:83/arx-api/

where *arx-management-ip* is either the out-of-band management-IP address (interface mgmt) or an in-band (VLAN) management-ip address (interface vlan). The permit command determines which of these address types are available for access. Use show interface mgmt and/or show interface vlan to find the IP addresses for each interface.

This URL accesses the same documentation through HTTPS:

https://*arx-management-ip*:843/arx-api/

where *arx-management-ip* is a valid management IP, as described above.

You can use the show statistics api command to find usage statistics for the API.

The API supports notifications of file and directory changes in any managed volume. To support these API queries, you create a notification rule in any desired managed volume.

Samples

```
bstnA(cfg)# management access console
bstnA(cfg-mgmt-access[Console])# ...
```

starts configuring authentication for console access. The console port is a serial port, labeled Console on the front panel.

```
bstnA(cfg)# management access all
bstnA(cfg-mgmt-access) ...
```

starts configuring authentication for all management-access points at once.

```
bstnA(cfg)# no management access telnet
bstnA(cfg)# ...
```

prevents any management access through Telnet.

Related Commands

permit
authentication
show management access

permit

Purpose Use the permit command to allow access to the current management-access point. You can permit access through the out-of-band MGMT interface, any in-band (VLAN) interface, or all management interfaces. Use no permit to deny access.

Mode cfg-mgmt-access

Security Role(s) crypto-officer

Syntax

```
permit {vlan | mgmt | all}
no permit {vlan | mgmt | all}
```

vlan | mgmt | all is a required choice:

vlan permits (or denies) access through any in-band (VLAN) management interface. Use the show interface vlan command to see all in-band management interfaces.

mgmt permits or denies access through the out-of-band management interface, labeled MGMT on the front panel. Use the show interface mgmt command to see the configuration for this interface. This option is unavailable on the ARX-VE, which has no out-of-band management interface. This option also may not apply to the ARX-1500 or ARX-2500, where the out-of-band management interface may be re-purposed as a standard client/server port (see the documentation for interface mgmt).

all permits or denies access through all of the above.

Default(s) For each management-access point:

Telnet - no permit

Ssh - permit

Snmp - no permit

Guidelines This does not apply to the Console, which always permits serial access.

Use show management access to see which management-access points are currently available.

Samples

```
bstnA(cfg)# management access telnet
bstnA(cfg-mgmt-access[Telnet])# permit all
bstnA(cfg-mgmt-access[Telnet])# ...
```

permits Telnet access through MGMT and VLAN-management interfaces.

```
bstnA(cfg)# management access ssh
bstnA(cfg-mgmt-access[SSH])# no permit mgmt
bstnA(cfg-mgmt-access[SSH])# ...
```

denies SSH access through the out-of-band MGMT interface:

Related Commands

management access
show interface vlan
show interface mgmt
show management access

radius-server

Purpose Use the radius-server command to configure/identify a RADIUS server as an authentication provider for the ARX. You can use this command multiple times to configure multiple RADIUS servers.

Use the no form of the command to remove one RADIUS-server configuration.

Mode gbl

Security Role(s) crypto-officer

Syntax

```
radius-server hostname-or-ip-address
no radius-server
```

hostname-or-ip-address (1-128 characters or an IP address) is the hostname or IP address of a RADIUS server.

Default(s) None

Guidelines This command places you in gbl-radius mode. From gbl-radius mode, use the key (gbl-radius) command to enter the shared-secret from the RADIUS server. You can change the timeout and/or retry interval with the timeout (gbl-radius) and retries (gbl-radius) commands. If the RADIUS server listens at a port other than the well-known port, you can use the auth-port (gbl-radius) command to adjust the port configuration.

To use a RADIUS server, you must configure the server itself as well as the ARX. For details about configuring the external RADIUS server, see Chapter 8, *Configuring Management Access*, in the ARX® CLI Network-Management Guide.

Important

Before you remove a RADIUS server through the no *command, ensure a backup authentication method is in place for management services.*

Sample

```
bstnA(gbl)# radius-server 192.168.25.201
bstnA(gbl-radius[192.168.25.201])# ...
```

starts configuring a RADIUS server at 192.168.25.201:

Related Commands

auth-port (gbl-radius)
key (gbl-radius)
retries (gbl-radius)
timeout (gbl-radius)
show radius-server
show management access

retries (gbl-radius)

Purpose Use the retries command to change the connection-retry interval for communication with RADIUS servers.

Use the no form of the command to reset the retry interval to the default.

Mode gbl-radius

Security Role(s) crypto-officer

Syntax

```
retries number
no retries
```

number (3-65535) is the number of times the switch attempts to connect to the RADIUS server before declaring a failure.

Default(s) 3

Guidelines By default, the switch retries a failed connection three times before it quits. If all retries fail, the request fails. A failed request causes a fallback to the next level of authentication service (secondary or tertiary), if one is configured.

Sample

```
bstnA(gbl)# radius-server 192.168.25.207
bstnA(gbl-radius[192.168.25.207])# retries 4
bstnA(gbl-radius[192.168.25.207])# ...
```

sets the number of retries for the RADIUS server at 192.168.25.207:

Related Commands radius-server
timeout (gbl-radius)

show management access

Purpose Use the show management access command to show the authentication services configured for system management.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show management access`

Guidelines Service is Console, Telnet, SSH, HTTP, HTTPS, SNMP, HTTP-API, or HTTPS-API. This is the management-access point, a point of entry for a system administrator. Use the management access command to edit one or all of these management-access configurations.

Primary is the primary authentication service for this management-access point. The choices are "AD" (for Active-Directory authentication, which allows Windows users to authenticate with their Windows credentials), "local" (for authentication service that runs on the ARX), or "RADIUS" (for authentication at a remote RADIUS server).

Secondary and

Tertiary are backup authentication services. If the primary service fails, the switch uses the secondary service; if the secondary authentication fails too, the switch falls back to the tertiary service. Use the authentication command to reset the primary, secondary, or tertiary service.

Allowed Interface is "VLAN," "Management," or both. "VLAN" indicates that administrators can gain access through any in-band (VLAN) management interface. "Management" indicates that access is allowed through the out-of-band management interface, labeled MGMT on the front panel. Use the permit command to permit VLAN and/or MGMT access.

Sample See Figure 17.1, below.

Related Commands management access
authentication
permit

Figure 17.1 *Sample Output: show management access*

```
bstnA> show management access

Service    Primary    Secondary    Tertiary    Allowed Interface
-----------------------------------------------------------------
Console    Local      None         None        N/A
Telnet     AD         RADIUS       Local       VLAN/Management
SSH        AD         Local        None        VLAN/Management
HTTP       None       None         None        None
HTTPS      AD         Local        None        VLAN/Management
SNMP       N/A        N/A          N/A         VLAN/Management
HTTP-API   None       None         None        None
HTTPS-API  None       None         None        None
```

show radius-server

Purpose Use the show radius-server command to display a summary of all RADIUS servers known to the switch.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show radius-server`

Guidelines The show command displays multiple RADIUS servers in the order they were configured. The switch uses them in this order; if the first server fails, the switch uses the second one, and so on.

Hostname is the name or IP address of the RADIUS server.

Authport is the destination port at the RADIUS server.

Acctport is the port used for accounting; currently not used.

Timeout is the current connection-timeout interval.

Retries is the current connection-retry interval.

Sample `bstnA(gbl)# show radius-server`

shows a summary of all RADIUS servers known to the switch.

See Figure 17.2 on page 17-16 for sample output.

Related Commands radius-server
auth-port (gbl-radius)
retries (gbl-radius)
timeout (gbl-radius)
show management access

Figure 17.2 *Sample Output: show radius-server*

```
bstnA# show radius-server

Hostname                        Authport  Acctport  Timeout  Retries
-------------------------------------------------------------------
192.168.25.201                  1812      1813      3        3
192.168.25.207                  5555      1813      10       4
```

show sessions

Purpose Use the show sessions command to view all current management sessions on the switch.

Mode (any)

Security Role(s) crypto-officer, network-engineer, or network-technician

Syntax `show sessions`

Guidelines Each login session is one table in the output, with the following fields:

Session ID identifies each session. CLI sessions have integer IDs, and GUI sessions have string IDs. The flag, (local session), appears for the current CLI session.

Username identifies the administrative user.

Access is the type of connection. This is the same as the management-access point that you choose with the management access command (for example, console, ssh, or https).

Connect Time shows how long the session has been running displayed in minutes and hours. If the session has been running less than a minute, a dash is displayed instead of the time.

Source IP is the IP address from which the user logged in.

You can use the clear session with one of the IDs to clear that administrative session, thereby logging off the administrator.

Sample `bstnA# show sessions`

shows the current users logged in to the switch. See Figure 17.3, below.

Related Commands management access
clear session

Figure 17.3 *Sample Output: show sessions*

```
bstnA# show sessions
Connected Sessions
------------------

Session ID:                 F3E9BE85537273B4351658A87C010BE3
Username:                   admin
Access:                     https
Connect Time:               0 days, 00:02:08
Source IP:                  192.168.101.57

Session ID:                 8395 (local session)
Username:                   admin
Access:                     ssh
Connect Time:               0 days, 00:00:52
Source IP:                  172.16.100.183

Session ID:                 9632
Username:                   admin
Access:                     ssh
```

```
Connect Time:                 0 days, 00:00:04
Source IP:                    172.16.100.183

bstnA# ...
```

show ssh-host-key

Purpose Immediately after you use ssh-host-key generate to change the SSH host keys on the switch, you should distribute the new keys to remote SSH clients. Use the show ssh-host-key command to show the public host keys.

Mode (any)

Security Role(s) crypto-officer

Syntax `show ssh-host-key [dsa | rsa | rsa1]`

dsa | rsa | rsa1 selects a particular SSH key; by default, all of them are shown.

dsa is the key for DSA over SSHv2,

rsa selects the key for RSA over SSHv2, and

rsa1 chooses RSA over SSHv1. By default, SSHv1 is not supported, use ssh-v1 enable to enable SSHv1 support.

Guidelines Use the ssh-host-key generate command to change all private/public host-key pairs on the switch. Then use this command to see the new public keys and begin distributing them to SSH clients. Install the appropriate public key (DSA, RSA, or RSA1) at each client; refer to the client's SSH documentation for instructions.

Sample

```
bstnA(cfg)# show ssh-host-key

ssh-dss
AAAAB3NzaC1kc3MAAACBAMGkquY++u13m5kMOveypKhx3BngJQws8/HxR/M+SHAx31UiWc
gSCL3DrB5ZrcQbD6ZKljAgqqkiDCDrCBg8EjTszpFlpe5k5ajpBtTvVQQPsi3o2LtutSoC
UEOBeoKqi9BIESqXoCQg0uuysnhWbQ3a7uRoQEHMO1a+AuHw1R7hAAAAFQCio48Mbumi4J
n+czKi311T+6GgSQAAAIAEJ3JRdANMfIGsQihWAXqVfnUcrhoLq5qJy3lxm41kWFf8P4L5
S30mfyGxg6ewW5GliHq3BDylRX1+Q5YLS5isvmQ+oCXjrg6agdoE2dXMUzizdkkIQRAxg1
l0lzvH+APxyNnPqmsa+fmPRC9X7l2nynlTynFGRdrZPFF+jI/32AAAAIBJPNit7Ax2xXe5
oKGvZ+BucSeAhL1sErEM+WwUg8YjK4zN7qooGyCVST85wgvGwGoqnv7pzVCM8LCwJnA2zV
/edBMSmfa8xgPIvkNYQDZ8xpeosTkVM6QeLhIeNjWi6vYInVs+HRrcKDq+k31xTKXmdAEB
m8iAfLZt15rkEHq+lg==


ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEAtB123IvN5D+nLB2eBH1XpC+OC+zvZpO/4v4nVwgX1M
aJoadWVIGSGp2rfO+rnEuh3UBNWP29uX1TWUqpyVqKYQLp/XHR7T6GYNnE3B4ACK58duLS
Z0c6kVIutsGWezpduvCjDhLCcm47V1506yhSSLARDvu4fPoG6r7zP1+TPK0=


1024 35
142987056756464496259115734760744315430018624820683986261988889498727 7
722832727925028180213482618739580704644443072907728765004995887397578 8
511113943366051258374865993944018298903508681768991335747976752760321 8
475683693785077295850663253229434184138312962732274986321738126633149 7
407257263010706185458244275 59
```

shows all three public host keys.

Related Commands ssh-host-key generate
ssh-v1 enable

show statistics authentication

Purpose Use this command to show the statistics on successful and failed management authentications.

Mode (any)

Security Role(s) crypto-officer, network-engineer, or network-technician

Syntax **`show statistics authentication`**

Default(s) None

Guidelines Use this command to display the authentication-login statistics. Use the clear statistics authentication command to clear the current authentication statistics and reset the counters.

For configuration information, use the show management access command.

Sample

```
bstnA# show statistics authentication

                          Success   Failure
                          -------   -------
Telnet
  Local database                0         0
  Active Directory              0         0
  RADIUS                        0         0

SSH
  Local database               73         0
  Active Directory              0         0
  RADIUS                        0         0

Console
  Local database                0         0
  Active Directory              0         0
  RADIUS                        0         0

API
  Local database                0         0
  Active Directory              0         0
  RADIUS                        0         0

HTTP/HTTPS
  Local database                0         0
  Active Directory              0         0
  RADIUS                        0         0

Totals                          73        0
```

displays the login authentication statistics.

Related Commands clear statistics authentication
show management access

ssh-host-key

Purpose

Some installations manage their SSH host keys with an external application, such as PuTTYgen (on Windows) or ssh-keygen (on Unix). You can use the ssh-host-key command to install a host key from one of these applications onto the ARX.

Mode

config

Security Role(s)

crypto-officer

Syntax

```
ssh-host-key {rsa | dsa}
```

dsa | rsa selects the type of SSH key to set.

dsa is the key for DSA, and

rsa selects the key for RSA over SSHv2.

Default(s)

None

Guidelines

The CLI prompts for the host key after you enter the command. Copy the private host key from your external application and paste it at this prompt.

The host key must have the following format:

```
-----BEGIN RSA PRIVATE KEY-----
key-text
-----END RSA PRIVATE KEY-----
```

where *key-text* has a <Return> character at the end of each line.

This is the format of a typical text file for a private key. PuTTYgen creates a text file with this format if you load a generated key and select **Conversions -> Export OpenSSH key**. This is also the format you find in /etc/ssh/ssh_host_*sa_key on a Linux box with SSH.

As an alternative to installing a pre-generated host key, you can use ssh-host-key generate to generate random host keys for the ARX.

Guidelines: Installing the Public Key on SSH Clients

The ARX generates a new public key to associate with this new private key. The new public key is accessible through show ssh-host-key. You copy the appropriate public key to trusted clients, as described below. The private key remains hidden on the switch.

After the new host key is installed, client machines warn their users that the switch's host key has changed. They must update to the new host key before they authenticate with the ARX. Right after generating new host keys, use the show ssh-host-key command to show the new public key, then copy and paste it onto the client machine. Refer to the client machine's SSH documentation for the specific configuration file.

SSH clients have the option to skip the host-key check, but we recommend checking the host-key at all times.

Sample

```
bstnA(cfg)# ssh-host-key rsa
Host Key: -----BEGIN RSA PRIVATE KEY-----
Host Key: MIICWgIBAAKBgQC2SZyBkHgOwnB9Stp+MUMOC0cshEyCsLAovBWFUDgGisjRroBO
Host Key: 5XYm+tmMBVmHzQayILI5XwCUl31dD1O0MKj09G9GFtSs376wVG8HJmfg2dpQmZWc
Host Key: /TVBxga6lc75y38aeXYWDgpOj4NU9V/jM6nz1Sy37g4EThrZjJolM9digwIBJQKB
Host Key: gB2PZXzLUb00SZfG8v+09h3z/bQxISnzI3w6LQDck3auuMgOdavt2874PvQciw8a
Host Key: VBzitR4PZ+C+nrUlFH4VuohP7vYItNY/JC95lTpLaY3sqrg8n1GmaYtVtzvtihPm
Host Key: 5Dq+a1hZIn9NPe3WsGghqFDzpubPzyx6uqpBwdYr+ppdAkEA7pYN47KnHIAxNDHT
Host Key: 9jvFXPWQj8HFduPIALQzz23MvNiB4r9QpvglQ83zJf7KK6plik2JKeD72/PED9wa
Host Key: m0wj7wJBAMOXoFLG3JM0r7J6S7R2hZ7H+pP9YAQb899olTg4XP7R2QB2JJa9bZoM
Host Key: 0c4hQY3G8GxZ9MKvdif9UOc9Yd6shq0CQQChNPSgxNHHJi8VbcaRngHk3UYARKgE
Host Key: OQqYsRwWiHWUWu/6EpBw0S3viye/uf45LfGUzJryyHLS4voYjczJy6mTAkBkcHTs
Host Key: PJrIIflbppWUBYLi1WT/4vn0RbSVUWFiDxr/ZNc3lp0qz/oC/6oCERPTLs581D9q
Host Key: 5Hrx8MiuGJoKj/IhAkAmF6NwzXu7UkWu+ToTh/RUfTpkA4JrF5yhZB5LubcdxYS5
Host Key: NilixpMbFWmXx5JFDBb2cKjLyTXKLIQ9ot9Bp3SN
Host Key: -----END RSA PRIVATE KEY-----
bstnA(cfg)#
```

installs a new SSH host key onto the “bstnA” switch. Note that the above sample is a single paste operation; the “Host Key” prompt repeats after each <Return> character in the key.

Related Commands show ssh-host-key
ssh-host-key generate

ssh-host-key generate

Purpose An attacker could use one of the ARX's SSH host keys to pose as the ARX for a "man-in-the-middle attack." To prevent this, you can periodically regenerate the host keys at the switch and distribute them to your SSH-client machines. Use the ssh-host-key generate command to generate new host keys.

Mode priv-exec

Security Role(s) crypto-officer

Syntax **`ssh-host-key generate`**

Default(s) None

Guidelines The CLI warns you that this forces all clients to update their host key for the ARX. Enter **yes** to proceed.

This generates three pairs of host keys, one for each encryption protocol (RSA for SSHv1, RSA for SSHv2, or DSA). Each pair has one private key (stored securely in the switch) and one public key (accessible through show ssh-host-key). You copy the appropriate public key to trusted clients, as described below. The private key remains hidden on the switch.

After the new host keys are generated, client machines warn their users that the switch's host key has changed. They must update to the new host key before they authenticate with the ARX. Right after generating new host keys, use the show ssh-host-key command to show the new public keys, then copy and paste the appropriate key onto the client machine. Refer to the client machine's SSH documentation for the specific configuration file.

SSH clients have the option to skip the host-key check, but we recommend checking the host-key at all times.

To install a pre-created host key on the ARX, use the ssh-host-key command in cfg mode.

Sample

```
bstnA# ssh-host-key generate
Warning: this command replaces the current SSH host keys with new key
material. Host public keys must be redistributed to SSH clients
following this action. Proceed? [yes/no] yes
bstnA#
```

generates new SSH host keys for the "bstnA" switch.

Related Commands show ssh-host-key
ssh-host-key

ssh-v1 enable

Purpose SSHv1 (as opposed to SSHv2) has well-documented security holes and is therefore not supported by default. Some installations only support SSHv1, so you can use the ssh-v1 enable command to enable it on the current switch.

Use the no form of the command to disable SSHv1 and return to v2-only support.

Mode cfg

Security Role(s) crypto-officer

Syntax
```
ssh-v1 enable
no ssh-v1 enable
```

Default(s) no ssh-v1 enable

Guidelines We do not recommend supporting SSHv1 for management access. This command exists for sites where there are SSH clients that do not support SSHv2.

SSHv2 is enabled in any case.

Samples
```
bstnA(cfg)# ssh-v1 enable
```
enables SSHv1 (in addition to SSHv2).

```
bstnA(cfg)# no ssh-v1 enable
```
disables SSHv1.

Related Commands

timeout (gbl-radius)

Purpose Use the timeout command to change the timeout interval for communication with RADIUS servers.

Use the no form of the command to reset the timeout interval to the default.

Mode gbl-radius

Security Role(s) crypto-officer

Syntax

```
timeout milliseconds
no timeout
```

milliseconds (3-65535) is the number of milliseconds to wait for a server response before timing out.

Default(s) 3

Guidelines After the timeout expires, the switch retries its request as many times as specified through the retries (gbl-radius) command. If all retries fail, the request fails, which causes fallback to the next level of authentication service (secondary or tertiary), if one is configured.

Samples

```
bstnA(gbl)# radius-server 192.168.25.207
bstnA(gbl-radius[192.168.25.207])# timeout 10
bstnA(gbl-radius[192.168.25.207])# ...
```

changes the timeout to 10 milliseconds for the RADIUS server at 192.168.25.207:

Related Commands radius-server
retries (gbl-radius)

18

CIFS Authentication

active-directory-forest

Purpose

A CIFS front-end service must understand the Active-Directory (AD) hierarchy in its network. That is, the CIFS service must know where to find the correct DC for each domain in the AD forest. The ARX database requires a representation of the AD forest, with the IP addresses of the DCs that control each domain.

Use the active-directory-forest command to start manually configuring one AD-forest representation on the ARX.

Use no active-directory-forest to remove one AD-forest configuration from the switch.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
active-directory-forest forest-name
no active-directory-forest forest-name
```

forest-name (1-256 characters) is the name you choose for the forest.

Default(s)

None

Guidelines

As an alternative to this command, you can automatically discover an AD forest in your network by using the active-directory update seed-domain command. After the forest is configured on the ARX, you can periodically use active-directory update forest to re-synchronize the ARX configuration with the actual AD-forest configuration.

The active-directory-forest command starts the process of manually describing an AD forest to the ARX. This command places you in gbl-forest mode, where you identify all of the domains in the AD forest, along with the DCs that manage them. Use the forest-root command to identify each DC for the forest root.

Domains in an AD forest have parent-child relationships based on their names. For example, the forest-root domain, “myco.com,” can have child domains named “usa.myco.com” and “britain.myco.com.” The “usa.myco.com” domain can have its own child domain named “ne.usa.myco.com.” From gbl-forest mode, use the child-domain command to identify a child domain.

A forest can also support trees that are outside the root’s domain. These are domains with two-way-transitive-trust relationships with one or more domains in the root tree. To continue the above example, a tree domain named “nonprofit.org” may also be trusted by “myco.com,” and it may have a child domain named “euro.nonprofit.org.” From gbl-forest mode, use the tree-domain command to identify an external tree domain. You use the child-domain command to assign a child to the tree domain; the name links each child to the correct parent.

Guidelines: Dynamic DNS

Kerberos authentication requires an up-to-date DNS database to map front-end CIFS services to their IP addresses. To facilitate DNS updates as CIFS services are added and removed, you can use the name-server command to identify a dynamic-DNS server for each domain in the forest. In many cases, the DC itself can also double as the name server for its domain.

Guidelines: Forest-to-Forest Trusts

For AD forests with two-way-trust relationships, where clients from one forest are allowed to access CIFS services from the other forest, you can use the active-directory forest-trust command to declare that trust relationship in the ARX configuration.

Guidelines: Removing an AD Forest

You cannot remove an AD-forest configuration if it contains any child domains or tree domains, or if it is part of a forest-to-forest trust. You can manually remove all of the forest's child domains (with no child-domain) and tree domains (no tree-domain), or an active-directory update forest operation may automatically remove all of them. To remove any forest trust, use no active-directory forest-trust.

After the last child domain, tree domain, an/or trust is removed, you can use no active-directory-forest.

◆ Note

You cannot remove an AD forest if any CIFS service is joined to its root domain (see domain-join). From gbl-cifs mode, use the no cifs *command to remove CIFS service, then remove the forest.*

Samples

```
bstnA(gbl)# active-directory-forest sierra
bstnA(gbl-forest[sierra])# ...
```

creates a forest named "sierra."

```
bstnA(gbl)# no active-directory-forest ADtrial
bstnA(gbl)# ...
```

removes a forest named "ADtrial."

Related Commands

active-directory update seed-domain
active-directory update forest
forest-root
child-domain
name-server
tree-domain
active-directory forest-trust
cifs
domain-join
enable (gbl-cifs, gbl-nfs)
show active-directory

active-directory forest-trust

Purpose Windows 2003 servers can support a two-way trust between AD forests. Clients from one forest can access CIFS services in the other, trusted, forest. You can use the active-directory forest-trust command to reflect this trust relationship in the ARX configuration.

Use the no form of the command to remove the forest-to-forest trust between two forests.

Mode gbl

Security Role(s) storage-engineer or crypto-officer

Syntax

```
active-directory forest-trust forest-1 forest-2
no active-directory forest-trust forest-1 forest-2
```

forest-1 (1-256 characters) is one of the two forests in the trust relationship. This must have a Windows 2003+ server as its forest-root; earlier versions of Windows do not support forest-to-forest trusts.

forest-2 (1-256 characters) is the other forest in the trust relationship, with the same restrictions as above.

Default(s) None

Guidelines This command is unnecessary if you use kerberos auto-realm-traversal.

Both AD forests must be configured as active-directory-forests on the ARX, must have Windows 2003+ DCs as their forest-roots, and must have a forest-function level of Windows 2003+ for this command to succeed. An error message explains any such configuration errors, and the command fails.

You can use the active-directory update seed-domain command to automatically discover an AD forest in your network and add it to the ARX configuration.

The forest-to-forest trust must exist externally, at the DCs, before you use this command to record the trust in the ARX configuration.

This records a two-way trust between the two forests. Clients from ***forest-1*** can access CIFS services in ***forest-2***, and clients in ***forest-2*** can also access CIFS services in ***forest-1***.

The forest-to-forest trust is direct, not transitive. Suppose a trust existed between forest A and forest B, and another trust existed between forest B and forest C. This does not imply that a client in forest A can access a CIFS service in forest C. A new trust must be established between forests A and C before the clients in one forest can access any services in the other. This reflects the forest-to-forest-trust behavior in Windows servers.

Because all forest trusts are direct, at least one CIFS service (see cifs) must be configured in one of the forests (***forest-1*** or ***forest-2***) for the trust to be relevant.

Guidelines: Selective Authentication

If any of the forest trusts in the Windows network use "Selective Authentication," the ARX must use a special algorithm to support clients in those forests. Use the kerberos auto-realm-traversal command to support remote forests with "Selective Authentication" set. As mentioned above, the kerberos auto-realm-traversal command makes the current command unnecessary; it enables automatic discovery of all trust relationships.

Samples

```
bstnA(gbl)# active-directory forest-trust medarch.org ri.com
```

establishes a two-way trust between two AD forests, 'medarcv' and 'ri.com.'

```
bstnA(gbl)# no active-directory forest-trust medarch.org ri.com
```

removes the above trust relationship. After this command, clients in the 'ri.com' forest can no-longer access any CIFS service in the 'medarcv' forest, nor can clients in 'medarcv' access services in 'ri.com.'

Related Commands

active-directory update seed-domain
active-directory-forest
forest-root
kerberos auto-realm-traversal
show active-directory
show active-directory status

active-directory update forest

Purpose A CIFS front-end service must understand the Active-Directory (AD) hierarchy in its network. That is, the CIFS service must know where to find the correct DC for each domain in the AD forest. The ARX database requires a full representation of the AD forest, with the IP addresses of at least one DC per domain.

If the actual forest changes, you can use the active-directory update forest command to discover the changes and add them to the ARX configuration.

Mode priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax

```
active-directory update forest forest-root proxy-user proxy
  [domain-controllers max-dcs] [site-name site]
  [verbose] [tentative]
```

forest-root (1-255 characters) is the name of the forest-root domain. This is often the name of the AD-forest itself, too.

proxy (1-32 characters) is a proxy-user with credentials for accessing the root domain's DC(s). The ARX queries the DC for the names of other domains in the same AD forest.

max-dcs (optional, 1-100) sets a maximum number of DCs used in each domain. The ARX queries its DNS server to discover all the DCs in each domain; if the DNS server returns more DCs than ***max-dcs***, the ARX takes the top DCs from the DNS list. The ARX uses the order returned from DNS.

site (optional, 1-64 characters) identifies the AD site for the ARX. This name is defined on a DC with the Active Directory Sites and Services plugin. The site name is case insensitive, so "portland" and "PORTLAND" are equivalent. If you omit this, the ARX software uses the AD site configured for the ip proxy-address subnet. Use this option if the AD's site configuration does not include the proxy-IP subnet.

verbose (optional) causes the command to show the results of the forest discovery as it progresses.

tentative (optional) makes the ARX perform the AD-forest discovery without creating the actual active-directory-forest configuration.

Default(s) ***max-dcs*** - all DCs returned from each DNS query.

site - the AD site configured on the external Active Directory for the proxy-IP subnet. This default requires that the proxy-IP subnet is defined in the AD; you can add the subnet on a DC with the Active Directory Sites and Services plugin.

Guidelines

You can only use this command for an AD forest that is already configured on the ARX. To automatically discover an AD forest and add it to the ARX configuration, use the active-directory update seed-domain command. You can also use active-directory-forest to manually add a forest to the ARX configuration, or to manually edit an existing forest's configuration.

The ARX uses a DNS query to find the DC(s) in this domain, then queries one of the DCs for the other domains in the forest. The DNS server that you use must be able to translate for all DCs in the target forest. Use the ip name-server command to add a DNS server to the ARX configuration.

The AD-update operation creates a report named "active-directory-*forest_name*.rpt," where *forest_name* is the forest that you identified in the command. The CLI displays the name of the report after you issue the command. Use show reports type AdUp to list all AD-update reports. To follow the progress of the AD-update operation, you can use tail reports *report-name* follow. Use show reports *report-name* to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands. If you want to truncate the report before it finishes, use the truncate-report command. See Figure 18.1 on page 18-9 for a sample report.

Use the show active-directory command to show the configuration of the AD forest, as recorded on the ARX. To see the current state of all DCs in the forest, use the show active-directory status command.

Guidelines: Domain Types

The report and the show commands all identify a domain type for each of the forest's domains. The update process categorizes all new domains as it finds them.

Domains in an AD forest have parent-child relationships based on their names. For example, the domain, "myorg.org," can have a child domain named "mygroup.myorg.org." In this relationship, "myorg.org" is called the *forest-root* domain and "mygroup.myorg.org" is a *child domain*.

An AD forest can also support trees that are outside the forest root's tree. These are domains with two-way-transitive-trust relationships with one or more domains in the root tree, but their domain names are entirely different. They are called *tree domains*. To continue the above example, a tree domain named "acme.com" may also be in the AD forest, and it may have a child domain named "coyote.acme.com."

Guidelines: Dynamic DNS

Kerberos authentication requires that the forest's DNS servers have up-to-date mappings of the ARX-CIFS services and their IP addresses. To facilitate DNS updates as the ARX-CIFS services are added and removed, you can use the gbl-forest name-server command to identify a dynamic-DNS server for each domain in the forest. The AD-update process does not discover dynamic-DNS servers.

In many cases, the DC itself can also double as the name server for its domain.

Guidelines: Trusts Between Forests

For AD forests with two-way-trust relationships, where clients from one forest are allowed to access CIFS services from the other forest, you can use the kerberos auto-realm-traversal command to automatically discover all such trust relationships. Alternatively, you can use the active-directory forest-trust command to declare a forest-to-forest trust relationship in the ARX configuration.

Samples

```
bstnA# active-directory update forest wells.me.org proxy-user ny_admin

Report File : active-directory-wells.me.org.rpt
```

updates an AD forest named "wells.me.org" and adds the forest to the ARX configuration. See Figure 18.1 for a sample report.

```
bstnA# active-directory update forest vt.com proxy-user ny_admin
tentative

Report File : active-directory-VT.COM.rpt
```

performs a tentative update of the "vt.com" forest. This re-discovers the forest without changing the ARX configuration. Details about the forest discovery appear in the AD-update report.

Related Commands

ip name-server
proxy-user
kerberos auto-realm-traversal
active-directory-forest
name-server
active-directory forest-trust
show active-directory
show active-directory status

Figure 18.1 *Sample Report: Active-Directory Update*

```
bstnA# show reports active-directory-wells.me.org.rpt
**** Active-Directory Update Report: Started at 10/10/2011 12:06:26 -0400 ****
**** Software Version: 6.02.000.14267 (Oct  3 2011 20:08:59) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

**** Command : active-directory update forest wells.me.org proxy-user ny_admin

DNS lookup wells.me.org from 172.16.108.139
    172.16.108.139
    172.16.108.136

Enumerate Forest wells.me.org from 172.16.108.139
    Domain Name                              Pre-Win2k Name        Domain Type
    ---------------------------------------  --------------------  -------------
    york.wells.me.org                        YORK                  child-domain
    adk.wells.me.org                         ADK                   child-domain
    wells.me.org                             WELLS                 tree-domain

Current Configuration:


  active-directory-forest WELLS.ME.ORG
    forest-root    wells.me.org                          WELLS             172.16.108.136
    child-domain   york.wells.me.org                     YORK              172.16.120.22
    child-domain   adk.wells.me.org                      ADK               172.16.110.8
    forest-root    wells.me.org                          WELLS             172.16.108.139
    child-domain   adk.wells.me.org                      ADK               172.16.110.5
    child-domain   york.wells.me.org                     YORK              172.16.120.5

    exit
```

```
Change Summary:

    Site:  (null)
        No Domain Controllers discovered for this site.
    No changes discovered.
**** Total processed:              0
**** Elapsed time:           00:00:00
**** AD Update Report: DONE at 10/10/2011 12:06:26 -0400 ****
```

child-domain

Purpose Use the child-domain command to manually add a child domain (or one of its DCs) to an Active Directory forest.

If this domain has no children of its own, you can use no child-domain to remove it. You can also use no child-domain to remove a redundant DC from the child domain's configuration.

Mode gbl-forest

Security Role(s) storage-engineer or crypto-officer

Syntax

```
child-domain domain-name dc [preferred]
no child-domain domain-name dc
```

domain-name (1-256 characters) is the child domain name. This must be in the same domain namespace as its parent domain; for example, if the parent domain is "nonprofit.org," this child domain could be "mycharity.nonprofit.org."

dc is the IP address (for example, 192.168.25.56) of the domain controller (DC) for this child domain. If redundant DCs are hosting the domain, repeat this command (using the same ***domain-name***) with the address of the redundant ***dc***.

preferred (optional) declares that you prefer to use this DC over others in the same domain. Whenever the ARX Kerberos process chooses from multiple DCs in the same domain, it chooses from the DCs labeled "preferred" first.

Guidelines This command may not be necessary if you use active-directory update seed-domain to automatically discover all child domains in the AD forest.

A child domain must have a parent, which must be the forest root (forest-root), a tree domain (tree-domain), or another child domain. As stated earlier, the child-domain name must begin with the name of its parent domain.

A child domain can be the parent to other child domains.

Redundant DCs may manage a child domain. Re-use the child-domain command (with the same domain name) for each redundant DC. If you prefer to use one or more DCs in the domain over the others, use the preferred flag to identify those DCs. The Kerberos process selects its active DCs from the preferred list. The process only stops using a DC if its connection fails; the kerberos health-check threshold command selects a timeout threshold for DC connections. The Kerberos process does not choose any DC outside the preferred list unless all of the preferred DCs have timed out. You can use show active-directory status to see which DCs are currently active for each domain.

To remove a DC from the preferred list, you can re-run the child-domain command without the optional preferred flag. You can use this flag to change the active DC at any time.

There are two circumstances under which you can use no child-domain: the domain has no child domains under it, or the domain has one or more redundant DCs configured and you want to remove one of them.

Samples

```
bstnA(gbl-forest[sierra])# child-domain ne.sierra.acopia.com 10.5.2.1
```

adds the ‘ne.siera.acopia.com’ child domain to the ‘sierra’ forest.

```
bstnA(gbl-forest[sierra])# child-domain ne.sierra.acopia.com 10.5.2.14
preferred
```

adds another DC to the ‘ne.siera.acopia.com’ domain, and declares it “preferred.” The ARX attempts to use this DC before it uses the one from the previous example.

```
bstnA(gbl-forest[sierra])# child-domain ne.sierra.acopia.com 10.5.2.14
```

takes the DC off of the “preferred” list for the ‘ne.siera.acopia.com’ domain. If there are no other preferred DCs for this domain, the ARX selects a DC from the non-preferred list (10.5.2.1 or 10.5.2.14, from these examples).

Related Commands

active-directory-forest
active-directory update seed-domain
kerberos health-check threshold
domain-join
forest-root
show active-directory
tree-domain

forest-root

Purpose Use the forest-root command to identify the forest-root domain (or one of its domain controllers) in the Active-Directory forest.

Use the no form of the command to remove the forest root, or one of the redundant domain controllers (DCs) that manages it.

Mode gbl-forest

Security Role(s) storage-engineer or crypto-officer

Syntax

```
forest-root domain-name dc [preferred]
no forest-root domain-name dc
```

domain-name (1-256 characters) is the domain name for the forest root.

dc is the IP address (for example, 192.168.25.56) of the forest root's DC. If redundant DCs are hosting the root domain, repeat this command (using the same ***domain-name***) with the address of each redundant ***dc***.

preferred (optional) declares that you prefer to use this DC over others in the same domain. Whenever the ARX Kerberos process chooses from multiple DCs in the same domain, it chooses from the DCs labeled "preferred" first.

Default(s) None

Guidelines This command may not be necessary if you use active-directory update seed-domain to automatically discover all domains in the AD forest.

Adding a forest root is the second step in manually creating an Active Directory forest. The forest-root command identifies the root domain for the forest, and one DC that manages the domain. If redundant DCs manage the root domain, you can use the command once for each DC.

An AD forest requires one forest root.

You can use child-domain to manually add one or more child domains to the forest root. The parent/child relationship between domains is established by their names: if "mycountry.gov" is the root domain, all of its child domains must reside in that domain (for example, "mystate.mycountry.gov" or "myprovince.mycountry.gov", but not "otherstate.org").

To add a disjoint tree of domains, use the tree-domain command. This is a tree of domains that is outside the root domain, but has a trust relationship with the domains in the forest root's tree. Just as with the forest root, you can use the child-domain command to add child domains to any tree domain.

If multiple DCs are configured to manage the forest root, you can use no forest-root to remove one of them at any time. Otherwise, the no forest-root command removes the entire root domain. To remove the forest-root domain, it cannot have any child domains, no CIFS services can be joined to the domain (through domain-join), and there can be no forest-trust relationships between this forest and any other forest (using no active-directory forest-trust).

Guidelines: Choosing the Active DC

If redundant DCs manage a domain, you can choose the DC(s) that you prefer for active use. Re-use the forest-root command (with the same domain name) for each redundant DC. You can use the optional preferred flag to identify preferred DCs. The Kerberos process uses DCs from the preferred list whenever possible. The process only stops using a DC if its connection times out; the kerberos health-check threshold command selects the timeout threshold. The Kerberos process does not choose any DC outside the preferred list unless all of the preferred DCs have timed out. You can use show active-directory status to see which DCs are currently active for each domain, including the forest-root domain.

To remove a DC from the preferred list, you can re-run the forest-root command without the optional preferred flag. You can use this flag to change the active DC at any time.

Samples

```
bstnA(gbl-forest[siera])# forest-root siera.acopia.com 10.5.1.23
preferred
```

sets the forest-root domain to 'siera.acopia.com,' and declares the DC at 10.5.1.23 to be preferred over other DCs in the same domain.

```
bstnA(gbl-forest[siera])# forest-root siera.acopia.com 10.5.1.23
```

re-sets the forest root, but removes the DC from the "preferred" list.

Related Commands

active-directory-forest
active-directory update seed-domain
kerberos health-check threshold
child-domain
domain-join
show active-directory
tree-domain
active-directory forest-trust

ip address (gbl-ntlm-auth-srv)

Purpose

This command is relevant to an NTLM Secure Agent, which is only necessary behind a cifs *service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the* domain-join *documentation, and avoid this CLI mode and command.*

At an installation that uses NTLM or NTLMv2 with unconstrained delegation, the ARX uses the *Secure Agent* as an NTLM proxy. You install the Secure-Agent software on a Windows DC (domain controller). When a client accesses CIFS on the ARX, the switch forwards the NTLM-authentication request to the Secure Agent. The Secure Agent accesses the password database on the Windows DC, and it tells the ARX whether or not the authentication succeeded.

Use this command to specify the IP address of the Windows DC where the ARX Secure Agent is installed

For multiple domains, enter the IP address of each DC that hosts an instance of Secure Agent.

Mode

gbl-ntlm-auth-srv

Security Role(s)

crypto-officer

Syntax

```
ip address server-address
no ip address
```

server-address (0.0.0.0-255.255.255.255) is the IP address of the Secure Agent's DC.

Use the no form of the command to remove the IP address/host.

Default(s)

None

Guidelines: Using Constrained Delegation Instead of a Secure Agent

As mentioned above, Secure Agents (and this command) are unnecessary if all of your CIFS services use constrained delegation. Constrained delegation is implemented at a local DC. You can use the probe delegate-to command to retrieve some necessary configuration information to set up constrained delegation for your CIFS services.

For sites where some DCs are older than Windows Server 2003 (which introduced constrained delegation), this CLI mode and command are required for NTLM support.

Guidelines: Running this Command

ARX Secure Agent software is installed on a Windows DC. You can install the Secure Agent on multiple DCs to support NTLM/NTLMv2 authentication for multiple domains.

Use this command to specify the IP address of the DC that runs an instance of Secure Agent. You can use the port command to specify the port at which the Secure Agent listens, assuming this port was also changed at the Secure Agent interface. See the *ARX Secure Agent Installation Guide* for more information about the Secure Agent.

Ensure there is a network link between the switch and each Windows domain controller that hosts an instance of Secure Agent.

Sample

```
gffstnA(gbl)# ntlm-auth-server ggh-dc
gffstnA(gbl-ntlm-auth-srv[ggh-dc])# ip address 192.168.25.109
```

specifies the IP address of the DC where the "ggh-dc" agent runs.

Related Commands ntlm-auth-server
port
show ntlm-auth-server

kerberos auto-realm-traversal

Purpose

On a Windows domain controller (DC), you can configure a forest-to-forest trust with *Selective Authentication*, where you list the clients in Forest A who have access privileges in Forest B. If the ARX is in Forest B, these privileged clients cannot use its services unless you enable *auto-realm traversal* on the ARX. Auto-realm traversal is an ARX algorithm that supports selective authentication. Auto-realm traversal also allows automatic support of all forest-to-forest trusts in your Active Directory, without explicitly stating them in the ARX configuration. This is a recommended feature.

You can use the no kerberos auto-realm-traversal command to stop supporting this feature, and make it impossible for clients to authenticate from an external forest with a "Selective Authentication" policy. This is not generally advisable, and should only be done on the advice of F5 Support.

Use the affirmative form of the command to return to the default Kerberos algorithm. This supports all Kerberos clients, including clients on the other side of a selective-authentication trust, and it automatically supports all forest-to-forest trusts.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no kerberos auto-realm-traversal
kerberos auto-realm-traversal
```

Default(s)

```
kerberos auto-realm-traversal
```

Guidelines

This is a system-wide parameter; the Kerberos software uses the same realm-traversal algorithm for all clients. If you use the default setting, you do not need to explicitly declare any forest-to-forest trusts with the active-directory forest-trust command. This is recommended.

If you choose to disable this feature, you can use the active-directory forest-trust command to inform the ARX of each forest-to-forest trust in your network. However, clients cannot connect to the ARX through a forest-to-forest trust with selective authentication.

Sample

```
bstnA(gbl)# kerberos auto-realm-traversal
```

supports clients on the other side of a forest-to-forest trust with "Selective Authentication" enabled.

Related Commands

active-directory forest-trust

kerberos health-check threshold

Purpose

The Kerberos software regularly sends a simple LDAP query to all of the DCs in its AD forest, and times the response from each DC. If a DC takes too long to respond, the Kerberos software declares it "offline" and sends an SNMP trap. You can use the kerberos health-check threshold command to determine the maximum allowable response time from any DC.

Use the no form of the command to return to the default response time.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

`kerberos health-check threshold` ***`milliseconds`***
`no kerberos health-check threshold`

milliseconds (optional 500-10,000) is the maximum amount of time allowed for a DC response.

Default(s)

2,000 (2 seconds)

Guidelines

This is a system-wide parameter; the maximum time applies to all DCs in all AD forests. For a complete list of all of these DCs, use the show active-directory command.

The LDAP query tests the DC's ability to perform lookups in the AD database. This is designed as an unobtrusive verification of the DC's basic functions.

The Kerberos process sends an LDAP request every 60 seconds. After the Kerberos software declares a DC offline, it continues to send the LDAP request at the same rate; if the DC responds within the allotted time, the Kerberos process declares it "online" again. The show active-directory status command displays the current state of all DCs known to the ARX.

Samples

```
prtlndA(gbl)# kerberos health-check threshold 5000
```
sets a 5-second threshold for DC responses.

```
stoweA(gbl)# no kerberos health-check threshold
```
sets the threshold to the default on the "stoweA" system.

Related Commands

active-directory update seed-domain
show active-directory
show active-directory status

name-server

Purpose

Kerberos authentication requires accurate DNS configuration for all CIFS services. To facilitate this, CIFS services can use *dynamic DNS* to register their host names with your DNS servers and update those servers whenever their hostname-to-IP mapping changes. Use the name-server command to identify a dynamic-DNS name server.

You can use no name-server to remove a name server.

Mode

gbl-forest

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
name-server domain-name ip-address
no name-server domain-name ip-address
```

domain-name (1-255 characters) is the name of the AD domain that is supported by the name server, below.

ip-address is the IP address of the name server. This is frequently the same as the DC that manages the domain, identified with forest-root, child-domain, or tree-domain.

Default(s)

None

Guidelines

Basic Domain Name Service (DNS) is defined in two RFCs, 1034 and 1035. The ARX also supports the Microsoft-specific authentication extensions specified in RFC 3645.

We recommend that every domain in the AD forest has at least one name server configured. (The show active-directory command shows all domains in an AD forest.) If a domain's DNS records do not correctly identify a front-end CIFS service, the domain's Windows clients cannot use Kerberos to authenticate with it.

You can use this command more than once to identify multiple name servers for the same AD domain. A CIFS service in that domain tries the first-configured service, then the second, and so forth until it successfully registers its IP address in an "A" record. After it succeeds, it stops registering; the external DNS servers are presumed to duplicate their "A" records independently, as configured by the DNS administrator.

From gbl-cifs mode, use dynamic-dns to DNS-register the front-end CIFS service. A CIFS service only registers with a name server when it is enabled (with enable (gbl-cifs, gbl-nfs)).

Samples

```
bstnA(gbl-forest[siera])# name-server siera.acopia.com 10.5.1.23
bstnA(gbl-forest[siera])# name-server siera.acopia.com 10.5.1.29
```

identifies two dynamic-DNS servers for the 'siera.acopia.com' domain.

```
bstnA(gbl-forest[siera])# no name-server siera.acopia.com 10.5.1.29
```

removes one name server from the 'siera.acopia.com' domain.

Related Commands

active-directory-forest
show active-directory
forest-root
child-domain
tree-domain
dynamic-dns

ntlm-auth-db

Purpose

This command is only necessary behind a cifs *service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the* domain-join *documentation, and avoid this CLI command.*

Use this command to configure an NTLM authentication database. The NTLM authentication database contains a list of valid usernames/passwords (up to 10). If a client is in the database, the ARX can use preset credentials to authenticate user access to back-end CIFS filers and shares.

Use the ntlm-auth-db command to instantiate a new NTLM authentication database.

Mode

gbl

Security Role(s)

crypto-officer

Syntax

`ntlm-auth-db` ***`name`***
`no ntlm-auth-db` ***`name`***

name (1-64 characters) is a name that you choose for the Windows NTLM authentication demo database.

Use the no form of the command to remove an NTLM authentication database.

Default(s)

None

Guidelines: Using Constrained Delegation Instead of a Secure Agent

As mentioned above, this command is unnecessary if all of your CIFS services use constrained delegation. Constrained delegation is implemented at a local DC. You can use the probe delegate-to command to retrieve some necessary configuration information to set up constrained delegation for your CIFS services.

For sites where some DCs are older than Windows Server 2003 (which introduced constrained delegation), either this command or ntlm-auth-server is required for NTLM support.

Guidelines: Running the Command

This command puts you into gbl-ntlm-auth-db mode, where you configure the NTLM demo database parameters. Use the user (gbl-ntlm-auth-db) command to configure up to 10 users/encrypted passwords associated with the NTLM authentication demo database. You need one valid username/password for each supported client (up to 10).

If you have more than one namespace that supports the same domain, you can assign the same database to both namespaces. Use the ntlm-auth-db (gbl-ns) command to assign an NTLM authentication database to a namespace.

If you remove an authentication database with the no form of the command, clients cannot access CIFS shares in a namespace that uses the database.

Samples

```
bstnA(gbl)# ntlm-auth-db MYDEMO
bstnA(gbl-ntlm-auth-db[MYDEMO])#
```

creates a new NTLM authentication database, “MYDEMO.”

```
bstnA(gbl)# no ntlm-auth-db OLDDEMO
```

removes an NTLM authentication database, “OLDDEMO,” from the switch.

Related Commands ntlm-auth-db (gbl-ns)
user (gbl-ntlm-auth-db)

ntlm-auth-server

Purpose

This command is only necessary behind a cifs *service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the* domain-join *documentation, and avoid this CLI command.*

The ARX Secure Agent (ASA) can be separately installed on a Windows domain controller (DC). If a front-end CIFS service uses unconstrained delegation, it requires the Secure Agent to assist with NTLM and/or NTLMv2 authentication for its clients. (This command has no effect on Kerberos authentication.) Use the ntlm-auth-server command to begin configuring a Secure Agent server at the ARX.

Use the no form of the command to remove a Secure-Agent-server configuration from the switch.

Mode

gbl

Security Role(s)

crypto-officer

Syntax

```
ntlm-auth-server name
no ntlm-auth-server name
```

name (1-128 characters) is a name that you choose for the Secure Agent server. Create a Secure Agent server for each DC that hosts a Secure Agent.

Default(s)

None

Guidelines

The CLI prompts for confirmation before creating a new NTLM-server configuration object; enter **yes** to continue.

◆ Note

*If the front-end CIFS Service uses **constrained delegation**, introduced with Windows Server 2003, the Secure Agent (and this command) is unnecessary. An administrator with "Domain Administrator" privileges can go to the DC and configure constrained delegation for this CIFS service. You can use the* probe delegate-to *command to find all back-end filers behind a CIFS service, which the DC administrator adds to the CIFS service's "delegate-to" list. Once the CIFS service is set up this way at the DC, the service's clients can authenticate with NTLM, NTLMv2, and/or Kerberos.*

For CIFS services with unconstrained delegation, one NTLM-server configuration object is required for each Secure Agent installed on each DC. Each DC should have Secure Agent Release 5.1.0 or later. Use the *ARX Secure Agent Installation Guide* for installation and/or upgrade instructions.

Guidelines (Cont.)

This command creates a new NTLM server object to represent a Secure Agent that is installed on a DC. The command places you in ntlm-auth-server mode, where you configure the Secure Agent server's attributes:

- Use the ip address (gbl-ntlm-auth-srv) command to enter the IP address of the DC running the Secure Agent.
- Use the port command to assign a port on which the Secure Agent listens for authentication requests from the switch.
- Use the windows-domain (gbl-ntlm-auth-srv) command to associate a domain name with this NTLM authentication server. This is a Windows domain that is managed by the server's DC host. Use the exact domain name that clients use for authentication.
- Use the password command to provide the password for accessing the Secure Agent authentication server.

After you configure the server's attributes, you assign it to one or more namespaces. Use the ntlm-auth-server (gbl-ns) command in gbl-ns mode to assign an NTLM-authentication server to a namespace. The Secure Agent assists the namespace with NTLM/NTLMv2 authentication for its CIFS clients. You can assign multiple NTLM-authentication servers to a single namespace, and you can assign a single NTLM-authentication server to multiple namespaces.

If you have redundant DCs for a given Windows Domain, we recommend installing a Secure Agent on all of them and using this command to add each of them to the ARX configuration. If you assign more than one to a namespace, the namespace software tries each of them until it successfully connects. The namespace software starts with the Secure Agent running the highest software revision.

If you remove a Secure Agent server with the no form of the command, NTLM/NTLMv2 clients cannot access CIFS shares in a namespace that uses this Secure Agent instance. Kerberos authentication is unaffected, though a fall-back to any version of NTLM fails if there is no Secure Agent for the requested domain.

The namespace's proxy-user (used for autonomous volume operations, such as import and file migrations) does not require an NTLM server to authenticate.

For more information about Secure Agent, see the *ARX Secure Agent Installation Guide*.

Samples

```
gffstnA(gbl)# ntlm-auth-server ggh-dc

This will create a new NTLM Server.

Create NTLM Server 'ggh-dc'? [yes/no] yes
gffstnA(gbl-ntlm-auth-srv[ggh-dc])#
```

instantiates a Secure Agent server, "server1."

```
bstnA(gbl)# no ntlm-auth-server dc1
```

removes the Secure Agent server, "dc1," from the ARX configuration.

Related Commands

show ntlm-auth-server
ip address (gbl-ntlm-auth-srv)
port
windows-domain (gbl-ntlm-auth-srv)
password

password

Purpose

This command is relevant to an NTLM Secure Agent, which is only necessary behind a cifs *service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the* domain-join *documentation, and avoid this CLI mode and command.*

At an installation that uses NTLM or NTLMv2 authentication with unconstrained delegation, the ARX uses the *Secure Agent* as an NTLM proxy. Use this command to enter a password to be shared by the Secure Agent and the ARX. The Secure Agent and the ARX use this password to verify each other's identity when they first establish a TCP connection.

Mode

gbl-ntlm-auth-srv

Security Role(s)

crypto-officer

Syntax

`password`

Default(s)

None

Guidelines: Using Constrained Delegation Instead of a Secure Agent

As mentioned above, Secure Agents (and this command) are unnecessary if all of your CIFS services use constrained delegation. Constrained delegation is implemented at a local DC. You can use the probe delegate-to command to retrieve some necessary configuration information to set up constrained delegation for your CIFS services.

For sites where some DCs are older than Windows Server 2003 (which introduced constrained delegation), this CLI mode and command are required for NTLM support.

Guidelines: Running the Command

The CLI prompts you for the password, then prompts again to re-enter the password. Enter 4-64 characters. The password does not appear in the show running/global-config output as plain text, but as a base64-encoded password.

The password you enter in this command must match the password configured at the ARX Secure Agent.

The switch uses Secure Agent software to assist in CIFS authentication requests received from clients to back-end CIFS storage devices. See the *ARX Secure Agent Installation Guide* for information about installing and managing Secure Agent software.

Sample

```
gffstnA(gbl)# ntlm-auth-server ggh-dc
gffstnA(gbl-ntlm-auth-srv[ggh-dc])# password
Password: ********
Validate Password: ********
gffstnA(gbl-ntlm-auth-srv[ggh-dc])#
```

Related Commands

ntlm-auth-server
show ntlm-auth-server

permit (gbl-mgmt-auth)

Purpose A Windows-management-authorization (*WMA*) group is a list of Windows clients who can use Windows-management applications (such as MMC or snapshots) in one or more namespaces. The permit command establishes a permissible management action for the current WMA group.

Use no permit to remove a permission from the group.

Mode gbl-mgmt-auth

Security Role(s) crypto-officer

Syntax

```
permit {share | session | open-file | snapshot | all} {monitor
| any}

no permit {share | session | open-file | snapshot | all}
```

share | session | open-file | snapshot | all chooses the object of this permission:

share selects CIFS-service shares, as well as non-shared volumes that are eligible to be added as CIFS shares. With **any** permissions (see below), members of this group can add and delete shares, too.

session represents client sessions that are currently connected to the CIFS service. A member client can potentially see all other client sessions if the group has the **monitor** permission (below). With the **any** permission, the members can also potentially disconnect other clients.

open-file chooses open files. If this has **monitor** permission, a group member can see these files; with **any** permission, a member can close files, too.

snapshot selects the directory for accessing a volume's snapshots. WMA group members, when they access the volume through CIFS, can only see the snapshot directory if you set the **monitor** permission for this. The **any** permission has the same effect.

all selects all of the above.

monitor | any shows what the member clients can do with the object(s):

monitor allows read-only access. This limits group members to listing (not changing) a CIFS service's shares, sessions, open files and/or snapshots.

any allows read-write access, so that member clients add, delete, close, and/or shut down the objects as appropriate. This has no additional affect on snapshots, which can only be monitored.

Default(s)

```
no permit share
no permit session
no permit open-file
```

Guidelines

Depending on the management permissions for the group, its members may be able to add new shares, view open files, and/or disconnect other client sessions. Use this command to set the specific permissions for this group; invoke the command multiple times to establish different permission levels for each object. Use the user (gbl-mgmt-auth) command to add Windows clients to the group.

Use the show windows-mgmt-auth command to display the configuration for a WMA group.

To apply a WMA group to a namespace, use windows-mgmt-auth (gbl-ns). You can apply multiple groups to the same namespace. Any CIFS service in front of this namespace uses these WMA groups if it has remote management/browsing enabled. From gbl-cifs mode, use browsing to enable remote management/browsing for a CIFS service.

The WMA group's members only see the results of this command if they connect *after* you invoke it.

Guidelines: Permissions Required for MMC Option to Delete Shares

MMC software finds all session connections to a share before it deletes the share. The MMC "delete share" operation therefore fails with an error if it does not have permission to find the share's session connections. Use this extra command to allow MMC to find session connections:

```
permit session monitor
```

(The **permit session all** command also works for MMC.)

This is in addition to the following command, which permits the client to remove the share:

```
permit share all
```

Samples

```
bstnA(gbl-mgmt-auth[fullAccess])# permit all any
```

allows the "fullAccess" group to perform any management actions on shares, open files, and client sessions.

```
bstnA(gbl-mgmt-auth[readOnly])# permit share monitor
bstnA(gbl-mgmt-auth[readOnly])# permit session monitor
```

permits members of the "readOnly" group to see all shares and client sessions, but not to change them in any way.

```
bstnA(gbl-mgmt-auth[snapViewers])# permit snapshot monitor
```

permits members of the "snapViewers" group to see the snapshot directory through CIFS.

```
bstnA(gbl-mgmt-auth[fullAccess])# no permit open-file
```

disqualifies the "fullAccess" group from viewing or changing open files though MMC or similar applications.

Related Commands

windows-mgmt-auth
user (gbl-mgmt-auth)
show windows-mgmt-auth
windows-mgmt-auth (gbl-ns)
browsing

port

Purpose

This command is relevant to an NTLM Secure Agent, which is only necessary behind a cifs *service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the* domain-join *documentation, and avoid this CLI mode and command.*

At an installation that uses NTLM or NTLMv2 authentication with unconstrained delegation, the ARX uses the *Secure Agent* as an NTLM proxy. Use this command to enter the TCP port (number) on which Secure Agent listens for authentication requests from the switch.

Use the no form of the command to remove the current port number.

Mode

gbl-ntlm-auth-srv

Security Role(s)

crypto-officer

Syntax

`port` ***`port`***
`no port`

port (optional; 1-65535) is the port number on which Secure Agent listens for authentication requests from the switch.

Default(s)

25805

Guidelines: Using Constrained Delegation Instead of a Secure Agent

As mentioned above, Secure Agents (and this command) are unnecessary if all of your CIFS services use constrained delegation. Constrained delegation is implemented at a local DC. You can use the probe delegate-to command to retrieve some necessary configuration information to set up constrained delegation for your CIFS services.

For sites where some DCs are older than Windows Server 2003 (which introduced constrained delegation), this CLI mode and command are required for NTLM support.

Guidelines: Running the Command

The Secure Agent server uses the specified port for NTLM-authentication transactions with the switch. The port number must match the port number specified when the Secure Agent software was installed on a Windows Domain Controller. If you do not enter a port number, the default port is used. See the *ARX Secure Agent Installation Guide* for more information.

Sample

```
gffstnA(gbl)# ntlm-auth-server ggh-dc
gffstnA(gbl-ntlm-auth-srv[ggh-dc])# port 25805
```

sets the Secure Agent server's port number to 25805.

Related Commands

ntlm-auth-server
ip address (gbl-ntlm-auth-srv)

show ntlm-auth-db

Purpose

This command is relevant to an NTLM authentication database, which is only necessary behind a cifs service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the domain-join documentation.

Use this command to display the configured NTLM-authentication demo databases.

Mode

(any)

Security Role(s)

crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

`show ntlm-auth-db [`*`name`*`]`

name (1-64 characters) is the name of an NTLM authentication demo database.

If no name is specified, this command shows names for all configured NTLM-authentication demo databases.

Guidelines

With the name argument, the command shows the username(s) and namespace(s) that use this database.

Use the user (gbl-ntlm-auth-db) command to add a user to the database.

Use the ntlm-auth-db (gbl-ns) command to assign the database to a namespace.

Sample

```
bstnA> show ntlm-auth-db

Name
--------------------------------------
ntlmMap2
ntlmMap3

bstnA> show ntlm-auth-db ntlmMap2

Name
------------------------------------------------------
ntlmMap2

   Users
   ------------------------------------------------------
   techspecialist


   Mapped to the Following Namespaces
   ------------------------------------------------------
cifsNamespace1
```

Related Commands

ntlm-auth-db
ntlm-auth-db (gbl-ns)
user (gbl-ntlm-auth-db)

show ntlm-auth-server

Purpose

This command is relevant to an NTLM authentication server, which is only necessary behind a cifs *service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the* domain-join *documentation.*

At an installation that does not support constrained delegation, the ARX uses the *Secure Agent* as an NTLM/NTLMv2 proxy. Use the show ntlm-auth-server command to display the configuration and statistics from one or more configured Secure Agent servers.

Mode

(any)

Security Role(s)

crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show ntlm-auth-server [name] [detailed]
```

name (optional, 1-128 characters) is the name of a Secure Agent server instance. If no name is specified, this command shows statistics for all configured Secure Agent servers.

detailed (optional) generates detailed output, with success and failure counts from all current connections to the Secure Agent(s).

Guidelines: Using Constrained Delegation Instead of a Secure Agent

As mentioned above, Secure Agents (and this command) are unnecessary if all of your CIFS services use constrained delegation. Constrained delegation is implemented at a local DC. You can use the probe delegate-to command to retrieve some necessary configuration information to set up constrained delegation for your CIFS services.

For sites that can follow this recommendation, you can use show statistics domain-controller to retrieve NTLM-authentication statistics.

For sites where some DCs are older than Windows Server 2003 (which introduced constrained delegation), a Secure Agent is required for NTLM support and this command is relevant.

Guidelines: Summary Output

For the summary output, this command includes the following categories:

Name is the name of this Secure Agent server instance, as configured with the ntlm-auth-server command.

Version identifies the Software Release of the Secure Agent software.

Server is the IP address of the Secure Agent server host system. If this is incorrect, you can use the ip address (gbl-ntlm-auth-srv) command to change it.

Domain Name is the associated Windows Domain. If this is incorrect, you can use the windows-domain (gbl-ntlm-auth-srv) command to correct it.

Pre-Win2k Domain is the "short" name of the above domain, for domains that existed before the release of Windows 2000. As above, you can use the windows-domain (gbl-ntlm-auth-srv) command to correct this in the ARX database.

Capabilities is "NTLM," "NTLM, NTLMv2, Session Key," or "Unknown." The "Session Key" flag indicates that the Secure Agent can support SMB signing (see cifs filer-signatures and signatures).

Status is one of the following:

- Reachable is the status for a healthy Secure Agent.
- Unreachable indicates a connection problem. Check the connection to the Secure Agent with show exports, expect traceroute, and/or ping.
- Password Mismatch means that the password configured at the ARX (with the password command) is different from the password configured in the Secure Agent's management applet. The *ARX Secure Agent Installation Guide* explains how to reset the password at the Secure Agent's management applet.
- If another error appears, contact F5 Support for analysis.

Guidelines: Detailed Output

The detailed keyword causes the command to download detailed statistics from the Secure Agent(s). These are high-level statistics followed by a table of counters for every current connection to a Secure Agent. If you selected all Secure Agents, a separate SECURE AGENT STATISTICS page appears for each of them. The high-level statistics include the following:

Agent IP is the IP address of the Secure Agent's host DC.

Agent Port is the TCP port where the Secure Agent is listening for ARX queries.

Uptime shows the amount of time since the last restart of the Secure Agent application.

Current Connections is the number of ARX systems currently connected to the Secure Agent.

Failed connection attempts and

Successful connection attempts summarize the connection statistics.

Account Scan Interval is the time between scans of the DC's SAM database, in seconds. The Secure Agent uses this scanned data to assist the ARX with its NTLM/NTLMv2 authentications.

Software Version is the release number for the Secure Agent software.

Capabilities is "NTLM," "NTLM, NTLMv2, Session Key," or "Unknown.", as above.

Connection #1 lists specific details about the first connection to the server including the source IP address, bytes transmitted, successful or failed client authentications, and filer response generation statistics. The "client" in these statistics is an ARX (typically, the current ARX), and the Source IP in the output is the ip proxy-address that the ARX used to connect to the Secure Agent.

Connection #2, and any additional connections shown, list details specific to that connection to the server.

Samples

`gffstnA> show ntlm-auth-server`

lists all configured Secure Agent servers. See Figure 18.2 on page 18-32 for sample output.

`gffstnA> show ntlm-auth-server ggh-dc`

shows high-levels statistics about one Secure Agent server, 'ggh-dc.' See Figure 18.3 on page 18-32 for sample output.

`gffstnA> show ntlm-auth-server detailed`

shows detailed statistics about all Secure Agent servers. See Figure 18.4 on page 18-32 for sample output.

Related Commands

ntlm-auth-server
ntlm-auth-server (gbl-ns)

Figure 18.2 *Sample Output: show ntlm-auth-server*

```
gffstnA> show ntlm-auth-server

   Name : ggh-dc
       Version : V5.01.005
       Server : 192.168.25.109
       Windows Domain : MEDARCH.ORG
       Pre-Win2k Name :
       Capabilities : NTLM, NTLMv2, Session Key
       Status : Reachable
```

Figure 18.3 *Sample Output: show ntlm-auth-server ggh-dc*

```
gffstnA> show ntlm-auth-server ggh-dc

******************** SECURE AGENT STATISTICS ******************
Agent IP   : 192.168.25.109
Agent Port : 25805

Uptime: 136 days, 9 hours, 9 minutes and 17 seconds
Current Connections: 18
Failed connection attempts: 1358881
Successful connection attempts: 1368652
Account Scan Interval: 300
Software version: Version 5.01.000.11899 (Sep 24 2009 20:55:35) [jc] [x86]
Capabilities: NTLM, NTLMv2, Session Key
```

Figure 18.4 *Sample Output: show ntlm-auth-server detailed*

```
gffstnA> show ntlm-auth-server detailed

******************** SECURE AGENT STATISTICS ******************
Agent IP   : 192.168.25.109
Agent Port : 25805

Uptime: 136 days, 9 hours, 9 minutes and 16 seconds
Current Connections: 18
Failed connection attempts: 1358881
Successful connection attempts: 1368651
Account Scan Interval: 300
Software version: Version 5.01.000.11899 (Sep 24 2009 20:55:35) [jc] [x86]
Capabilities: NTLM, NTLMv2, Session Key
Total authentication statistics since last reset:
  Successful authentications: 3573266
  Failed authentications: 177
  Successful filer requests: 3159
  Failed filer requests: 0

Connection #1
  Source IP: 10.46.13.210
  Duration: 814 (seconds)
  bytes Received: 168
  bytes Transmitted: 624
  Successful Client NTLM Authentications: 12
  Successful Client NTLMv2 Authentications: 0
  Failed Client Authentications:
    No Such User     : 0
    Bad Password     : 0
    Locked out       : 0
    Account Disabled : 0
    Account Expired  : 0
    Password Expired : 0
    Time Restricted  : 0
```

```
    API Error        : 0
  Filer Response Generation:
    NTLM Success count: 0
    NTLM Failure count: 0
    NTLMv2 Success count: 0
    NTLMv2 Failure count: 0

Connection #2
  Source IP: 10.54.143.170
  Duration: 26047 (seconds)
  bytes Received: 60
  bytes Transmitted: 276
  Successful Client NTLM Authentications: 1
  Successful Client NTLMv2 Authentications: 0
  Failed Client Authentications:
    No Such User     : 0
    Bad Password     : 0
    Locked out       : 0
    Account Disabled : 0
    Account Expired  : 0
    Password Expired : 0
    Time Restricted  : 0
    API Error        : 0
  Filer Response Generation:
    NTLM Success count: 0
    NTLM Failure count: 0
    NTLMv2 Success count: 2
    NTLMv2 Failure count: 0

Connection #3
  Source IP: 10.46.39.202
...

******************** SECURE AGENT STATISTICS ******************
Agent IP   : 192.168.25.102
Agent Port : 25805

Uptime: 152 days, 12 hours, 9 minutes and 46 seconds
Current Connections: 26
Failed connection attempts: 2410143
Successful connection attempts: 2409825
Account Scan Interval: 300
Software version: Version 5.00.000.11658 (Apr 22 2009 23:01:57) [jc] [x86]
Capabilities: NTLM

Connection #1
  Source IP: 10.54.74.253
  Duration: 3918 (seconds)
  bytes Received: 0
  bytes Transmitted: 0
...
```

show windows-mgmt-auth

Purpose A Windows-management-authorization (*WMA*) group is a list of Windows clients with permissions to use Windows-management applications (such as MMC). Use this command to show the members and permissions for this type of group.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show windows-mgmt-auth [`*`name`*`]`**

name (optional; 1-64 characters) is the name of a specific WMA group.

If no name is specified, this command shows all WMA groups.

Guidelines The output for this command shows the following information for each WMA group:

Windows Authorization Policy is the name of the WMA group, specified in the windows-mgmt-auth command.

The user list appears next, as a table with one row per user:

User Name and

Domain Name describe one client in the group. Use user (gbl-mgmt-auth) to add or remove a client. If there is a pre-Windows 2000 name for the domain, it appears in the Domain Name field in parentheses.

The last table shows the management permissions for the group, one permission per row. Each row has two fields:

Managed Object is Share, Session, Open-file, Snapshot, or All. The permission (below) applies to this is the type of object.

Permitted Operation is "Monitor" or "Any." "Monitor" means that WMA-group clients can view the Managed Objects (shares, client sessions, open-files, or snapshots). "Any" means that the clients can change them, too (except snapshots, which are always read-only). Use permit (gbl-mgmt-auth) to add or remove a permission.

Samples `bstnA> show windows-mgmt-auth`

shows all WMA groups on this system. See Figure 18.5 for sample output.

`bstnA> show windows-mgmt-auth fullAccess`

shows one WMA group, "fullAccess." See Figure 18.6 on page 18-35 for sample output.

Related Commands windows-mgmt-auth
user (gbl-mgmt-auth)
permit (gbl-mgmt-auth)

Figure 18.5 *Sample Output: show windows-mgmt-auth*

```
bstnA> show windows-mgmt-auth

Windows Authorization Policy: fullAccess
```

```
User Name                                    Domain Name
-------------------------------------------- ------------------------------
juser                                        MEDARCH.ORG
jquser                                       MEDARCH.ORG

Managed Object         Permitted Operation
---------------------- ----------------------
All                    Any


Windows Authorization Policy: readOnly

User Name                                    Domain Name
-------------------------------------------- ------------------------------
mhoward_md                                   MEDARCH.ORG
zmarx_cpa                                    MEDARCH.ORG
lfine_md                                     MEDARCH.ORG
choward_md                                   MEDARCH.ORG

Managed Object         Permitted Operation
---------------------- ----------------------
Share                  Monitor
Session                Monitor
Snapshot               Monitor


Windows Authorization Policy: snapViewers

User Name                                    Domain Name
-------------------------------------------- ------------------------------
juser                                        MEDARCH.ORG
jquser                                       MEDARCH.ORG

Managed Object         Permitted Operation
---------------------- ----------------------
Snapshot               Monitor
```

Figure 18.6 *Sample Output: show windows-mgmt-auth fullAccess*

```
bstnA> show windows-mgmt-auth fullAccess

Windows Authorization Policy: fullAccess

User Name                                    Domain Name
-------------------------------------------- ------------------------------
juser                                        MEDARCH.ORG
jquser                                       MEDARCH.ORG

Managed Object         Permitted Operation
---------------------- ----------------------
All                    Any
```

tree-domain

Purpose

Some Active-Directory (AD) domains are outside the domain namespace, but have a trust relationship with domains in the forest. These are called *tree domains*. Use the tree-domain command to add a tree domain to the AD forest.

Use the no form of the command to remove a tree domain, or one of the redundant DCs that manages it.

Mode

gbl-forest

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
tree-domain domain-name dc [preferred]
no tree-domain domain-name dc
```

domain-name (1-256 characters) is the tree domain name. This does not require any parent in the forest.

dc is the IP address (for example, 192.168.25.56) of the domain controller (DC) for the tree domain. If redundant DCs are hosting the domain, repeat this command (using the same ***domain-name***) with the address of the redundant ***dc***.

preferred (optional) declares that you prefer to use this DC over others in the same domain. Whenever the ARX Kerberos process chooses from multiple DCs in the same domain, it chooses from the DCs labeled "preferred" first.

Default(s)

None.

Guidelines

This command may not be necessary if you use active-directory update seed-domain to automatically discover all child domains in the AD forest. Before you use this command, use the above command (or forest-root) to add a root to the Active Directory forest.

The tree-domain command adds a disjoint tree to the forest, outside of the forest's domain namespace. For example, if the forest has a root domain of "telco.com" with two child domains, "na.telco.com" and "asia.telco.com," you could establish a tree domain outside the namespace, "myphone.org." The tree domain is presumed to have a two-way-transitive-trust relationship with one of the domains in the root forest.

If redundant DCs manage the tree domain, re-use the tree-domain command (with the same Domain name) for each redundant DC. If you prefer to use one or more DCs in the domain over the others, use the preferred flag to identify those DCs. The Kerberos process uses the DC(s) from the preferred list as its active DC(s). The process only stops using a DC if its connection times out; the kerberos health-check threshold command selects the timeout threshold. The Kerberos process does not choose any DCs outside the preferred list unless all of the preferred DCs have timed out. You can use show active-directory status to see which DCs are currently active for each domain.

To remove a DC from the preferred list, you can re-run the tree-domain command without the optional preferred flag. You can use this flag to change the active DC at any time.

Guidelines (Cont.)

You can use the child-domain command to add a child to this tree domain. The parent/child relationship is created through the domain names ("aparent.com" can be the parent to "achild.aparent.com").

If the tree is managed by redundant DCs, you can use no tree-domain to remove one of them from the configuration. Otherwise, the "no" command removes the tree domain entirely; you must remove all of the tree's child domains before you can do this.

Samples

```
bstnA(gbl-forest[medarch])# tree-domain ne.medarch.org 172.16.124.73
preferred
```

adds the 'ne.medarch.org' tree domain to a forest. This also declares this DC to be preferred over others in the same domain.

```
bstnA(gbl-forest[medarch])# tree-domain ne.medarch.org 172.16.124.73
```

has the same effect as the previous example, but removes the DC from the "preferred" list.

Related Commands

active-directory-forest
active-directory update seed-domain
kerberos health-check threshold
child-domain
domain-join
forest-root
show virtual service

user (gbl-mgmt-auth)

Purpose

A Windows-management-authorization (*WMA*) group is a list of Windows clients with permissions to use remote Windows-management applications, such as MMC. The user command adds a client to the current WMA group.

Use no user to remove a user from the group.

Mode

gbl-mgmt-auth

Security Role(s)

crypto-officer

Syntax

`user` *`name`* `windows-domain` *`fqdn`* `[pre-win2k-name` *`old-style-domain`*`]`
`no user` *`name`* `windows-domain` *`domain-name`*

name (1-64 characters) is a valid Windows user.

domain-name (1-64 characters) is the user's domain name, exactly as the client types it for authentications. We recommend using the fully-qualified domain name (FQDN) for this field. If the client uses the FQDN (such as "mygroup.org") to authenticate, they can use Kerberos; if the client uses the older NetBIOS name from the first part of the FQDN (such as "MYGROUP"), the client can only use some form of NTLM.

old-style-domain (optional, 1-15 bytes, no periods) is the legacy, NetBIOS name of the Windows domain, if it is different from the default (below). The default should suffice for most situations; this option should only rarely be necessary. Clients who authenticate with this pre-Windows 2000 name can also use MMC; for example, if the pre-Windows 2000 name is "GROUP" for the "mygroup.org" domain, an MMC client could authenticate as "GROUP/juser" or "mygroup.org/juser."

Default(s)

old-style-domain - The old-style name discovered with active-directory update seed-domain. If the old-style name was never discovered for this domain, the ARX uses the first part (before the first "."), up to 15 characters) of the FQDN in the ***domain-name***.

Guidelines

Depending on the management permissions for the group, they may be able to add new shares, view open files, and/or disconnect other client sessions. Use the permit (gbl-mgmt-auth) command to set the specific permissions for this group. Use this command multiple times to add multiple users to the group.

Use the show windows-mgmt-auth command to display the configuration for a WMA group.

To apply a WMA group to a namespace, use windows-mgmt-auth (gbl-ns). You can apply multiple groups to the same namespace. Use browsing to enable Windows management for a CIFS service.

Samples

```
bstnA(gbl-mgmt-auth[fullAccess])# user jquser windows-domain
MEDARCH.ORG
```

adds a user to the WMA group, "fullAccess."

```
bstnA(gbl-mgmt-auth[readOnly])# no user test5 windows-domain
MEDARCH.ORG
```

removes a user from the "readOnly" group.

Related Commands

windows-mgmt-auth
permit (gbl-mgmt-auth)
show windows-mgmt-auth
windows-mgmt-auth (gbl-ns)
browsing

user (gbl-ntlm-auth-db)

Purpose

This command is relevant to an NTLM authentication database, which is only necessary behind a cifs *service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the* domain-join *documentation, and avoid this CLI mode and command.*

Use this command to enter up to 10 usernames and passwords associated with the NTLM authentication database. If a client entry is in the database, the ARX can use the preset credentials to authenticate connections for access to back-end CIFS filers.

Mode

gbl-ntlm-auth-db

Security Role(s)

crypto-officer

Syntax

`user` ***`dbuser`***
`no user` ***`dbuser`***

dbuser (1-64 characters) is one authorized Windows user.

Default(s)

None

Guidelines

The CLI prompts for the user's password twice. You need one valid username and password for each supported client (up to 10). Use the no form of the command to remove a user from the current database.

Samples

```
bstnA(gbl-ntlm-auth-db[MYDEMO])# user sales
Password: ********
Validate Password: ********
```

creates the user, "sales," associated with the MYDEMO database

```
bstnA(gbl-ntlm-auth-db[MYDEMO])# no user sales
```

removes the user, "sales" associated with the "MYDEMO" database.

Related Commands

ntlm-auth-db

windows-domain (gbl-ntlm-auth-srv)

Purpose

This command is relevant to an NTLM Secure Agent, which is only necessary behind a cifs *service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the* domain-join *documentation, and avoid this CLI mode and command.*

At an installation that uses NTLM or NTLMv2 with unconstrained delegation, the ARX uses the *Secure Agent* as an NTLM proxy. Use this command to assign the domain name to a Secure Agent server.

Use the no form of the command to remove the domain name for this Secure Agent, thereby disabling it.

Mode

gbl-ntlm-auth-srv

Security Role(s)

crypto-officer

Syntax

```
windows-domain domain-name [pre-win2k-name old-style-domain]
no windows-domain
```

domain-name (1-64 characters) is the Windows Domain to which you assign a Secure Agent instance (authentication server). When a client authenticates to an ARX CIFS service, the client uses this Secure Agent if they use this domain name.

old-style-domain (optional, 1-15 characters, no periods) is the name of the legacy Windows domain, if it is different from the default (below). The default should suffice for most situations; this option should only rarely be necessary. Clients who authenticate with the pre-Windows 2000 name can also use this Secure Agent.

Default(s)

old-style-domain - The old-style name discovered with active-directory update seed-domain. If the old-style name was never discovered for this domain, the ARX assumes that the old-style name is the first part (before the first "."), up to 15 characters) of the FQDN in the ***domain-name***.

Guidelines: Using Constrained Delegation Instead of a Secure Agent

As mentioned above, Secure Agents (and this command) are unnecessary if all of your CIFS services use constrained delegation. Constrained delegation is implemented at a local DC. You can use the probe delegate-to command to retrieve some necessary configuration information to set up constrained delegation for your CIFS services.

For sites where some DCs are older than Windows Server 2003 (which introduced constrained delegation), this CLI mode and command are required for NTLM support.

Guidelines: Running the Command

The domain name is the Windows domain where the Secure Agent software is installed (for each domain controller that hosts Secure Agent).

Use the longest domain name that clients use for their authentications. For example, if clients authenticate with "porthos.musketeer.gov" domain as well as the shorter "porthos" name, you should create an authentication server for "porthos.musketeer.gov." The ARX uses the single authentication-server object for the domain name and all shorter variations of the same domain name.

For domains with a completely-different pre-Win2k name, such as "GUARD3" as an the early name for "porthos," use the pre-win2k-name option. The separate authentication server should use the same ip address (gbl-ntlm-auth-srv) to point to the same DC.

If you remove the domain name with the no form of this command, clients cannot access CIFS shares in a namespace that uses this instance of Secure Agent.

Sample

```
gffstnA(gbl)# ntlm-auth-server serverFQDN
gffstnA(gbl-ntlm-auth-srv[serverFQDN])# windows-domain ggh.medarch.org
```

assigns a Secure Agent instance to the "ggh.medarch.org" Windows Domain, and to the old-style "NHCIFS" domain name.

Related Commands

ntlm-auth-server
ip address (gbl-ntlm-auth-srv)
show proxy-user

windows-mgmt-auth

Purpose

A Windows-management-authorization (*WMA*) group is a list of Windows clients with permissions to use Windows-management applications (such as MMC) in one or more CIFS namespaces. Use the windows-mgmt-auth command to create a WMA group.

Use no windows-mgmt-auth to remove a WMA group.

Mode

gbl

Security Role(s)

crypto-officer

Syntax

`windows-mgmt-auth` ***`name`***
`no windows-mgmt-auth` ***`name`***

name (1-64 characters) is a name you choose for the group.

Default(s)

None

Guidelines

This command places you in gbl-mgmt-auth mode, where you then add one or more Windows clients and the management permissions for all of them. Use the user (gbl-mgmt-auth) command to add a Windows user to the group. Use the permit (gbl-mgmt-auth) command to add management permissions for these users, such as permission to add CIFS shares, close CIFS files, or view snapshot directories from Windows.

Any CIFS client can list all configured CIFS shares from a CIFS service, whether or not he or she is a member of a WMA group. WMA-group members, given the proper permissions, can also see all volumes behind a CIFS service; this provides a menu for adding new shares to the service.

Use the show windows-mgmt-auth command to display the configuration for a WMA group.

To apply a WMA group to a namespace, use windows-mgmt-auth (gbl-ns). You can apply multiple groups to the same namespace. Use browsing to enable Windows management for a CIFS service.

Samples

```
bstnA(gbl)# windows-mgmt-auth fullAccess
bstnA(gbl-mgmt-auth[fullAccess])#
```

starts the configuration of WMA group, “fullAccess.”

```
bstnA(gbl)# no windows-mgmt-auth testMMC
```

removes a WMA group from the system.

Related Commands

user (gbl-mgmt-auth)
permit (gbl-mgmt-auth)
show windows-mgmt-auth
windows-mgmt-auth (gbl-ns)
browsing

19

NFS Access Lists

anonymous-gid

Purpose

When an NFS client accesses a share as *root* (a UNIX superuser), an NFS-access list typically re-maps the user's identity to that of the *anonymous* user, which has very low access rights. This security feature is called *root squashing*. Use the anonymous-gid command to change the Group ID (GID) for the *anonymous* user.

Use the no form of the command to revert to the default GID for *anonymous*.

Mode

gbl-nfs-acl

Security Level

crypto-officer

Syntax

`anonymous-gid` ***`id`***
`no anonymous-gid`

id (1-65535) is a group ID number that the access list uses when it squashes root access.

Default(s)

65534

Guidelines

When permit rules have root-squash enabled, they translate the Group ID (GID) of a *root* user to an *anonymous* GID. To set the user ID (UID), use anonymous-uid. Use show nfs-access-list to see the current GID and UID used for *anonymous*.

Sample

```
bstnA(gbl)# nfs-access-list eastcoast
bstnA(gbl-nfs-acl[eastcoast])# anonymous-gid 100
bstnA(gbl-nfs-acl[eastcoast])# ...
```

sets the *anonymous* GID to 100.

Related Commands

nfs-access-list
show nfs-access-list
anonymous-uid

anonymous-uid

Purpose When an NFS client accesses a share as *root* (a UNIX superuser), an NFS-access list typically re-maps the user's identity to that of the *anonymous* user, which has very low access rights. This security feature is called *root squashing*. Use the anonymous-uid command to change the User ID (UID) for the *anonymous* user.

Use the no form of the command to revert to the default UID for *anonymous*.

Mode gbl-nfs-acl

Security Level crypto-officer

Syntax

```
anonymous-uid id
no anonymous-uid
```

id (1-65535) is a User ID number that the access list uses when it squashes root access.

Default(s) 65534

Guidelines When permit rules have root-squash enabled, they translate the User ID (UID) of a *root* user to an *anonymous* UID. To set the group ID (GID), use anonymous-gid. Use show nfs-access-list to see the current GID and UID used for *anonymous*.

Sample

```
bstnA(gbl)# nfs-access-list eastcoast
bstnA(gbl-nfs-acl[eastcoast])# anonymous-uid 100
bstnA(gbl-nfs-acl[eastcoast])# ...
```

sets the *anonymous* UID to 100.

Related Commands

anonymous-gid
nfs-access-list
permit (gbl-nfs-acl)

deny

Purpose

Use this command to deny NFS access for one subnet.

Use the no form of the command to remove a deny rule from the current access list.

Mode

gbl-nfs-acl

Security Level

crypto-officer

Syntax

```
deny ip-address mask
no deny ip-address mask
```

ip-address (0.0.0.0-255.255.255.255) is the address of the subnet to be denied access.

mask (0.0.0.0-255.255.255.255) is the netmask (network part of the ip address).

Default(s)

None

Guidelines

You may have a situation where most of a large subnet should be permitted access to NFS, but some portions of the subnet should be denied access. From gbl-nfs-acl mode, use this command to add a deny rule for one subnet. Use the permit command to allow access for a subnet. Order is important; if a client matches both rules, only the first one is enforced.

No permit or deny rule is enforced until you exit gbl-nfs-acl mode. This ensures that all permit and deny rules are enforced at once, after you finish setting them up in the desired order. Use the exit or end command to exit the mode.

Samples

```
bstnA(gbl)# nfs-access-list eastcoast
bstnA(gbl-nfs-acl[eastcoast])# deny 192.168.77.0 255.255.255.0
bstnA(gbl-nfs-acl[eastcoast])# deny 192.168.202.0 255.255.255.0
bstnA(gbl-nfs-acl[eastcoast])# permit 192.168.0.0 255.255.0.0
read-write
bstnA(gbl-nfs-acl[eastcoast])# ...
```

denies access to two Class C subnets, but then permits access to any other IP *inside* their Class B supernet, 192.168.0.0/16.

Related Commands

nfs-access-list
permit (gbl-nfs-acl)

description (gbl-nfs-acl)

Purpose

Use the optional description command to include a descriptive string for an access list.

Use the no form of the command to remove the description from the current access list.

Mode

gbl-nfs-acl

Security Level

crypto-officer

Syntax

`description` *`text`*
`no description`

text (1-255 characters) is your description for the current access list. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s)

no description

Guidelines

The description appears in the detailed output for show nfs-access-list.

Sample

```
bstnA(gbl)# nfs-access-list eastcoast
bstnA(gbl-nfs-acl[eastcoast])# description “allowable subnets in MA,
NY, & DC”
bstnA(gbl-nfs-acl[eastcoast])# ...
```

adds a description for the current NFS access list.

Related Commands

nfs-access-list
show nfs-access-list

ip address (gbl-nis-dom)

Purpose Use the ip address command to identify one NIS server for the current NIS domain.

Use the no form of the command to remove one NIS server from the list.

Mode gbl-nis-dom

Security Role(s) storage-engineer or crypto-officer

Syntax

```
ip address address
no ip address address
```

address identifies the NIS server (for example, 192.168.70.128). This address must be on a server/proxy-IP subnet (see ip proxy-address) or reachable through a gateway on that subnet (via static route: see ip route to create a static route).

Default(s) None

Guidelines You can enter this command multiple times to identify up to four servers for the same NIS domain. The switch uses the first server for all of its netgroup lookups, falling back on the rest in case the first server fails or times out. The servers are used in the order that they are defined; if the first server is unreachable, the switch tries the second server, and so on. Use show nis domain to see the current order of NIS servers.

Samples

```
bstnA(gbl-nis-dom[mydom.org])# ip address 192.168.78.55
```

uses a NIS server for "mydom.org," at 192.168.78.55.

```
bstnA(gbl-nis-dom[mydom.org])# no ip address 10.10.25.1
```

removes one NIS server from the list.

Related Commands

nis domain
show nis domain
ip proxy-address
ip route

nfs-access-list

Purpose Use this command to create an access list (optional) for NFS services. An *access list* is a list of IP hosts to which you permit or deny access to the NFS service. You can use an NFS access list (or create additional ones) for any number of NFS services.

Use the no form of the command to remove an access list.

If you *do not* include an access list for an exported volume, *all* clients can access the share. See export (gbl-nfs) for more information.

Mode gbl

Security Level crypto-officer

Syntax

```
nfs-access-list list-name
no nfs-access-list list-name
```

list-name (1-64 characters) is a name you choose for the access list.

Default(s) None

Guidelines The CLI prompts for confirmation before creating a new NFS access list; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new objects.)

When you configure an NFS service, you can optionally apply one NFS access list to the service. This command places you in gbl-nfs-acl mode, from which you can configure various permit and deny rules for specific subnets and/or NIS netgroups. For example, you could permit access from the subnet at 192.168.101.0 but deny access from all other subnets. By default, all subnets are denied any access. Use the permit (gbl-nfs-acl) and deny commands to add permit and deny access rules, respectively. To use NIS netgroups from your back-end servers, use nis domain to create a NIS domain on the switch, use nis domain (gbl-nfs-acl) to apply it to the access list, and use permit netgroup to allow access to the hosts in one netgroup.

When you use the no form of the command to remove an access list, you must first remove all references to the access list before you can remove the list itself.

Samples

```
bstnA(gbl)# nfs-access-list eastcoast

This will create a new NFS ACL.

Create NFS ACL ''eastcoast''? [yes/no] yes
bstnA(gbl-nfs-acl[eastcoast])# ...
```

creates a new access list and places you gbl-nfs-acl mode.

```
bstnA(gbl)# no nfs-access-list testacl
```

removes an NFS access list from the switch configuration.

Related Commands

anonymous-gid
anonymous-uid
nis domain (gbl-nfs-acl)
deny
permit (gbl-nfs-acl)
description (gbl-nfs-acl)
show nfs-access-list
export (gbl-nfs)

nis domain

Purpose

Use the nis domain command to identify a Network Information System (NIS) domain to be used in one or more NFS access lists.

Use no nis domain to remove the NIS-domain configuration from the ARX.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

`nis domain` ***`domain`***
`no nis domain` ***`domain`***

domain (1-256 characters) is the NIS-domain name (for example, “company” or “company.com”).

Default(s)

None

Guidelines

This command places you into gbl-nis-dom mode, where you use the ip address (gbl-nis-dom) command to identify at least one NIS server for the domain. After you specify one or more NIS servers, the ARX looks up all of the netgroups in the domain, then performs DNS lookups for all hostnames in those netgroups. The results are cached on the switch, to prevent excessive traffic between the switch and the DNS server; the show nis netgroup command shows the contents of this cache. Use nis update to refresh this cache by performing all of the necessary lookups.

The switch ignores all users and groups in these netgroups. As a proxy, the ARX does not authenticate specific users and groups; the back-end filers perform user/group authentications, and the switch passes the results back to the client.

You can use nis domain (gbl-nfs-acl) to use this NIS domain in an access list. The permit netgroup command (see permit (gbl-nfs-acl)) adds a permit rule for one netgroup. Use show nis domain to view all NIS domains and their configured NIS servers.

Samples

```
bstnA(gbl)# nis domain wwmed.com
bstnA(gbl-nis-dom[wwmed.com])# ...
```

creates a new NIS domain, “wwmed.com.”

```
bstnA(gbl)# no nis domain testnis
```

removes a NIS domain.

Related Commands

ip address (gbl-nis-dom)
nis update
show nis netgroup
show nis domain
nis domain (gbl-nfs-acl)
permit (gbl-nfs-acl)

nis domain (gbl-nfs-acl)

Purpose

Use the gbl-nfs-acl nis domain command to set the NIS domain for the current NFS access list.

Use no nis domain to remove the NIS domain from the access list.

Mode

gbl-nfs-acl

Security Role(s)

crypto-officer

Syntax

`nis domain` ***`domain`***
`no nis domain`

domain (1-256 characters) is the NIS-domain name (for example, “myorg.org” from lnx3.myorg.org). This domain must be pre-mapped to a NIS server with the nis domain command, from gbl mode.

Default(s)

None

Guidelines

If you plan to use NIS groups in your NFS access list, you must identify the NIS domain with this command. Use show nis domain for a list of configured NIS domains, or use nis domain to create a new one. To view the netgroups in a domain, use show nis netgroup.

The permit netgroup command (permit (gbl-nfs-acl)) permits all hosts in a netgroup to pass the NFS access list.

Samples

```
bstnA(gbl-nfs-acl[westcoast])# nis domain sfmed.com
```

sets the NIS domain to “sfmed.com” for the “westcoast” access list.

```
bstnA(gbl-nfs-acl[midwest])# no nis-domain
```

removes the NIS domain setting from the “midwest” access list.

Related Commands

permit (gbl-nfs-acl)
show nis domain
nis domain

nis update

Purpose The ARX caches a database of NIS netgroups and all of their DNS-resolved IP addresses. Use the nis update command to refresh this cache by querying the NIS server(s) and the local DNS server(s).

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer or crypto-officer

Syntax **`nis update [`*`domain`*`]`**

domain (1-256 characters) focuses the update on a single NIS domain. Use show nis domain to view all NIS domains. If this is omitted, the switch refreshes its cache for all configured NIS domains.

Default(s) None

Guidelines This command creates one report per updated domain. Each report is named "nis-update.*domain-name*.rpt." Use show reports to list all reports, including NIS-update reports. To follow the progress of the NIS-update operation, you can use tail reports report-name follow. Use show reports file-name to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands. If you want to truncate the report before it finishes, use the truncate-report command.

To schedule this command to run on a regular basis, use the at command.

The show nis netgroup command shows the netgroups in a NIS domain, or the hosts in a particular netgroup. This is the contents of the current NIS-netgroup cache. The show nis domain command shows when the most-recent NIS update occurred.

In a redundant pair, the NIS update works independently on each peer. The output of show nis netgroup and the NIS reports therefore only apply to the current peer. The benefit of these redundant updates is that failovers do not incur any extra down time for NIS.

Samples

```
bstnA# nis update
```

updates all NIS domains on the switch.

```
bstnA# nis update wwmed.com
```

updates a single NIS domain, "wwmed.com."

Related Commands

at
show nis domain
show nis netgroup
nis domain

Figure 19.1 *Sample Report: nis-update.wwmed.com.rpt*

```
bstnA# show reports nis-update.wwmed.com.rpt
**** NIS Update Report: Started at Fri Mar  5 03:00:34 2010 ****
**** Software Version: 5.02.000.12543 (Mar  2 2010 20:13:33) [nbuilds]
**** Hardware Platform: ARX-4000
```

```
**** Report Destination:
**** NIS Domain: wwmed.com
**** NIS Server: 192.168.25.201
**** NIS Server: 192.168.25.204
**** NIS Server: 192.168.25.205

**** Legend:
****   HN = Hostname not found.
****   NP = Netgroup parsing error.
****   NG = Netgroup not found.
****   WG = Watched netgroup has changed contents.
****   NE = NIS server error.

Status            Hostname/Netgroup
----------------  ------------------------------------------------------------
[HN            ]  not_a_real_host1 in group: bad_hosts
[HN            ]  not_a_real_host2 in group: bad_hosts
[HN            ]  ommegang in group: chassis
[HN            ]  bismark in group: chassis
[HN            ]  buzzsaw in group: dnashosts

...

[HN            ]  london in group: sixthousands
[HN            ]  montreal in group: sixthousands
[HN            ]  lasvegas in group: sixthousands

Netgroups Processed:                       2,396
Hosts Processed:                          48,043
Hostnames Not Found:                         539
Netgroup Parsing Errors:                       0
Netgroups Not Found:                           0
Watched Netgroup Changes                       0

**** Elapsed time:           00:00:23
**** NIS Update Report: DONE at Fri Mar  5 03:30:15 2010 ****
```

permit (gbl-nfs-acl)

Purpose Use this command to permit NFS access for one trusted subnet or NIS netgroup.

Use the no form of the command to remove a permit rule for a subnet or netgroup.

Mode gbl-nfs-acl

Security Level crypto-officer

Syntax: permit (for a subnet)

```
permit ip-address mask [read-only] [ root {squash | allow} ]
no permit ip-address mask [read-only] [ root {squash | allow} ]
```

ip-address (0.0.0.0–255.255.255.255) is the address of the subnet to be allowed access.

mask (0.0.0.0–255.255.255.255) is the netmask (network part of the ip address).

read-only (optional) limits users to read-only access.

root (optional) applies this rule to root-user access and requires one of the following keywords:

squash disables root access and remaps it to the configured UID and GID settings (provides more security).

allow enables root-user access.

Syntax: permit netgroup

```
permit netgroup group-name [read-only] [ root {squash | allow} ]
no permit netgroup group-name
```

group-name (1-1024) is the address of the NIS netgroup to be allowed access.

read-only (optional) limits users to read-only access.

root (optional) applies this rule to root-user access and requires one of the following keywords:

squash disables root access and remaps it to the configured UID and GID settings (provides more security).

allow enables root-user access.

Default(s) read-write;

root squash.

Guidelines

By default, a new NFS access list denies access to all subnets and NIS netgroups. This command enables you to selectively allow access for a trusted subnet or netgroup.

An access list can only use netgroups from a single NIS domain; use nis domain (gbl-nfs-acl) to choose one for the current access list.

An NFS-access list can support a maximum of 2048 rules. Each host in a NIS netgroup results in a separate rule, so use show nfs-access-list ... resolve-netgroups to verify that the netgroups are not too large for the access list.

A new permit rule squashes root access by default. That is, if a client logs in as the *root* user and accesses the NFS share, the ARX translates the client's user ID to an *anonymous* ID with limited access privileges. In gbl-nfs-acl mode, you can change the anonymous User/Group IDs through the anonymous-gid and anonymous-uid commands.

The no form of the command always removes the rule, whether or not you specify any options (such as read-write, root allow, and so on). This facilitates copying and pasting a rule from the show nfs-access-list output to the CLI and placing a no in front of it.

Guidelines: Rule Order

The order that you enter rules determines the order in which they are compared to client IPs. If a client's IP address matches more than one rule in the access list, the ARX uses the first matching rule and ignores the rest.

For example, consider the two permit rules below. Clients in 192.168.10.x match the first rule, while clients outside that subnet in the same Class B network (192.168.x.x) match the second rule. This allows read-only access to clients in the Class B network but full read-write access to clients in the smaller Class-C subnet:

```
permit 192.168.10.0 255.255.255.0 read-write
permit 192.168.0.0 255.255.0.0 read-only
```

If the rules were reversed, clients in the Class C subnet would match the read-only rule before reaching the read-write rule that was intended for them.

No permit or deny rule is enforced until you exit gbl-nfs-acl mode. This ensures that all permit and deny rules are enforced at once, after you finish setting them up in the desired order. Use the exit or end command to exit the mode.

Samples

```
bstnA(gbl)# nfs-access-list eastcoast
bstnA(gbl-nfs-acl[eastcoast])# permit 172.16.100.0 255.255.255.0
read-write
bstnA(gbl-nfs-acl[eastcoast])# ...
```

permits read-write access to clients at 172.16.100.0.

```
bstnA(gbl)# nfs-access-list eastcoast
bstnA(gbl-nfs-acl[eastcoast])# permit 172.16.204.0 255.255.255.0
read-only root allow
bstnA(gbl-nfs-acl[eastcoast])# ...
```

allows *root* access from clients at 172.16.204.0. To control security, access is read-only for this rule.

```
bstnA(gbl-nfs-acl[eastcoast])# permit netgroup nurses
```

permits access to all hosts in the "nurses" netgroup.

Related Commands

nis domain
anonymous-gid
anonymous-uid
deny
nfs-access-list
show nfs-access-list

show nfs-access-list

Purpose Use this command to display a summary of all configured NFS access lists or specify an individual list to view its details.

Mode (any)

Security Level operator

Syntax `show nfs-access-list [list-name [resolve-netgroups]]`

list-name (optional; 1-64 characters) is the access list you want to view.

resolve-netgroups (optional) expands the NIS netgroups. This shows every resolved host in every NIS netgroup, in order. Without this option, netgroups are summarized on a single line and counted as a single rule.

Default(s) None

Guidelines The show nfs-access-list command displays the following information:

Access List Name: The names of all configured access lists.

Anon UID: The anonymous User ID number assigned to *root* when root squashing is enabled (the default). You can change this with anonymous-uid.

Anon GID: The anonymous Group ID number assigned to *root* when root squashing is enabled (the default).Use anonymous-gid to edit this.

Num Rules: The number of permit and/or deny rules applied to this access list. This counts each NIS netgroup as a single rule; use the resolve-netgroups command to find the total count, including every host in every netgroup.

Num References: The number of times this access list is used by an NFS service(s).

Guidelines: Detailed Output

If you enter a list name, the output also shows the description (gbl-nfs-acl) for the access list, if any, and the list's exact rules. The order is important; if a client matches two rules in the list, the switch follows the first rule and ignores the second.

Two additional fields appear at the bottom if you use the resolve-netgroups flag:

Number of entries in access list: The total number of rules, including each host in the expanded netgroups. An error appears above this field if the number of rules exceeds the maximum, 2048. Each host in a netgroup requires a rule, so large netgroups can cause an access list to exceed its maximum. Only the first 2048 rules are used.

Status line: A line summarizes the results of resolving all NIS netgroups and each hostname in the netgroup:

- "All Netgroup(s) were successfully resolved." This indicates success.
- "Netgroup(s) were not completely resolved." This may indicate that one or more hosts in the NIS netgroups cannot be resolved at the local DNS server. Use show ip domain to find the DNS servers that the switch is using.
- "Domain was not found." The NIS domain for this access list, set with nis domain (gbl-nfs-acl), is not supported at the NIS server. Use show nis domain to find the NIS server(s) used by the ARX.
- "One or more Netgroup(s) were not found." A NIS netgroup contained another NIS netgroup as one of its members; the NIS server did not define one or more of these member netgroups.
- "The switch has not resolved all netgroups for this domain." Either the first update is not yet finished (for example, right after finishing the NIS configuration or rebooting the switch), or every NIS update has failed since the most-recent reboot. Use nis update to invoke another NIS update.

Samples

`bstnA# show nfs-access-list`

displays a list of all configured access lists. For example, see Figure 19.2.

`bstnA# show nfs-access-list eastcoast`

displays detailed information for this access list, including the permit and deny rules for the associated subnet(s) and/or netgroups. See Figure 19.3 on page 19-19 for sample output.

`bstnA# show nfs-access-list eastcoast resolve-netgroups`

expands all NIS netgroups into their component rules, one rule per host in the netgroup. See Figure 19.4 on page 19-19 for sample output.

Related Commands

anonymous-gid
anonymous-uid
deny
permit (gbl-nfs-acl)
description (gbl-nfs-acl)
nfs-access-list

Figure 19.2 *Sample Output: show nfs-access list*

```
bstnA# show nfs-access-list
Access List Name    Anon UID Anon GID    Num Rules      Num References
  ------------------------------------------------------------------------
```

```
eastcoast           100      100               8                 2
westcoast           65534    65534             2                 0
```

Figure 19.3 *Sample Output: show nfs-access-list eastcoast*

```
bstnA# show nfs-access-list eastcoast

Access List Name: eastcoast
  Description:               allowable subnets in MA, NY, & DC
  NIS Domain:                wwmed.com
  Anonymous UID:             100
  Anonymous GID:             100
  Number of References:      2

permit 172.16.100.0 255.255.255.0 read-write root squash
permit 172.16.204.0 255.255.255.0 read-only  root allow
permit 172.16.0.0 255.255.0.0 read-write root squash
permit netgroup surgeons read-write root allow
permit netgroup medtechs read-only  root squash
deny  192.168.77.0 255.255.255.0
deny  192.168.202.0 255.255.255.0
permit 192.168.0.0 255.255.0.0 read-write root squash
```

Figure 19.4 *Sample Output: show nfs-access-list eastcoast resolve-netgroups*

```
bstnA# show nfs-access-list eastcoast resolve-netgroups

Access List Name: eastcoast
  Description:               allowable subnets in MA, NY, & DC
  NIS Domain:                wwmed.com
  Anonymous UID:             100
  Anonymous GID:             100
  Number of References:      2

permit 172.16.100.0 255.255.255.0 read-write root squash
permit 172.16.204.0 255.255.255.0 read-only  root allow
permit 172.16.0.0 255.255.0.0 read-write root squash
permit 10.51.201.71 255.255.255.255 read-write root allow
permit 10.51.201.72 255.255.255.255 read-only  root squash
permit 10.51.201.73 255.255.255.255 read-only  root squash
permit 10.51.201.74 255.255.255.255 read-only  root squash
deny  192.168.77.0 255.255.255.0
deny  192.168.202.0 255.255.255.0
permit 192.168.0.0 255.255.0.0 read-write root squash

Number of entries in access list: 10
All Netgroup(s) were successfully resolved.
```

show nis domain

Purpose Use this command to display a summary of all configured NIS domains, or to show details for one NIS-domain configuration.

Mode (any)

Security Level operator

Syntax `show nis domain [domain-name]`

domain-name (optional; 1-256 characters) is the NIS domain you want to view. If you omit this, the output displays a summary of all NIS domains configured on the switch.

Default(s) None

Guidelines The summary form of the show nis domain command displays the following information:

NIS Domain is the name of the NIS domain. Use the nis domain command to configure a new NIS domain, or change an existing one.

Last Update is the date and time that the ARX last updated its internal NIS database. This occurs when each NIS domain is first configured, and whenever someone issues the nis update command.

Status summarizes the results of the most-recent NIS update. This is "Success," "Updating," or "Failed."

Servers are the NIS servers for this NIS domain. You can use the ip address (gbl-nis-dom) command to identify more servers that support this domain. This shows the order in which the servers are used; if the first server fails, the switch tries the second, and so on.

The detailed output adds several more fields:

Last Successful Update is the date and time for the last NIS update that ended with a "Success" status.

Netgroups is the number of netgroups defined for this NIS domain.

Netgroup Resolution Errors is the number of netgroup entries that the switch failed to parse. These are typically malformed lines in the NIS server's configuration file for netgroups.

Hosts is the number of hosts found in all the netgroups.

Hosts Resolved is the number of hosts that were successfully resolved to IP addresses. These are DNS resolutions, made by an external DNS server; use show ip domain to see the DNS server(s) used by this switch. If this number is lower than the number for Hosts, above, some hosts were not resolved.

Samples

```
bstnA# show nis domain
```

displays a list of all configured NIS domains. See Figure 19.5 for sample output.

```
bstnA# show nis domain wwmed.com
```

displays detailed information for the "wwmed.com" NIS domain. Figure 19.6 shows sample output.

Related Commands

nis domain
ip address (gbl-nis-dom)
nis update
show ip domain

Figure 19.5 *Sample Output: show nis domain (all)*

```
bstnA# show nis domain

NIS Domain                      Last Update   Status    Servers
-------------------------------  ------------  --------  ---------------
wwmed.com                       26 Jan 03:24  Success   192.168.25.201
                                                        192.168.25.204
                                                        192.168.25.205
```

Figure 19.6 *Sample Output: show nis domain wwmed.com*

```
bstnA# show nis domain wwmed.com

NIS Domain:                     wwmed.com
Server(s):                      192.168.25.201
                                192.168.25.204
                                192.168.25.205
Last Update:                    Fri Jan 26 03:24:16 2007
Last Update Status:             Success
Last Successful Update:         Fri Jan 26 03:24:16 2007
Netgroups:                      2397
Netgroup Resolution Errors:     0
Hosts:                          48046
Hosts Resolved:                 47507
```

show nis netgroup

Purpose Use this command to list the netgroups in a given NIS domain, or to show all of the hosts in a given netgroup.

Mode (any)

Security Level operator

Syntax **`show nis netgroup`** ***`domain`*** **`[`*****`netgroup`***`]`**

domain (1-256 characters) is the NIS domain you want to view.

netgroup (optional; 1-1024 characters) specifies a single netgroup. If you enter this, the command shows all hosts in the netgroup.

Default(s) None

Guidelines The summary form of the show nis netgroup command displays an alphabetical list of all netgroups defined for the domain. The ARX finds these at the back-end NIS servers; use show nis domain for a list of NIS servers.

The detailed form of the command shows a table with one row for each host found in the netgroup. The Hostname is the name found in the netgroup, and the IP Address is resolved at the local DNS server. Use show ip domain for a list of local DNS servers.

The switch ignores all users and groups in the netgroup. As a proxy, the ARX does not authenticate specific users and groups; the back-end filers perform user/group authentications, and the switch passes the results back to the client.

Samples `bstnA# show nis netgroup wwmed.com`

displays a list of all netgroups in the "wwmed.com" domain. See Figure 19.7 for sample output.

`bstnA# show nis netgroup wwmed.com medtechs`

shows every host in the "medtechs" netgroup. Figure 19.8 shows sample output.

Related Commands nis domain
ip address (gbl-nis-dom)
nis update
show ip domain

Figure 19.7 *Sample Output: show nis netgroup wwmed.com*

```
bstnA# show nis netgroup wwmed.com

Netgroup
-----------------------------------------------------------------
acopia_hosts
auto_1
...
medtechs
oems
onethoussands
```

```
sixthousands
surgeons

Total Netgroups: 2396
```

Figure 19.8 *Sample Output: show nis netgroup wwmed.com medtechs*

```
bstnA# show nis netgroup wwmed.com medtechs

Netgroup successfully resolved.

Hostname                                          IP Address
------------------------------------------------- -----------------
bench2.wwmed.com                                  10.51.201.72
bench3.wwmed.com                                  10.51.201.73
bench4.wwmed.com                                  10.51.201.74

Total Resolved Hosts: 3
```

20

External Filer

An *external filer* is either a Network-Attached Storage (NAS) device or a file server.

cifs connection-limit

Purpose

Some Tier-2 CIFS filers can only support a limited number of simultaneous CIFS connections. You can use the cifs connection-limit command to set a maximum number of CIFS connections from the ARX to the current filer.

For filers and servers that can support a large number of CIFS connections, you can use the no cifs connection-limit command. This removes any restriction on CIFS connections between the ARX and the filer.

Mode

gbl-filer

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
cifs connection-limit maximum-connections
no cifs connection-limit
```

maximum-connections (at least 250 for the ARX-4000; at least 150 for all other platform types) is the limit on the number of CIFS connections to the current filer.

Default(s)

no limit on the number of CIFS connections

Guidelines

This command is only necessary for filers that can support a limited number of CIFS connections.

If you set a limit that is lower than the current number of CIFS connections, the CLI prompts you with a choice: you can drop all connections immediately or block new connections until enough of the current connections drop off. If you choose to wait, you always have the option to later use the drop filer-connections command, which forces the connections to terminate immediately. To show the current connections to the filer, use show filer connections. The show statistics filer connections command shows connection statistics over time, and the clear statistics filer connections command clears the statistics.

All external filers with the same IP address must have the same number of connection limits. That is, if your configuration includes more than one “external-filer” object for the same back-end filer, this command sets the same connection limit for both objects.

Each NSM processor always has at least 2 CIFS connections per filer.

Sample

```
bstnA(gbl-filer[smb1])# cifs connection-limit 500
```

limits the ARX to 500 CIFS connections between itself and the “smb1” filer.

Related Commands

external-filer
drop filer-connections
show filer connections
show statistics filer connections
clear statistics filer connections

cifs-port

Purpose For filers that do not use a well-known port for CIFS, use the cifs-port command to set the CIFS port.

Use the no form of this command to use well-known CIFS ports.

Mode gbl-filer

Security Role(s) storage-engineer or crypto-officer

Syntax

```
cifs-port port
no cifs-port
```

port is the TCP port number (1-65535).

Default(s) TCP 445, falling back to 139. See the *Guidelines* below.

Guidelines By default, the ARX sends its CIFS messages to port 445 or 139 at the external filer. Port 445 supports raw CIFS communication over TCP, port 139 supports CIFS through NetBIOS; the ARX tries port 445 first, then falls back to 139 if that fails. For filers that do not listen at either of these well-known ports, use the cifs-port command to set the port.

Whenever the ARX falls back to port 139 to make a NetBIOS connection, it uses a "Source Computer" name of "ACOPIA_SWITCH." Every proxy-IP address identifies itself by this same name. You can see this name through MMC, by examining "Active Sessions" on the back-end filer.

This command is not needed for external filers that support NFS exports.

Sample

```
bstnA(gbl-filer[fs2])# cifs-port 1111
```

sets the CIFS port to 1111.

Related Commands external-filer

description (gbl-filer)

Purpose

Use the optional description command to set a descriptive string for the current external filer. This appears in the show command.

Use the no form of the command to delete the description.

Mode

gbl-filer

Security Role(s)

storage-engineer or crypto-officer

Syntax

`description` ***`text`***
`no description`

text (1-255 characters) is your description. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s)

no description

Guidelines

The description appears in the output for show external-filer.

Sample

```
bstnA(gbl-filer[das1])# description “shares with financial data (LINUX filer, rack 14)”
```

specifies a description for the current external filer.

Related Commands

external-filer
show external-filer

external-filer

Purpose

Use the external-filer command to begin external-filer configuration. This is a logical representation of an external file server or NAS filer connected to the ARX.

Use the no form of this command to remove the external-filer configuration.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

`external-filer` ***`name`***
`no external-filer` ***`name`***

name is a name that you choose for the external filer. Use up to 64 characters. You cannot use "all;" this is a reserved keyword.

Default(s)

None.

Guidelines

When you create a new external filer, this command prompts for confirmation. Enter **yes** to continue.

This command puts you into gbl-filer mode. From that mode, you must use ip address to identify the external filer. If the filer is a Windows cluster, you must also use spn to provide the Service-Principal Name (SPN) for the cluster.

You can also use the description (gbl-filer) command to describe the filer.

The ARX uses the following NFS settings for the external filer:

- NFS port from a call to the filer's portmapper service,
- NFS-mount port from a call to the filer's portmapper service,
- Unix User ID 0 (*root*) to read the filer's directory trees, and
- a null NIS domain.

An external filer's default CIFS port is 445, or 139 for filers that do not support CIFS over raw TCP; use cifs-port to change this, if necessary.

Use show external-filer to show the filer configuration. The show statistics namespace ... fastpath command shows read/write statistics for each of the filer's shares. For statistics on NFS and/or CIFS traffic between the switch and this filer, use show statistics filer. You can use show filer connections to see all current connections from the switch to the filer.

Guidelines: Preparation for ARX Snapshots

If the filer supports snapshots (point-in-time copies), you can use several commands to prepare it for use in an ARX snapshot. The filer-type command identifies the filer's vendor. The proxy-user (gbl-filer) command associates a proxy user with this filer, so that the ARX can uses its credentials to access the filer's CLI or API. (An ARX volume coordinates its snapshots by logging into its filers and issuing snapshot commands.) You can also use the ip address ... management command to specify a management IP for accessing the CLI (by default, the ARX volume uses the primary-IP address). The manage snapshots command declares that the filer can support snapshots or checkpoints.

Guidelines: Changing the Filer's IP Address

You can change the primary by re-running the ip address command with a new address, before the filer is used in any volume. After any of the filer's shares are used in any volume, you can only change the address with ip address ... change-to, followed by a failover or reboot (reload).

Samples

```
bstnA(gbl)# external-filer das1
This will create a new filer.

Create filer 'das1'? [yes/no] yes
bstnA(gbl-filer[das1])#
```

creates an external filer named "das1."

```
bstnA(gbl)# no external-filer das1
```

removes the configuration for the external filer named "das1."

Related Commands

ip address
ip address ... change-to
show external-filer
cifs-port
filer-type
proxy-user (gbl-filer)
manage snapshots
show statistics namespace ... fastpath
show statistics filer
show filer connections

ext-filer-ip-addrs activate

Purpose You can use the ip address ... change-to command to prepare for changing one or more IP addresses for your external-filers, then use this command to reboot the ARX and activate all the address changes. If the ARX has a redundant peer, this command reboots both peers.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer or crypto-officer

Syntax **`ext-filer-ip-addrs activate`**

Default(s) None

Guidelines This command causes the current ARX (and its redundant peer, if there is one) to reboot. This causes a service outage, so you should run the command during non-busy hours. The CLI prompts for confirmation before rebooting; enter **yes** to proceed.

After the reboot(s) finish, all IP-address changes made by the ip address ... change-to command go into effect. The ip address ... change-to command has no effect on client/server traffic until you invoke this command.

If there are no IP-address changes pending, the command prompts you that there is nothing to do. No reboot occurs in this case. As mentioned above, you can use the ip address ... change-to command to prepare for changing a filer's IP address.

The ip address ... change-to command is only necessary after one of the filer's shares or exports is used in an ARX volume. If none of the filer's shares are in active use, you can re-run the ip address command to change the address, and the ext-filer-ip-addrs activate command is unnecessary.

Samples

```
bstnA# ext-filer-ip-addrs activate
Activates all new external filer primary IP addresses and reboots the
chassis.

Are you sure? [yes/no] yes
```

activates all IP-address changes made by the ip address ... change-to command and reboots the ARX. This would also reboot the redundant peer if one was configured.

Related Commands external-filer
ip address ... change-to

filer-type

Purpose An ARX volume can coordinate *snapshots* (point-in-time copies) of all of its back-end shares. To support these coordinated snapshots, the ARX volume logs into each filer's CLI and runs snapshot or checkpoint commands. Use the filer-type command for a filer that supports snapshots, to identify the filer vendor. An ARX volume can use this information to determine the proper CLI-command set to invoke snapshots or checkpoints.

Use the no form of this command to designate no filer type. This makes it impossible for an ARX volume to use the filer for coordinated snapshots.

Mode gbl-filer

Security Role(s) storage-engineer or crypto-officer

Syntax

```
filer-type network-appliance [management-protocol {ssh | rsh}]
filer-type emc [nas-db-path path]
filer-type data-domain
filer-type bluearc
filer-type windows [cluster] [port winrm-port]
no filer-type
```

network-appliance declares that NetApp is the vendor for this filer.

> **management-protocol {ssh | rsh}** (optional) chooses the shell protocol for accessing the filer's CLI. The default is SSH. For NetApp filers that do not allow SSH access, you can use **management-protocol rsh**.

emc indicates that this is an EMC Celerra server.

> **nas-db-path *path*** (optional, 1-64 characters) defines a unique NAS-database path for an EMC Celerra server. Specifically, this is the value of the NAS_DB environment variable on the server. The default for this path is "/nas;" contact EMC Support for the proper setting on your EMC server.

data-domain means that the filer is an EMC Data Domain system. The CLI uses SSH to access the Data Domain CLI.

bluearc indicates that the filer is Hitachi HNAS, powered by BlueArc. The CLI uses SSH to access the BlueArc CLI.

windows means that the filer is a Windows server that supports snapshots and management through WinRM (see *Guidelines*, below).

> **cluster** (optional) asserts that this is in a Windows server 2008 (or later) cluster. If the Windows cluster has NetBIOS disabled, there is no way for the ARX software to automatically determine that it is a cluster, and this information is necessary to support snapshots. You can use this option to overcome this issue and directly inform the ARX.
>
> **port *winrm-port*** (optional, 1-65535) specifies the port where the server's WinRM listener is waiting for instructions.

Default(s)

`no filer-type`

nas-db-path *path* - the *path* default is "/nas"

cluster - by default, this flag is down and the ARX attempts to automatically discover if the filer is in a cluster

port *winrm-port* - the default for *winrm-port* is 80, the HTTP port.

Guidelines

This command identifies the vendor of the current filer. You must use this command to prepare the filer for coordinated-snapshot support with a snapshot rule. Additionally, you must create a proxy-user with credentials for logging into the filer's management interface, and you must use the proxy-user (gbl-filer) command to associate that proxy user with this filer. You can also use the ip address ... management command to specify a management IP for accessing the CLI (by default, the ARX volume uses the primary-IP address). The manage snapshots command declares that the filer is eligible for ARX snapshots.

Guidelines: NetApp Settings

Each NetApp volume to be used for ARX snapshots must have the nosnapdir option turned off. From the NetApp CLI, you can use the vol options command to turn this option off. For example: v**ol options vol2 nosnapdir off**.

If you use **management-protocol ssh** (the default), please verify that the idle timeout value, **ssh.idle.timeout**, is set properly. The ARX cannot create snapshots on the NetApp if this option is set to 0 (zero). From the NetApp CLI, use **options ssh** to check the current timeout setting, and use **options ssh.idle.timeout 600** to set the idle timeout to 600 seconds (the default for most releases).

Guidelines: Windows and WinRM

A Windows filer must have Windows Remote Management (WinRM) installed and configured for the ARX to manage its snapshots. The minimum configuration is a WinRM listener that allows HTTP negotiations (encrypted or unencrypted), and allows up to 5 minutes before timing out a snapshot operation. Use the **port** option if the listener is waiting at any IP port other than 80. The ARX volume communicates with the WinRM listener to create, list, and remove snapshots from the Windows Server. The ARX uses Kerberos authentication over HTTP; use active-directory update seed-domain to prepare for Kerberos authentication.

You can use the following command at the Windows filer's DOS prompt to create a 5-minute timeout for snapshots: **winrm set winrm/config @{MaxTimeoutms="300000"}**.

Samples

```
bstnA(gbl-filer[nas10])# filer-type network-appliance
```

designates the “nas10” filer as a NetApp filer that supports ARX snapshots.

```
bstnA(gbl-filer[nas2])# filer-type network-appliance
management-protocol rsh
```

establishes RSH as the shell protocol for accessing the “nas2” filer, another NetApp device.

```
bstnA(gbl-filer[fs2])# filer-type windows
```

designates the “fs2” filer as a Windows server that supports ARX snapshots. This sends its WinRM instructions through the default port, 80. If “fs2” is part of a Windows server 2008 (or later) cluster, it presumes that NetBIOS is enabled on that cluster, and the ARX can therefore automatically discover its cluster membership.

```
bstnA(gbl-filer[nasE4])# filer-type emc nas-db-path /nas/rdf/500
```

declares that “nasE4” is an EMC Celerra server with a non-default NAS_DB value of “/nas/rdf/500.”

```
bstnA(gbl-filer[das1])# no filer-type
```

indicates that the “das1” filer does not support ARX snapshots or CLI/API access.

Related Commands

filer
snapshot rule
proxy-user (gbl-filer)
manage snapshots
show external-filer

ignore-name

Purpose

Some directories are only related to a filer's local operation, such as local snapshot and checkpoint directories, and should not be presented to the client. Use the ignore-name command to identify each virtual directory to be ignored by the ARX. The ARX also ignores all files with the same name.

Use the no form of this command to remove a directory from the "ignore list."

Mode

gbl-filer

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
ignore-name directory
no ignore-name directory
```

directory (1-256 characters) identifies the directory to be ignored (for example, ".snapshot"). Do not use any slashes (/) in the directory name; "mydir" is valid, but "mydir/subdir" is not. You can use a * at the end of the directory name as a wildcard, as long as the directory name starts with a "." and contains at least three characters. These wildcards only apply to the root of the share: ".ignore*" ignores /.ignore2, but does not ignore /docs/.ignore7.

Default(s)

.snapshot, ~snapshot, .etc, .ckpt*, $RECYCLE.BIN, System Volume Information, SIS Common Store

Guidelines

The filer command includes directories into a namespace volume. This command omits any subdirectories that should not be included. You can ignore up to eight directories per external filer.

As mentioned above, this command also makes the ARX ignore files of the same name.

Ignore only virtual directories like "~snapshot" and "lost+found." If you ignore a real directory (or file) that exists below the share's root, clients cannot delete the ignored directory/file's parent. For example, if you ignore "/ignoreMe" directories and a "/var/export" share contains "/var/export/myDir/ignoreMe," clients cannot delete "/var/export/myDir."

The directory is ignored for all shares on the filer; for example, if a ".ckpt1" directory appears in each of three shares, it is ignored in all three of them.

The CLI prompts for confirmation if you use a wildcard. Enter **yes** to continue.

Several filer vendors have sets of directories that they commonly use for backups, snapshots, and storage maintenance. All except "lost+found" are ignored by default:

- EMC: .etc, lost+found, .ckpt*
- EMC Data Domain: .snapshot
- Network Appliance: .snapshot, ~snapshot
- Windows: "System Volume Information", "SIS Common Store"
- Hitachi HNAS, powered by BlueArc: .snapshot, ~snapshot

The show external-filer command shows all ignored directories.

Samples

```
bstnA(gbl-filer[nas10])# ignore-name ~snapshot
```

ignores all "~snapshot" directories on the "nas10" filer.

```
bstnA(gbl-filer[nasE1])# ignore-name .ckpt*
Wildcard matches are only made in the root directory of a share.
Is this the intended behavior? [yes/no] yes
```

ignores all directories starting with ".ckpt" on the "nasE1" filer (.ckpt1, .ckpt2, ...ckpt33, and so on).

```
bstnA(gbl-filer[das1])# no ignore-name .bin
```

stops ignoring all ".bin" directories.

Related Commands

filer
show external-filer

ip address

Purpose

Use the ip address command to identify the filer or file server associated with the current external-filer configuration object.

Use the no form of the command to remove an IP address from the current external-filer configuration.

Mode

gbl-filer

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
ip address address [secondary | management]
no ip address address [secondary | management]
```

address identifies the external filer (for example, 192.168.70.65). This address must be on a server/proxy-IP subnet (see ip proxy-address) or reachable through a gateway on that subnet (via static route: see ip route to create a static route).

secondary (optional) flags a physical-IP addresses in a multi-homed filer. Some filers are configured redundantly, accepting all packets at a shared virtual-IP address and responding to some UDP packets from their physical-IP addresses. For these filers, enter the virtual-IP address without any flag, then add up to four physical addresses with the **secondary** flag. Note that CIFS is a TCP-only protocol, so this option is unnecessary if you only use the filer for CIFS.

management (optional) is for an out-of-band (OOB) or external management address. The ARX uses this address to access the filer's CLI. A volume with a snapshot rule accesses the CLI to issue snapshot or checkpoint commands at the filer. If you do not specify any IP address with this flag, it uses the primary-IP address.

Default(s)

None

Guidelines

This command is required for configuring a NAS device or file server as an external filer. To verify that the address is valid in the current configuration, use the show exports command before you use this one.

You can repeat this command up to four times with the **secondary** flag, to identify secondary addresses for the filer. The **management** flag indicates an interface for accessing the filer's CLI; this is unnecessary if the CLI is accessible through the primary interface.

To change the primary IP address of a filer whose storage is in use, you can use the ip address ... change-to command followed by the reload command. If the filer does not have any exports or shares used by any volume, you can change the IP address by re-running this command with the new address.

Samples

```
bstnA(gbl-filer[das1])# ip address 192.168.25.19
```

designates 192.168.25.19 as the IP address for the “das1” filer.

```
bstnA(gbl-filer[nas3])# ip address 192.168.25.133
bstnA(gbl-filer[nas3])# ip address 192.168.25.134 secondary
bstnA(gbl-filer[nas3])# ip address 192.168.25.135 secondary
```

specifies that 192.168.25.133 is the primary address for “nas3,” then identifies two secondary addresses.

```
bstnA(gbl-filer[nasE1])# ip address 192.168.25.51
bstnA(gbl-filer[nasE1])# ip address 192.168.25.52 management
```

provides a primary address and a separate management address for the “nasE1” filer.

Related Commands

external-filer
ip proxy-address
ip route
show exports
show external-filer

ip address ... change-to

Purpose

Use the ip address ... change-to command to prepare for changing the current external-filer's IP address. The new IP address becomes the primary IP address for the filer the next time someone uses the ext-filer-ip-addrs activate command.

Use the no form of the command to remove the new IP address from the current external-filer configuration. Use this if you want the external filer's IP address to remain the same in case someone invokes ext-filer-ip-addrs activate.

Mode

gbl-filer

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
ip address address change-to new-address
no ip address address change-to new-address
```

address identifies an existing external filer. You can use show external-filer for a full list of external filers and their IP addresses.

new-address is the address to assign to this filer the next time someone invokes ext-filer-ip-addrs activate. You cannot use an existing secondary address (set with the ip address command) or any other IP address that is already in use on the ARX.

Default(s)

None

Guidelines

The CLI prompts for confirmation before reserving a new IP address for the external-filer; enter **yes** to proceed.

This command has no effect on client/server traffic until someone uses ext-filer-ip-addrs activate to activate the change. After you run this command, you can change the IP address of the filer itself, and then use the ext-filer-ip-addrs activate command to reboot the ARX or ARX pair. You only need to activate once for multiple IP-address changes.

This command is only necessary after one of the filer's shares or exports is used in an ARX volume. If none of the filer's shares are in active use, you can re-run the ip address command to immediately change the address.

Guidelines: IP Address Settings that are Unaffected by this Command

The ip address ... change-to command has no effect on the following IP-address settings:

- The *quorum-disk* share, used for redundancy. If the filer is the host for a quorum-disk share, you must re-run the quorum-disk command with the new IP address to re-connect to the quorum disk. Run this command on both ARX peers in the redundant pair.
- The management-IP address used for snapshots and other filer-CLI access. If the filer has shares with snapshots, and the management IP changes, you must re-run the ip address ... management command to continue to support an ARX snapshot rule at this filer.

The ext-filer-ip-addrs activate command is not necessary for these commands to take effect.

Samples

```
bstnA(gbl-filer[nas3])# ip address 192.168.25.133 change-to
192.168.25.198
This operation will add the IP address for the external filer as the
new IP address.
This change will not take effect until after 'ext-filer-ip-addrs
activate' is run.
Change the external filer's IP address ? [yes/no] yes
```

assigns a new IP address to the external-filer named “nas3.” The new address, 192.168.25.198, takes affect the next time someone runs ext-filer-ip-addrs activate.

```
bstnA(gbl-filer[nas3])# no ip address 192.168.25.133 change-to
192.168.25.198
```

nullifies the above command. The IP address for “nas3” will remain the same if anyone runs ext-filer-ip-addrs activate.

Related Commands

external-filer
ip address
show external-filer
ext-filer-ip-addrs activate

manage snapshots

Purpose

An ARX volume can coordinate *snapshots* (point-in-time copies) of all of its back-end shares. A limited number of filers and filer-OS releases support snapshots. Use the manage snapshots command to indicate that the current filer supports coordinated snapshots.

Use the no form of this command to indicate that the current filer does not support snapshots.

Mode

gbl-filer

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
manage snapshots
no manage snapshots
```

Default(s)

```
no manage snapshots
```

Guidelines

This command declares that the filer can support snapshots. This is the final command to prepare the filer for coordinated-snapshot support with a snapshot rule. Before you enter this command, use the filer-type command to enter some vendor-specific information about the filer. Additionally, you must create a proxy-user with the credentials for accessing into the filer's CLI or API, and you must use the proxy-user (gbl-filer) command to associate that proxy user with this filer. For some filers, you can use the ip address ... management command to specify a management IP for accessing the CLI (by default, the ARX volume uses the primary-IP address).

You can use the no manage snapshots command to suspend snapshot support on an entire filer. While this is disabled, no snapshot rule can take a snapshot on the filer. In a volume where other back-end filers support snapshots, this creates a *sparse snapshot* from which some of the volume's shares are *excluded*. You may want to create one or more place-rules to ensure that the snapshot-capable filers contain all the files that require backups. Run the affirmative form of the command, manage snapshots, to reinstate ARX snapshots on the filer.

Samples

```
bstnA(gbl-filer[nas10])# manage snapshots
```

allows the ARX to take snapshots on the "nas10" filer.

```
bstnA(gbl-filer[das1])# no manage snapshots
```

indicates that the "das1" filer does not support ARX snapshots.

Related Commands

external-filer
snapshot rule
filer-type
proxy-user (gbl-filer)
show external-filer
filer

nfs tcp connections

Purpose For filers that support NFSv3 over TCP, you may be able to improve import performance by increasing the number of TCP connections. Before you add any of the filer's shares to a namespace volume, you can use the nfs tcp connections command to increase the number of connections.

Mode gbl-filer

Security Role(s) storage-engineer or crypto-officer

Syntax **`nfs tcp connections {1 | 2 | 4 | 8 | 16 | 32}`**

1 | 2 | 4 | 8 | 16 | 32 is the number of connections to use. The ARX-1500, ARX-2500, and ARX-VE each support a maximum of 8 connections.

Default(s) 1

Guidelines By default, the ARX creates 1 TCP connection to an NFSv3/TCP filer. It establishes this connection when one of its shares is incorporated into a namespace volume (using the filer command). For some NFS filers, import speed may improve if you use multiple TCP connections. Consult your vendor documentation and F5 Support before changing this default.

The number of NFSv3/TCP connections to the filer persists as long as any of its NFS shares are used in a namespace. The namespace uses these connections for all of its communication with the filer. In case of a prior connection, this command has no effect until the next ARX reboot. To change the number of connections, use this command and then reload the ARX.

The output of the show filer connections command shows all of the current connections to a given filer.

To change the maximum CIFS connections to the filer, use cifs connection-limit.

Sample

```
bstnA(gbl-filer[nas1])# nfs tcp connections 16
```

uses 16 TCP connections to the "nas1" filer.

Related Commands

external-filer
cifs connection-limit
show filer connections
show external-filer

probe authentication

Purpose

The ARX may use Kerberos, NTLM, or NTLMv2 to authenticate with the CIFS filers behind it. Use the probe authentication command to test any of these authentication methods against any desired CIFS filer. This is useful for testing authentication methods in your network, qualifying an external filer or DC for inclusion in the ARX configuration, or for troubleshooting filers and DCs that are already used in the configuration.

Mode

priv-exec

Security Role(s)

crypto-officer or storage-engineer

Syntax: Not-Yet Configured Filer

```
probe authentication {kerberos | ntlm | ntlmv2}
  host ip-address [spn spn]
  [user user-name windows-domain domain | proxy-user proxy]
```

kerberos | ntlm | ntlmv2 chooses the authentication protocol to test.

ip-address identifies the filer or Domain Controller (DC).

spn (optional, 1-256 characters) identifies the filer by its Service-Principal Name (SPN). Use this option for a Windows cluster; use the SPN that is shared between the cluster nodes, if there is one. If there is a "$" at the end of the SPN name, omit it in this command.

user *user-name* windows-domain *domain* | proxy-user *proxy* are credentials to test. The ARX uses this identity to authenticate with the filer(s).

> **user *user-name*** (1-64 characters) and **windows-domain *domain*** (1-1024 characters) comprise a valid Windows user and domain name. If you are probing Kerberos authentication, use an FQDN for the ***domain*** name. The CLI prompts for the user's password after you enter the command.
>
> **proxy-user *proxy*** (1-32 characters) is an already-configured proxy-user with permissions to access the filer(s). Use show proxy-user for a list of proxy users. For a Kerberos probe, the windows-domain (gbl-proxy-user) for the proxy user must be an FQDN.

Syntax: Configured Filer

```
probe authentication {kerberos | ntlm | ntlmv2}
  external-filer ext-filer-name proxy-user proxy
```

external-filer *ext-filer-name* (1-64 characters) identifies an external filer that has already been configured on the ARX. For a list of configured external filers, use show external-filer.

The remaining options are explained above.

Guidelines

This command applies only to back-end shares that support CIFS.

The output contains the following fields:

IP address is the address used to connect to the filer.

SPN is the Service-Principal Name used in the CIFS connection.

User is the Windows username used for authentication.

Protocol is the authentication protocol, chosen in the command.

Status is "STATUS_SUCCESS" if the authentication succeeded, or an error otherwise.

You can use this command to test authentication methods in your Windows network. The show exports and probe exports commands are tests to qualify share(s) for inclusion in a namespace.

You can also use this command to test a configured filer that seems to be malfunctioning. The show namespace command has error messages that sometimes suggest failures at the filer; this command can help to isolate the failure(s).

Samples

```
bstnA# probe authentication ntlmv2 host 192.168.25.20 proxy-user
acoProxy2
  IP address  : 192.168.25.20
  SPN         : VM-SWP2003s2-04@MEDARCH.ORG
  User        : MEDARCH\jqpublic
  Protocol    : NTLMv2
  Status      : STATUS_SUCCESS
```

tests NTLMv2 authentication against a back-end filer at 192.168.25.20. The proxy user, "acoProxy2," authenticated successfully through NTLMv2.

```
bstnA# probe authentication kerberos external-filer fs1 proxy-user
acoProxy2
  IP address  : 192.168.25.20
  SPN         : vm-swp2003s2-04$@MEDARCH.ORG
  User        : jqpublic@MEDARCH.ORG
  Protocol    : Kerberos
  Status      : STATUS_SUCCESS
```

tests Kerberos authentication on a host that has been defined as an external filer, fs1. This uses the credentials defined for a proxy-user, "acoProxy2."

Related Commands

show exports
probe exports

probe exports

Purpose

The ARX independently migrates files from one back-end filer to another, and requires credentials with adequate permissions to do so. Use the probe exports command to test your Windows credentials (or test Unix *root* credentials) against a filer. The test attempts to write and delete a file at each of the filer's shares. This is useful for qualifying the shares for inclusion in an external filer, or for troubleshooting filer shares that are already used in a namespace.

Mode

priv-exec

Security Role(s)

crypto-officer or storage-engineer

Syntax: Not-Yet Configured Filer

```
probe exports host {ip-address | hostname} [spn spn]
  [share share-name] [user user-name windows-domain domain-name
  | proxy-user proxy]
```

ip-address* | *hostname (1-1024 characters) identifies the filer. To use a ***hostname***, the switch must be configured to perform DNS lookups; refer to the ip name-server documentation for instructions.

spn (optional, 1-256 characters) identifies the filer by its Service-Principal Name (SPN). Use this option for a Windows cluster; use the SPN that is shared between the cluster nodes, if there is one. If there is a "$" at the end of the SPN name, omit it in this command.

share-name (optional, 1-1024 characters) identifies one share on the filer.

user *user-name* windows-domain *domain-name* | proxy-user *proxy* are Windows credentials to test. If you omit these options, the command uses *root* as its identity and probes for Unix/NFS permissions. The ARX uses this identity to authenticate with the filer(s), for permission to probe the shares.

> **user *user-name*** (1-32 characters) and **windows-domain *domain-name*** (1-1024 characters) comprise a valid Windows user and domain name. The CLI prompts for the user's password after you enter the command.
>
> **proxy-user *proxy*** (1-32 characters) is an already-configured proxy-user with permissions to access the filer(s). Use show proxy-user for a list of proxy users.

Syntax: Configured Filer

```
probe exports external-filer ext-filer-name [share share-name]
  [user user-name windows-domain domain-name |
  proxy-user proxy]
```

external-filer *ext-filer-name* (1-64 characters) identifies an external filer that has already been configured on the ARX. You can use **all** to probe all configured filers. For a list of configured external filers, use show external-filer.

The remaining options are explained above.

Guidelines

The output lists all CIFS shares or NFS exports, depending on whether or not you used Windows credentials (***user-name*** or ***proxy***), and examines the access privileges at each of them. For CIFS shares, it performs a write test using the credentials of the ***user-name*** or ***proxy*** user. It also verifies that the user or proxy user has adequate permissions at the filer, equivalent to those of a Backup Operator. The results of these two tests, labeled Write and Privs, must be "OK" for the user or proxy user. For NFS shares, it performs a mount test as the Unix *root* user; the mount must succeed for NFS shares.

You can use this command to qualify a filer's shares before including them in a namespace. The show exports command runs some less-intrusive tests to qualify the share(s). For a share that passes all tests, you can use filer, a gbl-ns-vol-shr command, to include the share in a namespace.

You can also use this command to test a configured filer that seems to be malfunctioning. The show namespace command has error messages that sometimes suggest failures at the filer; this command can help to isolate the failure(s).

To test various CIFS-authentication methods against a filer, use the probe authentication command.

Samples

```
bstnA# probe exports host 192.168.25.20 share histories user jqpublic
windows-domain MEDARCH.ORG
Password: jqpasswd
```

runs a security test on one CIFS shares at 192.168.25.20. See Figure 20.1 on page 20-23 for sample output.

```
bstnA# probe exports external-filer nas1 share insurance proxy-user
acoProxy2
```

runs the same test on a host that has been defined as an external filer, nas1. This uses the credentials defined for a proxy-user, "acoProxy2." For sample output, see Figure 20.2 on page 20-24.

```
prtlndA# probe exports external-filer das-p2
```

runs an NFS-security test on another filer, das-p2. This uses *root* as its identity. For sample output, see Figure 20.3 on page 20-24.

Related Commands

show exports
probe authentication
show external-filer
filer
show namespace

Figure 20.1 *Sample Output: probe exports host ... user ... windows-domain ...*

```
bstnA# probe exports host 192.168.25.20 share histories user jqpublic windows-domain MEDARCH.ORG
Password: jqpasswd
Export probe of filer "192.168.25.20"

CIFS Credentials:
    User              jqpublic
    Windows Domain    MEDARCH.ORG

Security:
```

```
  CIFS

    Description Key: OK (success) NO (failure) -- (not applicable)

    Share                          Write   Privs
    -----------------------------  -----   -----
    histories                      OK      OK
```

Figure 20.2 *Sample Output: probe exports external-filer ... proxy-user*

```
bstnA# probe exports external-filer nas1 share insurance proxy-user acoProxy2
Export probe of filer "nas1" at 192.168.25.21

CIFS Credentials:
    User              jqpublic
    Windows Domain    MEDARCH.ORG
    Pre-Win2k         MEDARCH

Security:

  CIFS

    Description Key: OK (success) NO (failure) -- (not applicable)

    Share                          Write   Privs
    -----------------------------  -----   -----
    insurance                      OK      OK
```

Figure 20.3 *Sample Output: probe exports external-filer ... (NFS)*

```
prtlndA# probe exports external-filer das-p2
Export probe of filer "das-p2" at 192.168.74.82

Security:
  NFS
                                                 Read   Write     Space
    Path (Owner)                          Status   Size   Size   Total  Avail   FSID   Access
    --------------------------------      ------- ------ ------ ------ ------ -------- ------
    /exports                              Mounted  32 kB  32 kB 113 GB 111 GB 1f00000007
```

proxy-user (gbl-filer)

Purpose

Use this command to assign SSH, RSH, or WinRM credentials (username and password) for the current filer. The ARX uses these credentials to log into the filer's management CLI/API and run snapshot or checkpoint commands. The ARX can also use these credentials to find physical paths (as opposed to virtual, shared paths) for storage in a file-history archive.

Use the no form of the command to remove the proxy-user configuration from this filer. This prevents any ARX volume from using this filer for snapshots, or for finding physical paths on this filer for its file-tracking records.

Mode

gbl-filer

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
proxy-user name
no proxy-user
```

name (1-32 characters) is the proxy user to use for accessing the filer's CLI.

Default(s)

None

Guidelines

From gbl mode, use the proxy-user command to add a proxy-user configuration to the ARX. Use the show proxy-user command to view all configured proxy-users and their associated usernames.

◆ Note

This has no relationship to the proxy-user (gbl-ns) command, which a namespace uses for policy-driven migrations. This command supports CLI logins, and the gbl-ns command supports policy migrations. You must assign both types of proxy users to support both features.

This proxy user does not require a Windows domain for a NetApp or EMC device, where it is only used for SSH or RSH access to the Unix CLI, but it requires an FQDN (such as "mydomain.myco.com") for a Windows server. The ARX uses Kerberos to authenticate with a Windows server, and Kerberos requires FQDNs.

Guidelines: Snapshot Support

You must use this command to prepare the filer for coordinated-snapshot support with a snapshot rule. Additionally, you must use filer-type to identify the vendor of the current filer and manage snapshots to confirm that the filer supports snapshots. You can also use the ip address ... management command to specify a management IP for accessing the CLI (by default, the ARX volume uses the primary-IP address).

Guidelines: File-Tracking Support

You should also use this command for a managed volume that is tracking its files. This type of managed volume has a snapshot rule with special contents: the volume's configuration and possibly its metadata. The rule stores these regularly on an external archive. With these options in place, you can query the archive for the volume's file paths at different times. Use the show virtual path-history and show file-history virtual-service commands for these queries. These queries can tell you a file's location yesterday, when you may have performed an incremental backup on your filers, and today, when the managed volume may have migrated the file to a new filer according to your administrative policies.

To find the physical paths for display in the above queries, the snapshot rule accesses the volume's backing-filer through its CLI or API. Every filer behind the managed volume needs this assignment, or the physical paths may be blank in those queries. The volume does not have to manage snapshots, but it does need the filer-type.

Samples

```
bstnA(gbl-filer[nas10])# proxy-user nas_admin
```

uses proxy user "nas_admin" for CLI-login credentials at the "nas10" filer.

```
bstnA(gbl-filer[das2])# no proxy-user
```

removes any CLI-login credentials designated for the "das2" filer. This means that the filer cannot support coordinated snapshots.

Related Commands

proxy-user
show proxy-user
snapshot rule
filer-type
manage snapshots
ip address **... management**
file-history archive

show exports

Purpose Use the show exports command to examine back-end filers that may or may not be configured as external filers. This command is useful for qualifying a filer's shares for inclusion in an external filer, or for troubleshooting filers that are already used in a namespace.

Mode (any)

Security Role(s) operator

Syntax: Not-Yet Configured Filer

```
show exports host {ip-address | hostname} [spn spn]
  [share share-name]
  [connectivity|capabilities|shares|time|attributes|paths]
  [user user-name windows-domain domain-name |
  proxy-user proxy]
```

ip-address* | *hostname identifies the filer. (To use a *hostname*, the switch must be configured to perform DNS lookups. Refer to the ip name-server documentation for instructions.)

spn (optional, 1-256 characters) identifies the filer by its Service-Principal Name (SPN). Use this option for a Windows cluster; use the SPN that is shared between the cluster nodes, if there is one. If there is a "$" at the end of the SPN name, omit it in this command.

share-name (optional, 1-1024 characters) identifies one share on the filer.

connectivity | ... | paths selects a particular topic for examination; if you omit this, several of these topics are examined. Each topic generates a different table with its results. See the *Guidelines*, below, for details about each table.

user *user-name* windows-domain *domain-name* | proxy-user *proxy* is a required choice for either of the following tests on a CIFS filer: **shares** or **attributes**. If you omit both of these options, the command only probes for NFS services and attributes. The ARX uses this identity to authenticate with the CIFS filer(s), for permission to examine the shares.

> **user *user-name*** (1-32 characters) and **windows-domain *domain-name*** (1-1024 characters) is a valid Windows user and domain name. The CLI prompts for the user's password after you enter the command. The domain name must be an FQDN (that is, "thisdomain.thatdomain.com" as opposed to "thisdomain") to test Kerberos authentication.
>
> **proxy-user *proxy*** (1-32 characters) is an already-configured proxy-user with permissions to access the filer(s). Use show proxy-user for a list of proxy users. As above, the proxy user's domain must be fully-qualified for the command to use Kerberos authentication.

Syntax: Configured Filer

`show exports external-filer` ***`ext-filer-name`*** `[share` ***`share-name`***`] [connectivity | capabilities | shares | time | attributes | paths] [user` ***`user-name`*** `windows-domain` ***`domain-name`*** `| proxy-user` ***`proxy`***`]`

external-filer ***ext-filer-name*** identifies an external filer that has already been configured on the ARX. You can use **all** to show all configured filers. For a list of configured external filers, use show external-filer.

The remaining options are explained above.

Guidelines

This command shows several tables that indicate whether or not a filer's shares can be used in a namespace. Use this together with probe exports, which runs a write test to verify that your Windows credentials are adequate for CIFS shares.

To test various CIFS-authentication methods against a filer, such as Kerberos, use the probe authentication command.

You can use this command to qualify a filer's shares before including them in a namespace volume. For a share that passes all tests, you can use filer, a gbl-ns-vol-shr command, to include the share in a volume.

You can also use this command to test a configured filer that seems to be malfunctioning. The show namespace command has error messages that sometimes suggest failures at the filer; this command can isolate the failure(s).

If you use any Windows credentials (**user/windows-domain** or **proxy-user**), the command ignores NFS shares, capabilities, and time. If you omit Windows credentials, the command focuses exclusively on NFS.

Guidelines: Output Format

Connectivity indicates whether or not the ARX can ping the filer IP from each of its proxy-IP addresses. The ARX sends multiple pings from each proxy-IP, each with a different-sized packet, to ensure that the MTU size is consistent between the proxy IP and the filer. All pings must succeed from *all* proxy IPs for the filer to be usable. For the pings to succeed, the filer must reside on the same subnet as the proxy IPs, or be reachable through a static route (where the next-hop router is on the proxy-IP subnet). Use ip proxy-address to establish the proxy-IP subnet. Use ip route to add a static route. Additionally, all switches and network equipment between the ARX and filer must have consistent MTU settings; use jumbo mtu to set the MTU at an ARX VLAN.

CIFS Credentials are the credentials presented to the back-end filer for authentication. This is a table with three fields: User name, Windows Domain, and Pre-Win2k domain. Use the **user** and **windows-domain** options to set these, or the **proxy-user** option. The Pre-Win2k domain is only used for NTLM authentication: if you previously ran active-directory update seed-domain, this was automatically discovered; if not, and you specified a **user** in the command, this is the first part of the Domain's FQDN; if you specified a **proxy-user**, someone can explicitly set this with an option in the windows-domain (gbl-proxy-user) command. If the credentials were omitted from the command, "anonymous" appears in the User field.

Capabilities shows the NFS or CIFS services available at the filer:

> NFS lists the three services required for NFS, one per row. Each service shows the protocol and version used for transport (TCP vs. UDP), along with the port number where the service is listening.
>
> CIFS shows the SMB and/or SMB2 capabilities, supported authentication protocols, whether or not the ARX can connect to the IPC$ share, and information about the server's SPN.

Guidelines: Output Format (Cont.)

Shares lists all the NFS or CIFS shares, each in its own table:

NFS shows the results of two tests on each share, one share per row.

- Path (Owner) is the export path at the back-end filer.
- Status is "Mounted" if the ARX succeeded in mounting the export from each of its proxy-IP addresses. The share is only usable in a namespace if the mounts succeed from all of the proxy IPs. The list of proxy IPs appears in the Connectivity section of this report. If any mount fails, this shows the error returned from the filer in place of all fields in the row.
- Read Size and
- Write Size are the negotiated NFS-read size and NFS-write size.
- Space Total and Avail are the total and free space reported by the filer.
- FSID is the file-system ID reported by the filer. This identifies the back-end storage volume behind this NFS export. Any two exports with the same FSID draw from the same storage pool, so their space sizes (above) match.
- Access shows the IP address(es) or NIS group(s) that are authorized to access this share. A "*" or "(everyone)" indicates that any client machine can access the share.

CIFS displays the disk space on the share and the serial number for the share's storage volume. Any two shares with the same serial number draw from the same storage pool, so their disk-space numbers match.

Time shows the current time at the filer. The skew field at the end of the output shows the skew between the filer's clock and that of the ARX. The time settings must be nearly synchronized (preferably through NTP; see ntp server) for time-based policies to function.

Attributes shows each CIFS share and its CIFS attributes. If all of the shares behind an ARX volume support one of these attributes, the volume can support it, too. You can set each CIFS attribute for a volume with cifs access-based-enum, compressed-files, named-streams, persistent-acls, sparse-files, or unicode-on-disk, respectively. The attribute settings are either X, -, or ?.

- "X" means that the back-end share supports this attribute.
- If "-," the share does not support this CIFS attribute.
- A "?" indicates that the ARX was unable to run the test for some reason. This can occur if the Windows credentials are insufficiently privileged for the query; for example, the **user** or **proxy-user** must have Administrator-level privileges to test for ABE.

Paths is another option for CIFS shares; this shows the physical disk and path at the root of each CIFS share. Each row contains a CIFS-share name and a path. This is useful for finding and verifying share-to-subshare relationships; see the documentation for filer-subshares. This command requires Administrator-level access to read the directory path for each share; this is an increase from the usual requirements for a user or proxy-user. You can therefore use the **paths** option to verify that the user or proxy-user belongs to the Administrator group on the target filer.

Samples

`bstnA> **show exports host 192.168.25.21 user jqpublic windows-domain MEDARCH.ORG**`
`Password: **jqpasswd**`

shows export information for all CIFS shares on 192.168.25.21. See Figure 20.4 on page 20-30 for sample output.

`bstnA> **show exports host 192.168.25.19 connectivity**`

shows only the Connectivity table for another filer. See Figure 20.5 on page 20-31 for sample output.

`bstnA> **show exports external-filer fs1 capabilities proxy-user acoProxy2**`

shows the Capabilities for a CIFS filer, fs1. For sample output, see Figure 20.6 on page 20-32.

`bstnA> **show exports external-filer fs7 capabilities proxy-user acoProxy2**`

shows the Capabilities for another CIFS filer, fs7, which supports SMB2. For sample output, see Figure 20.7 on page 20-32.

`bstnA> **show exports host 192.168.25.19 time**`

shows the filer's time. See Figure 20.8 on page 20-33 for sample output.

`bstnA> **show exports host 192.168.25.29 paths proxy-user acoProxy2**`

shows the physical paths for every CIFS share at 192.168.25.29. See Figure 20.9 on page 20-33 for sample output, which shows a share and three subshares on the back-end filer.

Related Commands

show external-filer
filer
show namespace
probe exports
probe authentication

Figure 20.4 *Sample Output: show exports host ... user ... windows-domain ...*

```
bstnA> show exports host 192.168.25.21 user jqpublic windows-domain MEDARCH.ORG
Password: jqpasswd
Export probe of filer "192.168.25.21"

Connectivity:

  Slot.Proc  Proxy Address    Ping (size: result)
  ---------  ---------------  -------------------
  2.1        192.168.25.31    64: Success  2000: Success  8820: Success
  2.2        192.168.25.32    64: Success  2000: Success  8820: Success
  2.3        192.168.25.33    64: Success  2000: Success  8820: Success
  2.4        192.168.25.34    64: Success  2000: Success  8820: Success
  2.5        192.168.25.141   64: Success  2000: Success  8820: Success
  2.6        192.168.25.142   64: Success  2000: Success  8820: Success
  2.7        192.168.25.143   64: Success  2000: Success  8820: Success
  2.8        192.168.25.144   64: Success  2000: Success  8820: Success
```

```
  2.9        192.168.25.145   64: Success  2000: Success  8820: Success
  2.10       192.168.25.146   64: Success  2000: Success  8820: Success
  2.11       192.168.25.147   64: Success  2000: Success  8820: Success
  2.12       192.168.25.148   64: Success  2000: Success  8820: Success

CIFS Credentials:
    User              jqpublic
    Windows Domain    MEDARCH.ORG

Capabilities:

  CIFS
    Server            TCP/445, TCP/139
    SMB:
      Security Mode   User level, Challenge/response, Signatures disabled
      Capabilities    Large Files, Large Read, Large Write
      Max Buffer      33028
    Auth Protocols    NTLMSSP, Kerberos
    IPC$ Share        Access OK
    Auth Method Used  Kerberos
    SPN Used          ntap820$@MEDARCH.ORG
    Discovered SPN    ntap820$@MEDARCH.ORG

Shares:

  CIFS
                                              Space
    Share                                Total        Avail      Serial Num
    -------------------------------    ----------   ----------      ----------
...
    insurance                            17 GB        16 GB       a801-8ca5
...
    -------------------------------    ----------   ----------      ----------
    Share                                Total        Avail      Serial Num
                                              Space

Time:

  CIFS
    Filer's time is the same as the switch's time.
```

Figure 20.5 *Sample Output: show exports host ... connectivity*

```
bstnA> show exports host 192.168.25.19 connectivity
Export probe of filer "192.168.25.19"

Connectivity:

  Slot.Proc  Proxy Address    Ping (size: result)
  ---------  ---------------  -------------------
  2.1        192.168.25.31    64: Success  2000: Success  8820: Success
  2.2        192.168.25.32    64: Success  2000: Success  8820: Success
  2.3        192.168.25.33    64: Success  2000: Success  8820: Success
  2.4        192.168.25.34    64: Success  2000: Success  8820: Success
  2.5        192.168.25.141   64: Success  2000: Success  8820: Success
  2.6        192.168.25.142   64: Success  2000: Success  8820: Success
  2.7        192.168.25.143   64: Success  2000: Success  8820: Success
  2.8        192.168.25.144   64: Success  2000: Success  8820: Success
  2.9        192.168.25.145   64: Success  2000: Success  8820: Success
  2.10       192.168.25.146   64: Success  2000: Success  8820: Success
  2.11       192.168.25.147   64: Success  2000: Success  8820: Success
  2.12       192.168.25.148   64: Success  2000: Success  8820: Success
```

Figure 20.6 *Sample Output: show exports external-filer... capabilities ... (CIFS Filer)*

```
bstnA> show exports external-filer fs1 capabilities proxy-user acoProxy2
Export probe of filer "fs1" at 192.168.25.20

CIFS Credentials:
    User              jqpublic
    Windows Domain    MEDARCH.ORG
    Pre-Win2k Domain  MEDARCH

Capabilities:

  CIFS
    Server            TCP/445
    SMB:
      Security Mode   User level, Challenge/response, Signatures disabled
      Capabilities    Large Files, Large Read, Large Write, Info Passthru
      Max Buffer      16644
    Auth Protocols    NTLMSSP, Kerberos
    IPC$ Share        Access OK
    Auth Method Used  Kerberos
    SPN Used          vm-swp2003s2-04$@MEDARCH.ORG
    Discovered SPN    vm-swp2003s2-04$@MEDARCH.ORG
```

Figure 20.7 *Sample Output: show exports external-filer... capabilities ... (CIFS SMB2 Filer)*

```
bstnA> show exports external-filer fs7 capabilities proxy-user acoProxy2
Export probe of filer "fs7" at 192.168.25.41

CIFS Credentials:
    User              jqpublic
    Windows Domain    MEDARCH.ORG
    Pre-Win2k Domain  MEDARCH

Capabilities:

  CIFS
    Server            TCP/445
    SMB:
      Security Mode   User level, Challenge/response, Signatures disabled
      Capabilities    Large Files, Large Read, Large Write, Info Passthru
      Max Buffer      16644
    SMB2.1:
      Security Mode   Signatures optional
      Capabilities    Leasing, Large MTU
      Max Read        1048576
      Max Write       1048576
      Max Transact    1048576
    Auth Protocols    NTLMSSP, Kerberos
    IPC$ Share        Access OK
    Auth Method Used  Kerberos
    SPN Used          VM-SMB2K8R2-02@MEDARCH.ORG
    Discovered SPN    VM-SMB2K8R2-02@MEDARCH.ORG
```

Figure 20.8 *Sample Output: show exports host ... time*

```
bstnA> show exports host 192.168.25.19 time
Export probe of filer "192.168.25.19"

Time:
  NFS
    Filer's time is the same as the switch's time.
```

Figure 20.9 *Sample Output: show exports host ... paths (CIFS Filer)*

```
bstnA> show exports host 192.168.25.29 paths proxy-user acoProxy2
Export probe of filer "192.168.25.29"

CIFS Credentials:
    User              jqpublic
    Windows Domain    MEDARCH.ORG
    Pre-Win2k Domain  MEDARCH

Paths:

  CIFS

    Share                          Directory
    ---------------------------    -----------------------------------
    CELEBS$                        e:\exports\prescriptions\VIP_wing
    Y2004                          e:\exports\prescriptions\2004
    Y2005                          e:\exports\prescriptions\2005
    prescriptions                  e:\exports\prescriptions
    ---------------------------    -----------------------------------
    Share                          Directory
```

show external-filer

Purpose Use the show external-filer command to display summary or detailed information about the external filers connected to the ARX.

Mode (any)

Security Role(s) operator

Syntax

```
show external-filer
show external-filer filer-name
show external-filer all
```

If no filer name is specified, this command shows IP addresses for all external filers.

filer-name (optional; 1-64 characters) identifies a specific filer. This shows a detailed report for the filer.

all (optional) shows detailed reports for all external filers.

Guidelines: Summary The summary form shows a table with one row per external filer. The table contains three columns:

Name is set with the external-filer command.

IP Address shows the filer's address. You can reset this with ip address.

Description is the string set with the description (gbl-filer) command. The flag "(*n*th secondary)" appears here for filers with multiple IP addresses; this marks one of the filer's secondary addresses (see ip address).

Guidelines: Detailed The detailed form shows all configuration information for the external filer, followed by a series of tables that describe how the ARX is using the filer. The configuration information includes these fields:

Description is the string set with the description (gbl-filer) command.

Filer IP is/are the filer's address(es). You can use the ip address command to change this list of addresses.

Filer IP (new) only appears if someone designated a new primary address for this filer. This becomes the filer's primary IP address the next time the ARX reboots or fails over. You can use the ip address ... change-to command to change this new address.

Configured SPN only appears if someone used spn to set a service-principal name (*SPN*) for the external filer. The ARX needs the correct SPN for Windows clusters and for Kerberos support. If this appears, the ARX is always using it (not the Discovered SPN, below) for cluster connections and Kerberos authentications.

Discovered SPN is the SPN returned from a filer query. The filer query occurs when the filer's first share is enabled in a managed volume. If the Configured SPN is set, the ARX does not use this SPN in its cluster connections or Kerberos authentications. "Not Discovered" appears here if there is neither a configured SPN or a successfully-discovered SPN.

Guidelines: Detailed (Cont.)

Filer Management Proxy User identifies the Unix proxy-user credentials that the ARX uses to access the filer's CLI, where the ARX software issues snapshot commands. You can use the proxy-user (gbl-filer) command to select a different proxy user.

Filer Management IP is a separate management-IP address, for filers that support an out-of-band IP for management. The ARX logs into the filer's CLI through this address. You can use the ip address ... management command to reset it.

Filer Manage Snapshots is "Yes" if the ARX is permitted to run snapshots on the filer, or "No" otherwise. To change this setting, use the manage snapshots command.

CIFS Port is port that the ARX uses to make its CIFS connections. The ARX can automatically choose between the two well-known ports, 445 and 139, or you can set it explicitly with the cifs-port command.

CIFS Connection Limit Overall is the maximum number of CIFS/TCP connections to the filer. Use the cifs connection-limit command to set this maximum.

CIFS Connection Limit / NSM Processor is the maximum number of CIFS/TCP connections per NSM.

NFS TCP Connections is the number of NFSv3/TCP connections to the filer. This is typically 1, but you can set a higher number before you incorporate the share into a managed volume. Use the nfs tcp connections command to set this number.

These next fields only appear only if their respective values are set:

- Filer Type is "NetApp," "EMC," "Data Domain," or "Windows," set with the filer-type command. The word "(explicit)" appears after the vendor name, to indicate that the command explicitly named the vendor for this filer.
- Windows Cluster only appears if you use the cluster option in the filer-type command. If this appears, the ARX uses snapshot commands on this filer that are formatted for a Windows server 2008 (or later) cluster.
- Management Protocol is the protocol the ARX follows to access the filer's management API: SSH or RSH for NetApp, or WINRM for Windows. To change this protocol for a NetApp filer, use the management-protocol option in the filer-type command.
- NAS DB Path applies to an EMC Celerra server whose environment variable, NAS_DB, has a non-default value. The ARX requires the exact value of this variable to invoke checkpoints at the file server. Most EMC servers use the default; for those that do not, you can use an option with the filer-type command to indicate the non-default value. This field does not appear unless someone has set a non-default NAS_DB path. with the filer-type command.
- Management Port is the WinRM port that the ARX uses to manage snapshots. This presumes that the back-end filer is a Windows server. This field only appears if you set the port option in the filer-type command.

Guidelines: CIFS Filer Capabilities Table

CIFS Filer Capabilities shows the SMB and/or SMB2 capabilities for this filer. This table only appears if the filer supports SMB or SMB2.

SMB describes the supported capabilities for the original SMB protocol. This shows "not available" if the filer does not support SMB. Otherwise, one or two lines describe the filer's SMB capabilities:

Max Buffer is the negotiated size for buffer memory allocated at each end of the SMB connection. This is the negotiated maximum size of any SMB packet.

The next line is a list of discovered SMB capabilities, if any:

- Large Files,
- Large Read,
- Large Write, and/or
- Info Passthru.

SMB2 describes the supported capabilities for the SMB2 protocol. This shows "not available" if the filer did not support SMB2 at the time its first share was enabled (with enable (gbl-ns-vol-shr)); use show exports if you want to re-probe for SMB2 support. If the filer supported SMB2 when its first share was enabled, the line describes the filer's SMB2 maximums:

- Max Read,
- Max Write, and
- Max Transact.

The next line is a list of SMB2 capabilities. The line only appears if one of these capabilities is found:

- Leasing, and/or
- Large MTU.

These capabilities are accepted from the filer, but not supported by the ARX or advertised to its clients.

Has not been probed by current software appears at the bottom in the following situation: the service existed on an ARX release that did not support SMB2, the ARX was upgraded to a release that supports SMB2, and the ARX software has not yet tested the filers for SMB2 support. The probe should be complete in about 15 minutes.

Guidelines: Managed Exports Table

Managed Exports is the list of filer shares that are used in managed volumes (as opposed to direct volumes). This table only appears if one of the filer's shares is imported into a managed volume.

NFS Export or CIFS Share is the name of the export or share at the back-end filer. A CIFS filer may export "ipc$," a standard CIFS share, if an ARX namespace uses it as a sam-reference filer.

- Namespace and
- Volume identify the managed volume that uses this share.

Automatically Striped Subshares is a list of CIFS subshares that a managed volume has created on this filer. You can configure a managed volume to pass a CIFS client from a front-end subshare directly to one of its corresponding back-end subshares, where the back-end filer can subject the client to the correct share-level ACL. To support this feature, some volumes replicate, or *stripe*, subshares between their filer shares. You can use the filer-subshares command to make a volume stripe its subshares.

(restore data share) appears for a share that has been used in a restore data operation, and is not assigned to any namespace.

(file-history archive location) appears for a share that is used as a location for a file-history archive, and is not assigned to any namespace.

(contains no metadata) indicates that the share is unused. It is not assigned to any namespace, nor is it used in a file-history archive, nor has it ever been used in a restore operation.

Guidelines: Replica-snap Exports Table

Replica-snap Exports is the list of replica-snap shares, which are also used in managed volumes. This table only appears if one of the filer's shares is used as a replica-snap share.

NFS Export or CIFS Share is the name of the export or share at the back-end filer.

- Namespace and
- Volume identify the managed volume that uses this replica-snap share.

Guidelines: Directories to Ignore

Directories and/or files to ignore for importing shows all directories/files to ignore during the import. You can use the ignore-name command to identify these names.

Guidelines: Direct Mapped Exports Table

Direct Mapped Exports appears only if at least one of the filer's shares is used in a direct volume. Each share in a direct volume has one or more virtual directories, called *attach points*. Each attach point is directly mounted (or mapped) to one path on a back-end share. When a client lists files in an attach point, he sees the contents of the physical directory at the back-end share.

Backend Path is the absolute path to the root of the back-end share. This is established by the filer command in gbl-ns-vol mode.

Namespace and

Volume identify the direct volume that uses this share.

Protocols are the namespace protocols, all of which must be supported by this back-end share. Use the gbl-ns protocol command to establish these for the namespace.

Volume Path is one of the direct volume's attach-point directories. Clients see this path under the Volume path, above. The attach command creates it.

Filer-Path is the actual filer directory to which the volume is attached. This is relative to the Backend Path described earlier. The attach command attaches the Volume Path to this real path.

Guidelines: Metadata-only Shares

Metadata-only shares are back-end shares devoted to managed-volume metadata. Use the metadata share command to add one of these to a namespace or volume. This table only appears if the filer is hosting at least one metadata-only share. It shows the same information as the Managed Exports table.

Samples

`bstnA> show external-filer`

shows IP addresses for all the external filers configured for storage. See Figure 20.10 for sample output.

`bstnA> show external-filer das1`

shows details about the external-filer named "das1." See Figure 20.11 on page 20-40 for sample output.

`bstnA> show external-filer fs2`

shows details about a CIFS filer named "fs2," where a managed volume has automatically striped some filer subshares. See Figure 20.12 on page 20-40 for sample output.

`bstnA> show external-filer nas2`

shows details about the external-filer named "nas2," which has direct exports for a direct volume. See Figure 20.13 on page 20-41 for sample output.

`bstnA> show external-filer nas10`

shows details about the external-filer named "nas10," which supports SMB2 and ARX snapshots. See Figure 20.14 on page 20-42 for sample output.

`bstnA> show external-filer nas11`

shows details about the external-filer named "nas11," which is the host for replica-snap shares. See Figure 20.15 on page 20-43 for sample output.

`prtlndA> show external-filer all`

shows details about all external-filers configured at another switch, "prtlndA." See Figure 20.16 on page 20-44 for sample output.

Related Commands

external-filer
ip address
description (gbl-filer)
cifs-port
ignore-name

Figure 20.10 *Sample Output: show external-filer*

```
bstnA> show external-filer
  Name                      IP Address     Description
  ------------------------  -------------  ---------------------------
  das1                      192.168.25.19  financial data (LINUX filer, rack 14)
  fs1                       192.168.25.20  misc patient records (Table 3)
  fs2                       192.168.25.27  bulk storage server (Table 3)
  fs3                       192.168.25.28  Hematology lab server (Table 8)
  fs4                       192.168.25.29  prescription records (Table 3)
  fs5                       192.168.25.71  docs, invoices, for scanners (Table 7)
  fs6                       192.168.25.30  records of lab animals - lab rats (rack C)
  fs7                       192.168.25.41  lab animal test results - lab rats (rack C)
  das2                      192.168.25.22  Solaris filer 2 (rack 16)
  das3                      192.168.25.23  Solaris filer 3 (rack 16)
  nas1                      192.168.25.21  NAS filer 1 (rack 31)
```

```
                              192.168.25.61   (1st secondary)
                              192.168.25.62   (2nd secondary)
  das7                        192.168.25.24  Redhat-LINUX filer 1
  das8                        192.168.25.25  Redhat-LINUX filer 2
  nas2                        192.168.25.44  NAS filer 2 (rack 31)
  nas3                        192.168.25.47  NAS filer 3 (rack 32)
  nas10                       192.168.25.49  NAS filer 10 (rack 38)
  nas11                       192.168.25.48  filer 11 (rack 38)
  nasE1                       192.168.25.51  NAS filer E1
  smb1                        192.168.25.48  Samba filer
```

Figure 20.11 *Sample Output: show external-filer das1*

```
bstnA> show external-filer das1

Filer "das1" Configuration

  Description:                              financial data (LINUX filer, rack 14)
  Filer IP:                                 192.168.25.19
  Discovered SPN:                           Not discovered
  Filer Management Proxy User:
  Filer Management IP:                      192.168.25.19
  Filer Manage Snapshots:                   No
  CIFS Port:                                auto-detect
  CIFS Connection Limit Overall:            none
  CIFS Connection Limit / NSM processor:    none
  NFS TCP Connections:                      1

Managed Exports
--------------------------------------------------------------------------------

  NFS Export: /exports/budget
    Namespace: wwmed
    Volume: /acct


  Directories and/or files to ignore for importing
  ------------------------------------------------
    .snapshot
    ~snapshot
    .etc
    .ckpt*
    $RECYCLE.BIN
    System Volume Information
    SIS Common Store
    .clusterConfig
```

Figure 20.12 *Sample Output: show external-filer fs2*

```
bstnA> show external-filer fs2

Filer "fs2" Configuration

  Description:                              bulk storage server (Table 3)
  Filer IP:                                 192.168.25.27
  Discovered SPN:                           vm-swp2003s1-5$@MEDARCH.ORG
  Filer Management Proxy User:              cifs_admin
  Filer Management IP:                      192.168.25.27
  Filer Manage Snapshots:                   Yes
  CIFS Port:                                445
  CIFS Connection Limit Overall:            none
  CIFS Connection Limit / NSM processor:    none
  NFS TCP Connections:                      1
```

```
  Filer Type:                                Windows
   Management Protocol:                        WINRM
   Management Port:                            80

CIFS Filer Capabilities
-----------------------
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2: not available

Managed Exports
--------------------------------------------------------------------------------

  CIFS Share: ipc$
    Namespace: medarcv

  CIFS Share: bulkstorage
    Namespace: medarcv
    Volume: /rcrds

  CIFS Share: backlot_records
    Namespace: medarcv
    Volume: /lab_equipment


  Directories and/or files to ignore for importing
  ------------------------------------------------
    .snapshot
    ~snapshot
    .etc
    .ckpt*
    $RECYCLE.BIN
    System Volume Information
    SIS Common Store
    .clusterConfig
```

Figure 20.13 *Sample Output: show external-filer nas2 (Behind a Direct Volume)*

```
bstnA> show external-filer nas2

Filer "nas2" Configuration

  Description:                           NAS filer 2 (rack 31)
  Filer IP:                              192.168.25.44
  Discovered SPN:                        Not discovered
  Filer Management Proxy User:
  Filer Management IP:                   192.168.25.44
  Filer Manage Snapshots:                No
  CIFS Port:                             auto-detect
  CIFS Connection Limit Overall:         none
  CIFS Connection Limit / NSM processor: none
  NFS TCP Connections:                   1

  Directories and/or files to ignore for importing
  ------------------------------------------------
    .snapshot
    ~snapshot
    .etc
    .ckpt*
    $RECYCLE.BIN
    System Volume Information
```

```
    SIS Common Store
    .clusterConfig

Direct Mapped Exports

  Backend Path             Namespace     Volume               Protocols
                                             Volume Path          Filer-Path
  ------------------------------------------------------------------------------
  /vol/datavol1/direct     medco         /vol                 nfsv3-tcp
                                              vol3/sales          export
                                              vol3/mtgMinutes     mtgs
```

Figure 20.14 *Sample Output: show external-filer nas10 (Supporting SMB2 and ARX Snapshots)*

```
bstnA> show external-filer nas10

Filer "nas10" Configuration

  Description:                              NAS filer 10 (rack 38)
  Filer IP:                                 192.168.25.49
  Discovered SPN:                           enterprise$@MEDARCH.ORG
  Filer Management Proxy User:              nas_admin
  Filer Management IP:                      192.168.25.49
  Filer Manage Snapshots:                   Yes
  CIFS Port:                                445
  CIFS Connection Limit Overall:            none
  CIFS Connection Limit / NSM processor:    none
  NFS TCP Connections:                      1
  Filer Type:                               NetApp
   Management Protocol:                       RSH

CIFS Filer Capabilities
-----------------------
  SMB: Max Buffer 33028
       Large Files, Large Read, Large Write
  SMB2.002: Max Read 65536, Max Write 65536, Max Transact 65536

Managed Exports
------------------------------------------------------------------------------

  CIFS Share: equipment
    Namespace: medarcv
    Volume: /lab_equipment

  CIFS Share: for_lease
    Namespace: medarcv
    Volume: /lab_equipment


  Directories and/or files to ignore for importing
  ------------------------------------------------
    .snapshot
    ~snapshot
    .etc
    .ckpt*
    $RECYCLE.BIN
    System Volume Information
    SIS Common Store
    .clusterConfig
```

Figure 20.15 *Sample Output: show external-filer nas11 (Host for Replica-Snap Shares)*

```
bstnA> show external-filer nas11

Filer "nas11" Configuration

  Description:                              filer 11 (rack 38)
  Filer IP:                                 192.168.25.48
  Discovered SPN:                           ntap-prov$@MEDARCH.ORG
  Filer Management Proxy User:              nas_admin
  Filer Management IP:                      192.168.25.48
  Filer Manage Snapshots:                   Yes
  CIFS Port:                                445
  CIFS Connection Limit Overall:            none
  CIFS Connection Limit / NSM processor:    none
  NFS TCP Connections:                      1
  Filer Type:                               NetApp
   Management Protocol:                       RSH

CIFS Filer Capabilities
-----------------------
  SMB: Max Buffer 33028
       Large Files, Large Read, Large Write
  SMB2: not available

Replica-snap Exports
--------------------------------------------------------------------------------

  CIFS Share: equipBkup
    Namespace: medarcv
    Volume: /lab_equipment

  CIFS Share: leasedBkup
    Namespace: medarcv
    Volume: /lab_equipment


  Directories and/or files to ignore for importing
  ------------------------------------------------
    .snapshot
    ~snapshot
    .etc
    .ckpt*
    $RECYCLE.BIN
    System Volume Information
    SIS Common Store
    .clusterConfig
```

Figure 20.16 *Sample Output: show external-filer all*

```
prtlndA> show external-filer all

Filer "das-p1" Configuration

  Description:                              Debian server 1, rack 3
  Filer IP:                                 192.168.74.81
  Discovered SPN:                           Not discovered
  Filer Management Proxy User:
  Filer Management IP:                      192.168.74.81
  Filer Manage Snapshots:                   No
  CIFS Port:                                auto-detect
  CIFS Connection Limit Overall:            none
  CIFS Connection Limit / NSM processor:    none
  NFS TCP Connections:                      1

Managed Exports
--------------------------------------------------------------------------------

  NFS Export: /exports/BU
    Namespace: nemed
    Volume: /acctShdw


  Directories and/or files to ignore for importing
  ------------------------------------------------
    .snapshot
    ~snapshot
    .etc
    .ckpt*
    $RECYCLE.BIN
    System Volume Information
    SIS Common Store
    .clusterConfig


Filer "das-p2" Configuration

  Description:                              Solaris server 1, rack 3
  Filer IP:                                 192.168.74.82
  Discovered SPN:                           Not discovered
  Filer Management Proxy User:
  Filer Management IP:                      192.168.74.82
  Filer Manage Snapshots:                   No
  CIFS Port:                                auto-detect
  CIFS Connection Limit Overall:            none
  CIFS Connection Limit / NSM processor:    none
  NFS TCP Connections:                      1

Managed Exports
--------------------------------------------------------------------------------

  NFS Export: /exports/bkup
    Namespace: nemed
    Volume: /acctShdw


  Directories and/or files to ignore for importing
  ------------------------------------------------
    .snapshot
    ~snapshot
    .etc
```

```
    .ckpt*
    $RECYCLE.BIN
    System Volume Information
    SIS Common Store
    .clusterConfig


Filer "nas-p1" Configuration

  Description:                              NTAP filer 1, rack 2
  Filer IP:                                 192.168.74.89
  Filer IP (1st secondary):                 192.168.74.88
  Discovered SPN:                           Not discovered
  Filer Management Proxy User:
  Filer Management IP:                      192.168.74.89
  Filer Manage Snapshots:                   No
  CIFS Port:                                auto-detect
  CIFS Connection Limit Overall:            none
  CIFS Connection Limit / NSM processor:    none
  NFS TCP Connections:                      1

  Directories and/or files to ignore for importing
  ------------------------------------------------
    .snapshot
    ~snapshot
    .etc
    .ckpt*
    $RECYCLE.BIN
    System Volume Information
    SIS Common Store
    .clusterConfig

Metadata-only shares
--------------------------------------------------------------------------------

  NFS Export: /vol/vol2/mdata_A
    Namespace: nemed
    Volume: /acctShdw

  NFS Export: /vol/vol2/mdata_B
    Namespace: insur_bkup
    Volume: /insurShdw


Filer "nas-p2" Configuration

  Description:                              EMC filer 2, rack 7
  Filer IP:                                 192.168.74.90
  Discovered SPN:                           medarch4$@MEDARCH.ORG
  Filer Management Proxy User:
  Filer Management IP:                      192.168.74.90
  Filer Manage Snapshots:                   No
  CIFS Port:                                445
  CIFS Connection Limit Overall:            none
  CIFS Connection Limit / NSM processor:    none
  NFS TCP Connections:                      1

CIFS Filer Capabilities
-----------------------
  SMB: Max Buffer 65535
       Large Files, Large Read, Large Write, Info Passthru
  SMB2.002: Max Read 65536, Max Write 65536, Max Transact 65536
```

```
Managed Exports
--------------------------------------------------------------------------------

  NFS Export: /bkup_B
  CIFS Share: BKUP_B
    Namespace: insur_bkup
    Volume: /insurShdw


  Directories and/or files to ignore for importing
  ------------------------------------------------
    .snapshot
    ~snapshot
    .etc
    .ckpt*
    $RECYCLE.BIN
    System Volume Information
    SIS Common Store
    .clusterConfig
```

spn

Purpose

The ARX attempts to automatically discover a filer's service-principal name (*SPN*) when the filer's first share is enabled in a managed volume. A Windows server in a cluster may provide its local SPN instead of the SPN for the cluster; you can use the spn command to manually set the cluster SPN. This setting overrides any previously-discovered SPN.

The auto-discovery method does not work for a Windows 2008 cluster, so this command is required for that server type. It is also required for a Windows 2003 cluster: the ARX needs a Windows 2003 cluster's shared SPN to support cluster failovers.

Use the no form of this command to revert to the discovered SPN, if there is one.

Mode

gbl-filer

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
spn hostname@domain
no spn
```

hostname@domain (1-256 characters) is the SPN for the external filer. Do not use any "$" or "/" characters.

hostname is the filer's hostname, or the Cluster Network Name.

@ is required, and it must be between the hostname and the domain. There can only be one "@" symbol in the SPN.

domain must be in FQDN format (for example, "ourdiv.ourco.com" instead of simply "ourdiv").

Default(s)

(the discovered SPN, if there is one)

Guidelines: Windows Clusters

A Windows 2008 cluster only supports connections to a name (as opposed to an IP address). The SPN contains the name of the virtual service, and is therefore required for any Windows 2008 cluster. You must manually set the SPN for a Windows 2008 cluster, which does not provide the data necessary to perform an automatic SPN discovery. You must also set the SPN for a Windows 2003 cluster, so that the ARX can stay connected across a cluster failover.

For any cluster, this SPN must be a *virtual SPN*, one that persists after a cluster failover. Therefore, you must create a virtual SPN on the cluster before using the cluster's storage behind the ARX. This may mean joining the cluster's virtual service to the local AD domain.

Guidelines: Kerberos Support

For other server types, the correct SPN is required for Kerberos authentications. If the SPN is incorrect, a managed volume's proxy-user cannot connect and cannot import the filer's shares.

The show external-filer command shows the discovered SPN for the filer. If this is empty or incorrect, you can use the spn command to override the discovery. Once you use the spn command, the show external-filer output contains two SPN values: the discovered SPN and your configured SPN.

For an external filer without a discovered SPN, the no spn command ensures that the volume cannot use Kerberos authentication with this filer. A CLI prompt warns you of this; enter **yes** to proceed.

Guidelines: SPN Changes When Filer's Hostname Changes

If the back-end filer changes its hostname, its SPN also changes. Invoke this command again to establish the new SPN after the filer name changes. Until the SPN is updated, no managed volume can perform autonomous operations on the filer, such as importing files and migrating them to or from one of the filer's shares, and CIFS clients have limited access through Kerberos.

Samples

```
bstnA(gbl-filer[fs3])# spn fs3@medarch.org
```

sets a SPN for the "fs3" external filer.

```
bstnA(gbl-filer[fs1])# no spn
The ARX cannot discover the service principal name for the filer.
Removal of the configured service principal name may prevent use of
Kerberos authentication when connecting to the filer.

Proceed? [yes/no] yes
```

reverts to the discovered SPN for the filer.

Related Commands

filer
show external-filer

21

Namespace

character-encoding nfs

Purpose

Character encoding is the mapping between binary numbers and written characters. Some character-encoding schemes use only a single byte for each character; these typically support alphabets without any Asian characters. Multi-byte encoding schemes encompass Asian character sets. Unicode (such as UTF-8) can also use more than one byte per character, and encompasses most character sets and languages. Use the character-encoding command to set the namespace's character encoding for NFS file names.

Use no character-encoding to reset NFS file names back to the single-byte default.

Modes

gbl-ns

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
character-encoding nfs {utf-8 | shift-jis | cp932 | euc-jp
  | ksc5601 | iso-8859-1}
no character-encoding nfs
```

utf-8 specifies UTF-8 (Unicode) character encoding.

shift-jis specifies Shift_JIS (Japanese) character encoding.

cp932 is Code Page 932, or Windows-31J (Japanese) character encoding. This is the Microsoft version of Shift_JIS.

euc-jp specifies EUC-JP (Extended Unix Code - Japanese) character encoding.

ksc5601 is KSC5601 (Korean) character encoding.

iso-8859-1 is ISO 8859-1 (Latin1, single-byte) character encoding.

Default(s)

ISO 8859-1 (Latin1)

Guidelines

NFS character encoding determines the character encoding for file names. Any NFS server and client must be set for the same NFS character encoding to communicate properly. The NFS character encoding should be well-established at any site before the ARX is installed.

CIFS uses Unicode, which can support any character set, so it does not require this command.

A multi-protocol (NFS and CIFS) namespace does not allow a CIFS client to name a file or directory with characters that are not supported by NFS character encoding. If NFS names support only single-byte characters, the namespace enforces the same restriction on CIFS names. During the initial import of multi-protocol shares, the volume uses the NFS-side name of each file (possibly a filer-generated name), and renames each directory so that it is valid in NFS. (You can use no import rename-directories to prevent directory renames on import.) We recommend UTF-8 character encoding for multi-protocol namespaces and the filers behind them.

The no form of the command returns the namespace to default character encoding.

If a single front-end nfs service has exports from more than one namespace, this setting must be the same for all of the exported namespaces. The export (gbl-nfs) command exports a namespace volume through NFS.

Sample

```
bstnA(gbl-ns[wwmed])# no character-encoding nfs
```

resets the character encoding for NFS file names to the default.

Related Commands namespace

cifs anonymous-access

Purpose

Some scanners and photocopiers have a "Save As" feature that allows the client to save a copy of a scanned file onto a remote CIFS server. Many scanners and copiers make queries to the CIFS server's hidden IPC$ share before they save the file. These devices make their queries as the "anonymous" Windows-user account. To allow anonymous queries to the current namespace's IPC$ share, use the cifs anonymous-access command.

Use no cifs anonymous-access in a namespace that does not support connections to the IPC$ share.

Modes

gbl-ns

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
cifs anonymous-access
no cifs anonymous-access
```

Default(s)

```
no cifs anonymous-access
```

Guidelines

This command does not apply to an NFS-only namespace.

There are only a small set of photocopiers that require this access before they save files to the namespace. The *F5 Data Solutions Compatibility Matrix* (included with this doc set) lists all the photocopiers that are known to require this feature.

The cifs anonymous-access command permits anonymous queries to the virtual IPC$ "share," but does not permit scanners and photocopiers to actually save files to any of the namespace's volumes. As always, the scanner/copier must provide a valid Windows username and password to perform the file-save operation.

Certain CIFS-client operating systems, such as Mac OS X v10.4, also require anonymous access to make certain queries.

If a single front-end cifs service has exports from more than one namespace, this setting must be the same for all of the exported namespaces. (The export (gbl-cifs) command exports a namespace volume through CIFS.) Therefore, if this command makes the current namespace inconsistent with the other namespace(s) behind the same CIFS service, the CLI prompts with an opportunity to make the same change in the other namespace(s). Enter **yes** to allow the CLI to propagate the change to the other namespace(s).

Samples

```
provA(gbl-ns[provMed])# cifs anonymous-access
```

allows certain scanners and copiers to perform anonymous queries in the "provMed" namespace. This is a necessary first step to supporting the "Save As" feature from those scanners.

```
bstnA(gbl-ns[insur])# no cifs anonymous-access
Synchronize conflicting parameter(s) for all the namespaces exported
along with this namespace in a service and continue? [yes/no] yes
```

denies any anonymous access to the IPC$ share in the "insur" namespace. Scanners and copiers cannot save files to volumes in this namespace. The CLI also disables this option in all other namespaces behind the same CIFS service.

Related Commands

namespace

cifs authentication

Purpose

Use the cifs authentication command to enable an authentication protocol (Kerberos, NTLM, and/or NTLMv2) for the namespace's CIFS clients.

Use no cifs authentication to disable a CIFS-authentication protocol for this namespace.

Mode

gbl-ns

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
cifs authentication {kerberos | ntlm | ntlmv2}
no cifs authentication {kerberos | ntlm | ntlmv2}
```

Default(s)

all CIFS-authentication protocols are disabled by default

Guidelines

This applies only to namespaces that support CIFS.

You can run this command once for each supported protocol. A namespace can support all three authentication protocols, or any combination of them.

Each CIFS client negotiates an authentication protocol from this set. The protocol is chosen during the client-side negotiation.

CIFS clients access namespace storage through a front-end CIFS service. If a single CIFS service has exports from more than one namespace, the authentication settings must be the same for all of the exported namespaces. (The export (gbl-cifs) command exports a namespace volume through CIFS.) Therefore, if this command makes the current namespace inconsistent with the other namespace(s) behind the same CIFS service, the CLI prompts with an opportunity to make the same change in the other namespace(s). Enter **yes** to allow the CLI to propagate the change to the other namespace(s).

Guidelines: Kerberos

The namespace's front-end CIFS service must join a Windows domain to support Kerberos authentication. Use the domain-join command to enable Kerberos for a front-end cifs service.

Guidelines: NTLM and NTLMv2

The domain-join command can also enable a CIFS service for NTLM or NTLMv2. If you (or an authorized Domain Administrator) go to the domain controller (DC) and set up constrained delegation for the CIFS service, clients can authenticate with NTLM, NTLMv2, or Kerberos. No further configuration is required. In this case, the namespace software connects to its back-end filers with Kerberos, whether or not the client uses a variant of NTLM.

Guidelines: Secure Agent

A CIFS service that is configured for unconstrained delegation, or is not even joined to its domain, requires more configuration to support NTLM or NTLMv2. For these configurations, you must install an application onto one or more DCs. This DC application, called the *Secure Agent*, enables the namespace to authenticate a client once and then access storage on multiple back-end filers. See the *ARX Secure Agent Installation Guide* for instructions on installing a Secure Agent onto a DC. Then you can use ntlm-auth-server and ntlm-auth-server (gbl-ns) to connect the namespace software to the Secure Agent application.

For the namespace to support NTLMv2 as well as NTLM, you must have a version of the Secure Agent software that also supports NTLMv2. This command issues a prompt to inform you of the minimal version required.

If the namespace is configured for both NTLM and NTLMv2, it prefers NTLMv2 in its connections with back-end servers. This is true even for client connections where the client connects through NTLM.

Samples

```
bstnA(gbl-ns[ns1])# cifs authentication kerberos
bstnA(gbl-ns[ns1])# cifs authentication ntlmv2

% INFO: To use NTLMv2 with the Acopia Secure Agent, all agent instances
must support NTLMv2 (requires agent version 5.1.0 or later).
```

enables Kerberos and NTLMv2 for clients of the “ns1” namespace.

```
bstnA(gbl-ns[insur])# no cifs authentication ntlm
Synchronize conflicting parameter(s) for all the namespaces exported
along with this namespace in a service and continue? [yes/no] yes
```

disables NTLM for clients of the “insur” namespace.

Related Commands

namespace
domain-join
ntlm-auth-server

cifs filer-signatures

Purpose Use the cifs filer-signatures command to enable (or perhaps require) SMB signing between this namespace and the external filers behind it. *SMB signing* is the process of placing a digital signature into each Server Message Block (SMB) exchanged between a CIFS server (each filer) and client (the namespace software). SMB signing prevents man-in-the-middle attacks at the cost of slower performance.

Use no cifs filer-signatures to disable SMB signing between the namespace and its filers. This breaks all CIFS communication with any filers that require SMB signing.

Mode gbl-ns

Security Role(s) storage-engineer or crypto-officer

Syntax

```
cifs filer-signatures [required]
no cifs filer-signatures
```

required (optional) obligates all CIFS filers to use SMB signing in their communication with the namespace. If any of the namespace's filers refuse to support SMB signing, the namespace cannot make any CIFS connections to the filer.

Default(s) `no cifs filer-signatures`

Guidelines This applies only to namespaces that support CIFS. Use the protocol command to set the file-access protocols for the namespace.

Without SMB signing, the default for the namespace, the namespace cannot access the CIFS storage on any of its filers that require it.

If you use this command with the required option, the namespace can only connect to filers that support SMB signing. The namespace refuses to make a CIFS connection to any filer that does not support SMB signatures.

The least-restrictive setting is to enable SMB signing without requiring it (using the cifs filer-signatures syntax). The namespace can then connect to any filer, whether it requires or refuses SMB signing. If the namespace software has a choice, it prefers not to use SMB signing for performance reasons.

To control the SMB-signing policy between the ARX and its clients, you can use the signatures command in gbl-cifs mode. To see the number of filers and/or clients who have used SMB signing, along with some success and failure statistics, use the show fastpath cifs-signatures command.

Samples

```
bstnA(gbl-ns[insur])# cifs filer-signatures
```

allows SMB signing for any filer behind the “insur” namespace. If a filer requires SMB signing, the namespace will comply. Otherwise, to improve CIFS performance, the namespace negotiates for no SMB signing.

```
bstnA(gbl-ns[ns1])# cifs filer-signatures required
```

requires SMB signing for all filers behind the “ns1” namespace. If any filer behind the namespace refuses SMB signing, the namespace cannot connect to any of its CIFS shares.

```
bstnA(gbl-ns[ns4])# no cifs filer-signatures
```

disables SMB signing for the “ns4” namespace. If any filer requires SMB signing for its CIFS connections, the namespace cannot access that filer’s CIFS shares.

Related Commands

namespace
signatures
show fastpath cifs-signatures

description (gbl-ns, gbl-ns-...)

Purpose

Use the optional description command to set a descriptive string for the current namespace, volume, or share. This appears in the show namespace command.

Use the no form of the command to delete the description.

Modes

gbl-ns

gbl-ns-vol

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

`description` ***`text`***
`no description`

text (1-255 characters) is your description. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s)

`no description`

Guidelines

Either description appears in the output for show global-config namespace and show namespace.

Samples

`bstnA(gbl-ns[wwmed])# description “namespace for World-Wide Medical network”`

specifies a description for the current namespace.

`bstnA(gbl-ns-vol-shr[medarcv~/rcrds~rx])# description “prescriptions since 2002”`

specifies a description for the current ARX share, “medarcv~/rcrds~rx.”

Related Commands

namespace **->** volume **->** share
show global-config **namespace**
show namespace

enable (gbl-ns, gbl-ns-vol)

Purpose Use the enable command to activate the current namespace or volume, or its shares.

Use no enable to stop access to the current namespace or volume.

Mode gbl-ns

gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
enable [shares [take-ownership]]
no enable [shares]
```

shares (optional) causes all of the namespace's or volume's shares to be enabled.

take-ownership (optional) causes the namespace or managed volume to take ownership of all back-end shares. Use this option only if you are sure that the shares are not in active use by a managed volume on another ARX. For example, some sites use filer applications to replicate all data from one site to another; if an ARX had managed volumes at the primary site, the ARX's ownership marker (a file) would be copied to the second site. An ARX at the second site could only import the share if you use the take-ownership option. The option has no effect on a direct volume, or on any direct volumes in the namespace.

Important

This option could possibly remove a share from another managed volume that is in service. Use the take-ownership *option only for cases where some shares are spuriously marked by another ARX. The CLI prompts for confirmation if you use this option; enter* ***yes*** *to proceed.*

Default(s) Disabled

Guidelines You must enable a volume for it to be accessible by clients. When you enable a namespace, the CLI enables all volumes in the namespace.

The enable command causes a managed volume to import external files and directories into its shares. For large directory structures, this takes some time. If there are any name collisions, they occur when you issue the enable command. The import happens asynchronously; you can monitor its progress with the show namespace or show namespace status commands.

Important

For shares backed by NetApp or EMC, you may need to access the filer directly and pre-create some qtrees or EMC tree quotas. This rare configuration issue only occurs if:
- this is a managed volume,
*- you want to support both free-space quotas (*freespace cifs-quota*), and*
- you also want to support filer-subshares *in this volume.*
In this case, a NetApp share requires one qtree per subshare, and an EMC import share must be an EMC File System with one quota tree per subshare. Pre-create the NetApp qtrees and/or EMC quota trees before you enable the share. See the Guidelines: Subshare Replication with Free-Space Quotas section of the filer-subshares *documentation.*

Guidelines (Cont.)

Each volume belongs to a volume group, which shares a memory pool as well as CPU time. The volume group is associated with several resource limits that are enforced as soon as the volume is enabled. Refer to the volume-group and reserve files commands for details.

The no enable command makes the volume(s) inaccessible to clients. When a volume is disabled, client applications get no response from it. Different applications react to this in different ways; some hang, others log error messages to an internal log. The shutdown is cleaner for your clients if you first perform no export (gbl-nfs), no export (gbl-cifs), and no browsing for all NFS and/or CIFS services that export the volume.

Guidelines: Direct Volume

A *direct volume* is a collection of directory attach points that is easier to configure than a managed volume but does not offer any policy features. Each attach-point directory in the direct volume is attached to an actual directory on a back-end filer. A direct volume keeps no metadata. You use the direct command to declare a volume to be a direct volume.

The enable command does not trigger an import in a direct volume, since there is no metadata to construct. The enable operation is therefore much faster. Also, the take-ownership option has no effect on a direct volume.

Guidelines: Sync Reports

A CIFS volume with filer-subshares or cifs access-based-enum enabled performs some additional processing during import. Specifically, the volume software discovers CIFS subshares (shares under the imported shares) and ABE settings, and it makes these settings consistent at every back-end filer. This process produces a report to show its results, named "syncSshrNewStorageReport_*timestamp*.rpt." You can use show reports to get a list of reports, and show reports *report-name* to read a particular report.

If the sync-subshare process fails for any reason, you can use the report to diagnose the problem. You can then rerun the sync process with the sync subshares from-namespace or sync subshares from-namespace commands.

Samples

```
bstnA(gbl-ns-vol[archives~/eng])# enable
```

enables the current volume, “archives~/eng.”

```
bstnA(gbl-ns[ns])# enable shares take-ownership
This command allows the switch to virtualize shares that are used by
other Acopia switches.
Allow switch to take ownership of all the shares in this namespace?
[yes/no] yes
```

enables all shares in the current namespace, “ns.” For all managed volumes in the namespace, this takes ownership of any filer shares that appear to be managed by another ARX.

```
bstnA(gbl-ns-vol[unused-ns~/vol2])# no enable
```

disables the “unused-ns~/vol2” volume. Clients can no longer access this volume.

Related Commands

namespace
show namespace

metadata cache-size

Purpose Use the metadata cache-size command to specify how much cache memory (in MB) to reserve for managed-volume metadata. This size is used for each volume-group domain used by the current namespace.

Use the no form of this command to revert to the default cache size.

Mode gbl-ns

Security Role(s) storage-engineer or crypto-officer

Syntax

```
metadata cache-size cache-size
no metadata cache-size
```

cache-size (64-1750) is the size, in megabytes (MB), of memory to reserve for the current namespace's metadata cache(s).

Default(s) 512

Guidelines Use this command only under the guidance of authorized F5 personnel.

The total cache requirement is the number of simultaneous file accesses multiplied by the amount of memory required for each file (approximately 300 bytes at press time). The cache size should therefore be small if users access only a small segment of files at a given time, and it should be large if clients access many different files. The proper size for your site can vary widely depending on a number of configuration settings (number of shares per volume, front-end-service configuration, and so forth). In most cases, this can only be determined by experimentation and close monitoring by trained F5 personnel.

Each of the namespace's volume groups has a separate memory cache, shared by all managed volumes assigned to the domain. Use the volume-group command to assign a volume to a group.

On the ARX-500, you can increase the maximum number of volume groups. The memory and CPU resources are more scarce on this platform than its newer counterparts, so it is shipped with a lower volume-group maximum. If you increase the maximum number of volume groups (using the max-volume-groups command), you should also consult with F5 to change this metadata cache size. The total memory used for this namespace is the total number of volume groups times this cache size. All platforms can use swap space (from their internal disks) to add to memory size if memory gets low. You can use the show processors usage command to monitor memory and swap-space usage on the system.

You must disable and re-enable the namespace for this command to take effect; use the [no] enable (gbl-ns, gbl-ns-vol) command in gbl-ns mode.

Sample

```
bstnA(gbl-ns[wwmed])# metadata cache-size 500
```

allocates 500 MB cache memory for each volume group used by the "wwmed" namespace.

Related Commands

enable (gbl-ns-vol-shr)
volume-group
max-volume-groups
show global-config **namespace**
show processors usage

namespace

Purpose

Use the namespace command to create a new namespace, or edit an existing one.

Use the no form of the command to delete a namespace.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

namespace *name*
no namespace *name*

name (1-30 characters) is a name you choose for the namespace. The name "all" is reserved and cannot be used.

Default(s)

None

Guidelines

The CLI prompts for confirmation before creating a namespace; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new objects.)

This places you in gbl-ns mode, where you must establish one or more managed volumes and/or direct volumes for the namespace. Each managed volume is like a file system in the namespace; it is composed of files and directories from various back-end filers. A direct volume contains shares with *attach points*, which are analogous to NFS mount points and mapped CIFS shares. A managed volume contains metadata and supports policy rules, a direct volume does not. Use the volume command to create either type of volume.

From gbl-ns mode, you must also set the file-access protocol (NFSv2, NFSv3 (over UDP), and/or CIFS), and you must configure any security parameters to properly authenticate clients. Use the protocol command to set the file-access protocol(s). Use the enable (gbl-ns, gbl-ns-vol) command to enable the namespace.

Guidelines: no namespace

You must remove all of the namespace's volumes before you can remove the namespace with no namespace. Removing a volume is a complex process, described in the documentation for the volume command. The remove namespace command removes all volumes for you; best practices dictate that you use that command instead. The remove service command removes the namespace and all other configuration objects that are exclusively dedicated to the namespace, such as external filers and global servers.

Samples

```
bstnA(gbl)# namespace ns
This will create a new namespace.

Create namespace 'ns'? [yes/no] yes
bstnA(gbl-ns[ns])#
```

creates a namespace called "ns."

```
bstnA(gbl)# no namespace myNameSpace
```

deletes the "myNameSpace" namespace.

Related Commands

volume
protocol
enable (gbl-ns, gbl-ns-vol)
volume-group
show namespace
show global-config **namespace**
remove namespace
remove service

ntlm-auth-db (gbl-ns)

Purpose *This command is only necessary behind a* cifs *service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the* domain-join *documentation, and avoid this CLI command.*

Use the ntlm-auth-db command to assign an NTLM-authentication database to the current namespace.

Use the no form of the command to remove an NTLM authentication database.

Mode gbl-ns

Security Role(s) storage-engineer or crypto-officer

Syntax **`ntlm-auth-db`** ***`name`***
`no ntlm-auth-db` ***`name`***

name (1-64 characters) is the name of an NTLM authentication database. Use the show ntlm-auth-db command for a list of configured NTLM databases.

Default(s) None

Guidelines This feature is designed for demonstrations with limited CIFS-service offerings. This is for CIFS services that either do not support Kerberos or support Kerberos with unconstrained delegation. For a namespace behind a non-Kerberos service, or a service with unconstrained delegation, the ntlm-auth-server is designed for long-term use.

Use this command only under the guidance of F5 personnel.

This command assigns an NTLM authentication database to a namespace. One database can be used in multiple namespaces. Use the gbl-mode ntlm-auth-db command to create an NTLM authentication demo database.

You must also use cifs authentication ntlm and/or cifs authentication ntlmv2 for the namespace to support NTLM or NTLMv2 authentication.

If a single front-end cifs service has exports from more than one namespace, this group of NTLM-Authentication DBs must be the same for all of the exported namespaces. (The export (gbl-cifs) command exports a namespace volume through CIFS.) If this command makes the current namespace inconsistent with the other namespace(s) behind the same CIFS service, the CLI prompts with an opportunity to make the same change in the other namespace(s). Enter **yes** to allow the CLI to propagate the change to the other namespace(s).

Use the show ntlm-auth-db command to display a list of configured NTLM authentication databases.

Sample

```
bstnA(gbl-ns[ns1])# ntlm-auth-db DEMO
```

assigns the NTLM authentication database, “DEMO,” to the current namespace.

Related Commands namespace
cifs authentication
ntlm-auth-db
show ntlm-auth-db

ntlm-auth-server (gbl-ns)

Purpose

This command is only necessary behind a cifs *service that uses unconstrained delegation (or is not joined to its domain). Best practices dictate that you use constrained delegation, as described in the* domain-join *documentation, and avoid this CLI command.*

Use the ntlm-auth-server command to assign an ARX Secure Agent (ASA) server to the current namespace. To support clients from multiple Windows domains, you can use this command multiple times.

Use the no form of the command to remove an ASA from the namespace. This prevents clients in the ASA's Windows Domain from using NTLM or NTLMv2 authentication.

Mode

gbl-ns

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
ntlm-auth-server name
no ntlm-auth-server name
```

name (1-128 characters) is the name of an ASA.

Default(s)

None

Guidelines

Use this command to assign an ARX Secure Agent (ASA) server to a namespace. The ASA facilitates NTLM and/or NTLMv2 authentication for the namespace's clients. Each ASA runs on a Windows-Domain Controller (DC) for a single Windows Domain. See the *ARX Secure Agent Installation Guide* for more information.

Install ASA Release 5.1.0 or later. If your DC is already running an older release, upgrade it to 5.1.0 or later. You can use the show ntlm-auth-server detailed command to find the release level of all ASAs used on this system.

After you install the ASA on a DC, use the ntlm-auth-server command to configure it on the ARX. Then use this command to make it available to the current namespace. Use the show ntlm-auth-server command to display the configured ASAs along with connection statistics.

◆ Note

If the namespace's front-end CIFS Service uses ***constrained delegation****, introduced with Windows Server 2003, the Secure Agent (and this command) is unnecessary. An administrator with "Domain Administrator" privileges can go to the DC and configure constrained delegation for this CIFS service. You can use the* probe delegate-to *command to find all back-end filers behind a CIFS service, which the DC administrator adds to the CIFS service's "delegate-to" list. Once the CIFS service is set up this way at the DC, the service's NTLM/NTLMv2 clients can authenticate without help from a Secure Agent.*

Guidelines: Multi-Domain Support and Redundant ASAs

A single namespace can support NTLM/NTLMv2 for clients from multiple Windows domains. For each supported domain, install an ASA on at least one DC.

Additionally, you can assign multiple ASAs from the same domain to a namespace, for redundant coverage. If the namespace has two or more ASAs for the same domain, it selects the ASA running the highest version of software.

Guidelines: Co-Existence with Kerberos

A namespace can authenticate its clients with NTLM, NTLMv2, and Kerberos concurrently. This facilitates network transitions from NTLM to Kerberos. Use cifs authentication to configure NTLM, NTLMv2, Kerberos, or any combination of those authentication protocols. Each client chooses a protocol from that set.

Guidelines: Same Set of ASAs for All Namespaces Behind a CIFS Service

If a single front-end cifs service has exports from more than one namespace, this group of ASAs must be the same for all of the exported namespaces. (The export (gbl-cifs) command exports a namespace volume through CIFS.)

If this command makes the current namespace inconsistent with the other namespace(s) behind the same CIFS service, the CLI prompts with an opportunity to make the same change in the other namespace(s). Enter **yes** to allow the CLI to propagate the change to the other namespace(s).

Guidelines: Removing an ASA

Removing an ASA (with no ntlm-auth-server) may stop NTLM and NTLMv2 support for the ASA's Windows domain; if this is the only ASA for the domain, clients from that domain can no longer authenticate through NTLM. This may also affect Kerberos clients in the domain if they drop down to NTLMv2 or NTLM. You can add a new ASA to the namespace at any time to support NTLM and/or NTLMv2 in a new Windows domain.

This is not an issue if the cifs services in from of the namespace all support constrained delegation. You can use show cifs-service all to see all CIFS services, the namespace behind each service, and whether or not the service supports constrained delegation.

Guidelines: An Alternative NTLM-Authentication Option

For demonstration purposes, there is an alternative to configuring an external ASA: you can use ntlm-auth-db to map a small group of clients to a single set of valid NTLM credentials. The switch uses the back-end credentials to authenticate to CIFS filers.

Sample

```
bstnA(gbl-ns[insur])# ntlm-auth-server server1
NFS and CIFS services may export volume paths from multiple namespaces
subject to certain restrictions. Parameter 'NTLM auth servers' for
namespace 'insur' conflicts with that for namespace 'medarcv' in
service 'ac1.medarch.org'. Synchronize conflicting parameter(s) for all
the namespaces exported along with this namespace in a service and
continue? [yes/no] yes
```

assigns the ASA, "server1," to the current namespace.

Related Commands

namespace
show ntlm-auth-server
cifs authentication
ntlm-auth-server

protocol

Purpose Use the protocol command to establish the protocol for accessing the files in this namespace (NFSv2UDP, NFSv3UDP, CIFS, or NFSv3TCP).

Mode gbl-ns

Security Role(s) storage-engineer or crypto-officer

Syntax

```
protocol {nfs2 | nfs3 | cifs | nfs3tcp}
no protocol {nfs2 | nfs3 | cifs | nfs3tcp}
```

nfs2 | nfs3 | cifs | nfs3tcp is a required choice:

nfs2 is NFS version 2 over UDP,

nfs3 is NFS version 3 over UDP,

cifs is CIFS, and

nfs3tcp is NFS version 3 over TCP.

Default(s) None

Guidelines You can repeat this command to configure multiple protocols for the namespace.

All of the namespace's shares must support all of the file-access protocols you set for the namespace. For example, if you set **nfs2** for the namespace, all of the shares in this namespace must support NFSv2.

After the first share is imported, the CLI only allows protocol changes that do not affect client service. Any of the following changes interrupt service:

- converting a single-protocol namespace to a multi-protocol namespace,
- removing CIFS,
- removing NFS2, or
- removing both of NFS3 and NFS3/TCP.

To make any of the above changes after the volume is imported, you must use nsck ... destage, change the protocol with this command, then bring the volume back online with enable (gbl-ns, gbl-ns-vol).

For other protocol changes, such as adding NFSv2 to a running NFSv3 namespace, you must disable the volume (no enable (gbl-ns, gbl-ns-vol)) and its front-end service (no enable (gbl-cifs, gbl-nfs)), change the protocol, then re-enable both. This causes a shorter service outage.

A namespace cannot support both **nfs2** and **nfs3tcp** unless it also supports **nfs3**.

Samples

```
bstnA(gbl-ns[ns])# protocol nfs2
```

sets the current namespace to support NFSv2 UDP.

```
bstnA(gbl-ns[ns1])# protocol cifs
```

sets the current namespace to support CIFS shares.

Related Commands namespace

proxy-user (gbl-ns)

Purpose Use this command to assign Windows credentials (username, password, and Windows domain) to the current namespace. The namespace's managed volumes use these credentials to import CIFS shares, periodically check their health, and migrate files between them.

Use the no form of the command to remove a proxy-user configuration from this namespace.

Mode gbl-ns

Security Role(s) storage-engineer or crypto-officer

Syntax

```
proxy-user name
no proxy-user
```

name (1-32 characters) is the Windows proxy user to associate with the current namespace.

Default(s) None

Guidelines From gbl mode, use the proxy-user command to add a proxy-user configuration to the ARX.

If the proxy user has an FQDN for its windows-domain (gbl-proxy-user), it uses Kerberos to authenticate with back-end filers. If it fails to get a Kerberos ticket, it drops down to NTLMv2, and then (if that fails, too) NTLM.

Use the show proxy-user command to view all configured proxy-users and their associated usernames and Windows domains.

◆ Note

This has no relationship to the proxy-user (gbl-filer) command, which a snapshot-supporting volume uses to log into the filer's CLI. This command supports CIFS imports and policy migrations, and the gbl-filer command supports coordinated snapshots. You must assign both types of proxy users to support both features.

Sample

```
bstnA(gbl-ns[medarcv])# proxy-user acoProxy2
```

associates proxy user "acoProxy2" with the current namespace.

Related Commands

proxy-user
show global-config namespace
show proxy-user

remove namespace

Purpose Use the remove namespace command to remove a namespace or volume, along with all of its associated metadata and front-end exports.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax `remove namespace` ***`name`*** `[volume` ***`volume`***`] [timeout` ***`seconds`***`] [sync]`

name (1-30 characters) identifies the namespace.

volume (optional, 1-1024 characters) focuses on a single volume.

seconds (optional, 300-10,000) sets a time limit for the removal of each namespace (or volume) component.

sync (optional) shows the operation's progress at the command line. With this option, the CLI prompt does not return until all components have been removed.

Default(s) None.

Guidelines The CLI prompts for confirmation before removing any configuration objects or metadata; enter **yes** to proceed.

This removes the following objects from a namespace or volume:

- all policy (rules, share farms, and filesets)
- all shares and metadata shares
- all front-end exports that reference the volume(s)
- all metadata behind the volume(s)

By default, this command generates a report to show all of the actions it takes to remove the volume(s), in order. The CLI shows the report name after you issue the command, and then returns. You can enter CLI commands as the namespace software removes the objects in the background. Use tail to follow the report as it is written. Use show reports file-name to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands. Use the sync option to send the status to the command line instead; the command does not generate a report if you use the sync option.

To remove a namespace and all other configuration objects dedicated to the namespace (including global servers and external filers), use remove service. To remove a share from a volume, use remove-share migrate or remove-share nomigrate. Use the remove namespace ... policy-only command to remove all policy objects (rules, share farms, and/or filesets) from a namespace or volume. The remove namespace ... volume ... exports-only command finds all front-end exports for a volume and removes them, leaving the volume itself intact.

Sample

```
prtlndA# remove namespace insur_bkup sync

Remove namespace 'insur_bkup'? [yes/no] yes
% INFO: Removing service configuration for namespace insur_bkup
% INFO: Removing CIFS browsing for namespace insur_bkup
% INFO: Removing volume policies for namespace insur_bkup
% INFO: destroy policy insur_bkup /insurShdw
% INFO: Removing shares for namespace insur_bkup
% INFO: no share backInsur
% INFO: Removing volume metadata shares for namespace insur_bkup
% INFO: no metadata share nas-p1 path /vol/vol1/mdata_B
% INFO: Removing volumes for namespace insur_bkup
% INFO: Removing NFS services for namespace insur_bkup
% INFO: Removing CIFS services for namespace insur_bkup
% INFO: no volume /insurShdw
% INFO: Removing namespace metadata shares for namespace insur_bkup
% INFO: Removing namespace insur_bkup
% INFO: no namespace insur_bkup
```

removes the “insur_bkup” namespace, showing progress on the command line instead of in a report.

Related Commands

remove service
remove-share migrate
remove-share nomigrate
remove namespace ... policy-only
remove namespace ... volume ... exports-only
wait-for remove

sam-reference

Purpose

A properly-empowered CIFS client can edit the list of users and groups who are permitted to access a file. For example, the owner of a file can allow read-access for a trusted group. A Windows server keeps a list of all available groups in its Security Account Management (SAM) database; a namespace proxies all of its SAM requests to one of the CIFS servers behind it. For installations that use Local Groups on some filers, it is possible for the switch to choose a filer that is unaware of all groups in the namespace. You can use the sam-reference command to choose the filer manually. This command is unnecessary unless at least one of the filers uses Local Groups.

Use no sam-reference to allow the namespace to choose the SAM-reference filer. This is appropriate to an installation that only uses globally-defined groups.

Mode

gbl-ns

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
sam-reference ext-filer [cluster cluster-name]
no sam-reference
```

ext-filer (1-64 characters) identifies the external filer to use. For a list of configured external filers, use show external-filer.

cluster-name (optional, 1-64 characters) is only relevant if the ARX is part of a disaster-recovery (DR) configuration. In a DR configuration, there is an active ARX cluster with one set of filers and a backup cluster with a mirrored set of filers. This determines which cluster uses this *ext-filer* for its SAM queries. Run the sam-reference command twice per namespace if you use DR: once to designate the SAM-reference filer for the active cluster, and again to determine the SAM-reference filer at the backup cluster. Use show cluster for a list of configured clusters. If you omit this option, the command assumes that this is the local cluster.

Default(s)

The switch chooses one of the namespace's filers as a SAM reference.

Guidelines

This command does not apply to an NFS-only namespace. This only affects CIFS clients who want to add group access to one of their files. The namespace must query a SAM database on a back-end filer to offer the list of available groups and users.

Each CIFS filer has its own SAM database, with its own names for groups and users. The namespace chooses one filer for all of its SAM queries. For installations that use Global Groups, this is not an issue. For installations that use Local Groups at one or more filers, it is possible for the namespace to choose a filer that does not support all possible groups.

◆ Note

SAM-database queries are not associated in any way with the file (per the standard CIFS and SAM protocols), so the client does not provide the information to determine which CIFS filer should respond to each query. The namespace therefore chooses one CIFS filer for all SAM queries.

Guidelines (Cont.)

For installations where Local Groups differ amongst the CIFS filers, one of the filers must define all of the groups for the namespace. If necessary, you may need to define some Local Groups at one of the back-end filers to create a superset. Use this command to identify that filer.

Conversely, this command is not needed if every filer behind the namespace supports the same set of local groups. This is a recommended configuration.

If a single front-end cifs service has exports from more than one namespace, this SAM-reference filer must be the same for all of the exported namespaces. By extension, the single SAM-reference filer must define all the Local Groups behind all of the namespaces behind the CIFS service. (The export (gbl-cifs) command exports a namespace volume through CIFS.)

The currently chosen SAM-reference filer (if any) appears in the output of the show global-config namespace command.

Samples

```
bstnA(gbl-ns[medarcv])# sam-reference fs2
```

establishes “fs2” as the SAM-reference filer for the “medarcv” namespace.

```
bstnA(gbl-ns[medarcv])# no sam-reference
```

allows the namespace to choose any CIFS filer behind the namespace as its SAM reference.

Related Commands

namespace
show sid-translation
show global-config **namespace**
show cluster

show namespace

Purpose Use the show namespace command to show summaries of all namespaces, or to include a single namespace name to view configuration details for that namespace.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show namespace
show namespace name [volume vol-path [share share-name]]
show namespace all
```

name (optional, 1-30 characters) identifies the namespace. If you omit this, the output lists all namespaces on the switch.

vol-path (optional, 1-1024 characters) is the name of the volume.

share-name (optional, 1-64 characters) is the name of the share.

all lists details for all configured namespaces.

Guidelines This shows the full configuration and status of a namespace. For status only (to monitor the progress of an import), you can use the smaller show namespace status command.

The show namespace output displays the following fields:

Description is set with the description (gbl-ns, gbl-ns-...) command.

Metadata Cache Size is the internal cache size in MB (512 MB is the default). To change the cache size, use the metadata cache-size command.

Proxy User is the username/password used by the ARX to access back-end CIFS shares. Use the proxy-user (gbl-ns) command to set this.

Filer SMB Signatures describes this namespace's setting for SMB signing, a CIFS security feature. This is Enabled (use SMB signing if and only if the filer requires it), Required (only connect to CIFS filers that agree to use SMB signing), or Not Enabled (only connect to filers that do *not* require SMB signing). You can change this with the cifs filer-signatures command.

NFS Character Encoding only appears in a namespace that supports NFS. This shows the character encoding used for all file and directory names sent to NFS clients. You can change this with the character-encoding nfs command.

SAM Reference Filer only appears if explicitly set with the sam-reference command. This shows the CIFS filer used to answer all CIFS-client queries to the Security Account Management (SAM) database.

Supported Protocols lists the protocols (NFSv2, CIFS, and so on) supported by the namespace. This is set with the protocol command.

Guidelines (Cont.)

CIFS Authentication appears only for a namespace that supports CIFS:

Protocols is a list of the authentication used (NTLM, NTLMv2, and/or Kerberos). Use cifs authentication to change this setting.

- Anonymous Access (IPC$ only) may also appear in this list. This indicates that CIFS clients can access the namespace's IPC$ volume with the "anonymous" Windows-user account; some scanners/photocopiers require this access for saving files to the namespace's CIFS volumes. You can use the cifs anonymous-access command to change this setting.
- None indicates that CIFS clients cannot authenticate to this namespace, and therefore cannot access any of the namespace's storage.

NTLM Servers (only appears if configured) is a list of external authentications servers, set by ntlm-auth-server (gbl-ns).

NTLM Database (only appears if configured) is set by ntlm-auth-db (gbl-ns).

Participating Switches is a list of ARXes that support the namespace. For each switch, the volume groups are shown in parenthesis; this is where the namespace software runs. Before you enable a volume (not after), you can specify its volume group with the volume-group command.

Guidelines: Metadata Shares (Namespace)

Metadata shares is shown only if the namespace uses a metadata-only share. Use the metadata share command to use a metadata-only share for the namespace. This is a table with the following columns (one row per configured share):

Filer is the IP address or DNS name for the filer,

Backend Path is the filer's share name or NFS-export path,

Contains Metadata is "yes" for the one export that holds metadata, and

Status is the current status of the share. For possible status values, see *Guidelines: Import Status* below.

Guidelines: Windows Management Authorization

Windows Management Authorization Policies appears for namespaces that support CIFS and management by authorized Windows clients. The Microsoft Management Console (MMC) bundled with Windows is a typical interface for remote management. The windows-mgmt-auth (gbl-ns) command assigns a group of authorized clients to a namespace. If any Windows-management-authorization (WMA) groups are assigned to the namespace, they are listed in this table.

Guidelines: Volumes

Volumes is a list of managed volumes in the namespace. Use the volume command to create a volume.

CIFS: is a list of supported CIFS options, if applicable. The following commands set these options: compressed-files, named-streams, persistent-acls, sparse-files, unicode-on-disk, cifs case-sensitive, and cifs file-system-name.

Volume freespace is the amount of free space advertised to this volume's clients. After the free-space amount is the calculation method for free space: "automatic," or "manual," or "clients use dir-master-only."

- automatic - Add up the free space from all shares behind the volume and present the sum to all clients. The namespace volume monitors the back-end-storage ID behind each share; if two or more shares have matching storage IDs (that is, they draw their free space from the same storage pool), only one of the shares has its free space counted. This is the default setting for a new volume.
- manual - Add up the free space from all shares, ignoring any shares that draw from the same back-end-storage pool. You can manually eliminate shares from this sum with freespace ignore. You can also adjust the free space that is counted for each share with freespace adjust. Use freespace calculation manual to set this.
- clients use dir-master-only - Give the client the free space for the *master* back-end share only, without adding any any free space from any of the *stripe* shares in the same volume. The volume chooses the master share when the client connects to the volume through a front-end share or subshare: the volume starts with the root directory of that front-end share, finds the master instance of that directory, and chooses the back-end share that hosts that master directory. This is useful for volumes where directory trees reside only on a single back-end filer most of the time. Use the freespace calculation dir-master-only command to set this.

Volume total space is the sum of total space on all shares behind the volume. This is the actual space. The volume's clients may see different space settings, as determined by the settings above.

CIFS quotas is Enabled if the volume supports path-based quotas on its back-end filers. This indicates that the volume shows free-space values based on these back-end quotas; each client sees his or her space quota, not the entire size of the volume. You can use the freespace cifs-quota command to enable or disable this feature.

Auto Sync Files: Enabled appears only if the volume is permitted to synchronize metadata that is found to be incorrect. This means that certain client-access errors will trigger an auto-sync operation. Use the auto sync files command to enable or disable this feature.

Guidelines: Volumes (Cont.)

Auto Sync Options: Rename-Files appears only if the volume is permitted to rename files when auto-synchronizing. That is, if auto-sync discovers a file whose name is the same as an already-imported file, this feature allows the volume to rename the newly-discovered file. As above, use the auto sync files command to enable or disable this feature.

Metadata size is the amount of metadata space allotted to this volume.

Metadata free space is the amount of free space left on this volume's metadata share. This does not appear if metadata is stored on the internal disks.

Filer Subshares: Enabled appears only if the filer-subshares flag is enabled. This applies only to managed volumes that support CIFS; it indicates that the volume can pass a client from a CIFS subshare at the front-end service through to the corresponding subshare at the back-end filer. This pass-through mechanism makes it possible to support share-level ACLs for the volume's subshares. The additional text, native-names-only, appears if the filer-subshares command was invoked with a flag of the same name.

Oplock support appears only for volumes that support CIFS. This shows whether or not this volume supports CIFS opportunistic locks (oplocks). This can be Enabled, Disabled, or Automatic, as set by the cifs oplocks-disable command.

Notify-change mode shows the degree of support for the CIFS *change notification* feature. This is "Normal" (tell clients only of changes in the top level of the directory requested, ignoring any request to track changes in its subtree), "Use Subtree Flag" (inform CIFS clients of all back-end file system changes that they request), or "No changes sent" (silently ignore all client requests for change notification). You can change this with cifs notify-change-mode.

CIFS path cache only appears for volumes that support CIFS. This is "Enabled" or "Not Enabled," depending on the cifs path-cache setting. If enabled, NSM processors keep a cache of the volume's CIFS paths as it learns them. This prevents repetitive queries to the volume process on an ACM processor. Otherwise, NSM processors query the volume process for each CIFS-client request.

CIFS access based enum also appears only for volumes that support CIFS. This is "Enabled" or "Not Enabled," possibly followed by some combination of "Auto-Enable," "Ignore Metadata Skew," and/or "Full Autosync." This output depends on the cifs access-based-enum setting. If enabled, the volume changes its behavior for sending directory listings to its CIFS clients; if a back-end share omits some files or directories, the volume assumes that the omissions are caused by ABE, and makes no attempt at correction. Additionally, the volume only allows a share to import if the share has ABE enabled at its back-end filer; the volume rejects shares with ABE disabled, or automatically enables ABE if the "Auto-enable" flag is raised. If "Not Enabled," the volume presumes that none of its filers have ABE enabled, and therefore may amend a directory listing with omissions.

Snapshots is either "Enabled" or "Not Enabled," depending on whether or not the volume contains at least one snapshot rule.

Migration method is either "Staged" or "Direct." This indicates the method for migrating files; either performing a network transfer to a staging area first, or performing the transfer directly to the final location. The direct method fails if the migration is interrupted by a snapshot. You can use the policy migrate-method command to change this setting. "Direct (default)" indicates that the setting is "Direct" only because that is the default, not because anyone set it.

State is usually Enabled or Disabled, as set by the enable (gbl-ns, gbl-ns-vol) command. This cycles through several stages during import. The namespace *imports* the files and directories from back-end storage when you create and enable the volume and its share(s). See *Guidelines: Volume State*, below, for an explanation of all possible volume states.

Host Switch is the ARX with the volume-group where this volume resides. This is typically the ARX peer where the volume was originally created. If you want to use show volume-group to show this volume's group, run the command on the ARX shown in this field.

Instance is an integer ID for the volume processes.

Volume Group is the volume group where this volume's processes run. Several volumes from the same namespace can run in a given volume group. The volume group is automatically assigned when the volume is enabled. You can optionally use the volume-group command to pre-set the volume group beforehand.

Processor shows the physical CPU where the volume group is currently running. A volume group can fail over between peers in a redundant pair. The processor appears in *slot.processor* format.

Files is the number of files in the volume, and the file credits that are remaining for the volume. This also shows the maximum possible number of file credits if auto reserve files is set. The auto-reserve feature adds file credits as the volume grows.

The volume's metadata shares appear next.

Guidelines: Metadata Shares (Volume)

Metadata shares is shown only if the volume uses a metadata-only share. Use the metadata share command to use a metadata-only share for the volume. This is a table with the same columns as shown above in *Guidelines: Metadata Shares (Namespace)*.

Guidelines: Shares

Share name (s) are the configured shares in the volume. Use the share command to add a share to the volume.

Note

Instead of using show namespace *for this information, you can use the shorter* show namespace status *or* show share status *command to display share, filer, path, and import status only.*

Guidelines: Shares (Cont.)

If the share has been designated as a replica-snap share, a replica of one of the other shares that is dedicated to snapshots, [replica-snap] appears next to the share name. Many of the detailed fields (below) do not appear for replica-snap shares because they are irrelevant to snapshots.

Description appears only if someone set a description (gbl-ns, gbl-ns-...) for the share.

Filer is the filer that hosts the back-end share. This is set by the filer command. If [Acopia Namespace] appears after the filer name, the filer is a managed-volume in a direct volume; this shows the namespace name and the Volume Path field (below) shows the volume name.

Volume Path only appears for managed-volume shares in direct volumes. This is the name of the volume behind this direct share, as set by the managed-volume command.

NFS Export is the name of the NFS export at the back-end filer. This is also set by the filer command.

CIFS Share (also set by filer) is the name of the CIFS share at the filer. In a multi-protocol namespace, both the NFS Export and CIFS Share appear for each namespace share.

Features is a list of multi-protocol features supported at the back-end share:

- "cifs-acls,"
- "unix-perm,"
- "shared-attr,"
- "asymmetric-ro,"
- "no-trail-period" (the filer creates a CIFS FGN for file/directory names that end in a period, so the volume will never migrate any such names to it),
- "fgn-see-thru" (the filer allows write access to files through their CIFS "8.3" names, so the share is not supported in a multi-protocol volume),
- "ignore-unix" (the share is backed by an NTFS qtree that ignores changes to Unix permissions without any error; the managed volume also quietly drops these changes),
- "cifs-case-blind," and/or
- "unix-acls" (The filer is a NetApp with the **cifs.perserve_unix_security** option enabled. The NetApp converts Unix security flags into a special CIFS ACL, so that CIFS clients can change the Unix security settings. The volume software records this setting at import time so that it can support those CIFS clients.)

Guidelines: Shares (Cont.)

CIFS Access-based Enum is "Exclude" if someone issued cifs access-based-enum exclude for the share. Otherwise, this field does not appear in the output. The "Exclude" flag indicates that the filer behind the share cannot support access-based enumeration (ABE), so this share is excluded from the volume's ABE-consistency checking.

SID Translation only appears for a CIFS volume that translates Security IDs (SIDs) for all files migrated to or from this share. This indicates that the filer uses Local Group support, so the SID for each group name is unique at this filer. Use the sid-translation command to enable SID translation for a share.

Ignore SID errors only appears for a volume that supports CIFS. A "Yes" indicates that the back-end server is configured to return an error for a file or directory with an unknown SID, but accept the file/directory anyway. A "No" indicates that SID errors from the filer indicate a rejection of the file or directory. You can change this with the ignore-sid-errors command.

Status is the current status of the share. This cycles through several stages during import. The namespace *imports* the files and directories from back-end storage when you create and enable the share. See *Guidelines: Share-Import Status*, below.

Volume Root Backing appears if the share is designated to hold new files created in the root of the volume. This is the first share imported into the volume, which is typically the first share configured.

Critical Share appears if the share has been tagged with the critical command. If a redundant switch loses contact with a critical share, it may fail over to its peer switch.

Migrate Retain Files: Yes only appears the share is set to keep copies of all files migrated away from it. The copies are kept in a hidden directory at the root of the share. Use migrate retain-files to edit this setting.

Strict Attr Consistency: No only appears if you disable strict-attribute consistency. This is recommended for multi-protocol (CIFS and NFS) shares, which may have directories with different CIFS-side and NFS-side names. If the volume cannot find the CIFS-side name of a directory, it may not be able to find all of its CIFS attributes. The volume must either rename these directories (see below) or be allowed to operate without strict-attribute consistency; it cannot import the share otherwise. You can use no strict-attribute-consistency to set this.

Import Sync Attributes: Yes only appears if (during import) the volume is allowed to synchronize the attributes of colliding directories in this share. That is, if a directory in this share has the same name as an already-imported directory but different file attributes (such as read/write permissions), the volume can change the attributes. This is strongly recommended for shares in multi-protocol namespaces. Use import sync-attributes to change this.

Import Rename Files: No only appears if (during import) the volume is prevented from renaming collision files in this share. A file is said to *collide* if it has the same name and path as an already-imported file or directory. Use modify to allow the volume to change files at all during import, and use no import rename-files to disallow file-name changes in this particular share.

Guidelines: Shares (Cont.)

Import Skip Managed Check: Yes is another import option. This only appears if someone used import skip-managed-check on this share. It means that the volume will not run a time-consuming test on any of the share's directories during import. The test confirms that the directory is not already managed by another ARX volume.

Import Rename Directories: No only appears if (during import) the volume is prevented from renaming collision directories in this share. Use modify to allow the volume to change directories at all during import, and use no import rename-directories to disallow directory-name changes in this particular share.

Import Rename Non Mappable Directories: Yes only appears if the volume is allowed to rename a multi-protocol directory with a CIFS-side name that is unmappable to the NFS character-encoding. Filers create unintelligible NFS-side names for these files; this indicates that the volume is allowed to rename these directories during import. Use import rename-directories unmapped-unicode to set this.

Import Priority shows the priority for this share over other shares in the same volume. If two shares have a conflicting file or directory and their import priorities are different, the share with the higher priority wins the conflict. The file or directory on the winning share is the *master*, and the other file or directory may have to change according to one of the import settings above. You can use the import priority command to change the import priority for any given share.

Free space on storage is the remaining space on the share. The expression, "(excluded from volume)," appears after the number if someone used the freespace ignore command in this share; it means that this share's free space is not being counted toward the volume's free space.

Freespace adjustment only appears if someone used freespace adjust to change the advertised free space for this share. This is the adjustment to the free space advertised for this back-end share. For example, if this is 1024 bytes, the volume adds 1024 bytes to the share's current free space.

Total space on storage is the sum of the share's free space (above) and used space.

Apparent size of storage is the size of the share that is advertised to clients. You can change this with the freespace apparent-size command; this field only appears if that command is set.

Policy Maintain Freespace is the amount of free space to maintain on the share. The policy engine does not migrate files to this share if it drops below this amount of free space. You can change this with the policy freespace (gbl-ns, gbl-ns-vol) command.

Policy Resume Freespace is another free-space threshold for this share. If the share drops below the free-space to maintain (above), the share is ineligible for any more migrations until it rises back up to this *resume* value. You can change this threshold with the same policy freespace (gbl-ns, gbl-ns-vol) command that you use for the *maintain* threshold.

Free files on storage (NFS shares only) is the maximum number of files that can be added to the back-end export. All file systems impose a limit on the maximum number of files on a share.

Virtual inodes (NFS direct shares only) is the total number of inodes (files) that can be supported at the direct share.

Guidelines: Shares (Cont.)

Transitions shows the number of times that the share has changed from offline to online, or from online to offline.

Last Transition is the date of the last transition.

Last Probe Status only appears if the most-recent probe of the share resulted in a failure. The volume probes the share at regular intervals to confirm its health; if this field appears, the share failed the latest probe and the failure status appears here. This often indicates a problem with the back-end filer or filer connection; contact F5 Support if you see this field and require any guidance.

Guidelines: Volume State

In the Volume's State field, any of the following messages may appear:

- Enabled - Someone used enable (gbl-ns, gbl-ns-vol) to enable the volume. If nothing appears after this, the volume is capable of processing client requests. Otherwise, one of the following messages appears in parenthesis, indicating an interim and/or degraded state:
  - Waiting for available metadata location - The volume is looking for an available share amongst its configured metadata shares (metadata share).
  - Metadata free space low,
  - Metadata free space very low, and
  - Insufficient metadata free space all indicate a progressively-worse problem with free space on the volume's metadata share: the free space is below 512 M, 256 M, or 128 M. Each volume requires multiple Gigabytes of space for its metadata. Choose another metadata share for the volume (using nsck ... migrate-metadata), or connect to the metadata share and clear any space outside of any .acopia directories.
  - Metadata filer failed - The volume's metadata share is offline. The volume cannot function without this share. Check and correct the share at the file server itself. If necessary, use nsck ... migrate-metadata to move the volume's metadata to a new share.
  - Metadata starting - Volume processing is starting (perhaps after an ARX reboot), and the volume is re-connecting to its metadata share.
  - No enabled shares - The volume is ready for processing, but none of its shares are running. Use enable (gbl-ns-vol-shr) to enable a share.
  - Read Only - The volume allows reads but no writes. This might be because it is configured as a shadow volume, or because no one has set the modify flag.
  - Instance not responding - Volume processing has failed.
- Starting - The volume's enabled shares are currently importing files and directories from back-end filers.
- Stopping - Someone has issued a no enable (gbl-ns, gbl-ns-vol) command to stop processing in this volume, and the volume is in the process of stopping. You cannot change the volume's configuration until the volume reaches the Disabled state.

Guidelines: Volume State (Cont.)

- Disabled - The volume is fully stopped after someone has issued a no enable (gbl-ns, gbl-ns-vol) command. Clients cannot access this volume. You can make changes to the volume's configuration.
- Offline (Migrating Metadata) - The volume's metadata is migrating from one filer to another. The volume comes back online after the migration finishes. You can use the nsck ... migrate-metadata command to migrate the volume's metadata.

Guidelines: Share-Import Status

In the Share's Status field, the following messages show the progress of a successful import:

- Pending (Uninitialized) - The share is configured, but either the volume or the share is not yet enabled. Use the enable (gbl-ns, gbl-ns-vol) command to enable the volume. For a new share in an already-enabled volume, use the enable (gbl-ns-vol-shr) command to enable the share.
- Pending Import - Scheduled for import. If this volume supports CIFS and has filer-subshares or cifs access-based-enum enabled, the volume is synchronizing all of the subshares and/or ABE settings it can find on the filers behind it.
- Pending Import: Unavailable - The volume lost contact with the filer after the share was scheduled for import. You can use no enable (gbl-ns-vol-shr) on this share to send it to the Offline: Disabled status; this allows an nsck ... rebuild of the volume.
- Importing: a items scanned, b items imported - Actively inventorying the back-end storage. Both counters, *a* and *b*, are the sum of files and directories.
- Import Interrupted - The ARX unexpectedly lost contact with the back-end share/export, disrupting the import process.
- Import Paused: Volume Disabled - Someone issued the no enable command while the volume was importing. To restart the import, use the enable (gbl-ns, gbl-ns-vol) command.
- Pending Online - The share is nearly online.
- Online: Direct - This volume has direct shares, and is ready for client access (no import is required).
- Online - Import completed successfully, back-end storage is ready to be used.

Guidelines: Disable/Removal Status

These messages indicate the status of a share that has been disabled and/or removed:

- Offline: Disabled - either the share or the entire volume is disabled. Use the enable (gbl-ns-vol-shr) command to enable the share; use the enable (gbl-ns, gbl-ns-vol) command to enable the entire volume.
- Offline: Unavailable - the volume lost contact with the filer.
- Pending Remove - The share has been marked for removal by the no share command. See share.
- Preparing Remove: n records visited - The switch is scanning the volume's metadata, separating out the metadata for this particular share and retrieving all master-directory attributes from the back-end filer. The *n* is the number of files and directories scanned.
- Removing: n records visited - The switch is currently removing the share from the namespace. The *n* is the number of file/directory records that have been removed from the volume's metadata.
- Remove Paused: Volume Disabled - the volume was disabled after the removal started. Use the enable (gbl-ns, gbl-ns-vol) command to enable the volume and restart the share removal.
- Remove Interrupted - The ARX unexpectedly lost contact with the back-end share/export, disrupting the remove process.

Guidelines: Import Errors

An Error at the beginning of the message indicates the import failed. There are a large number of specific import errors, to help with diagnosis and recovery. See Table 21.1 on page 21-49 for a full list of possible errors and suggestions for troubleshooting each error.

Once the error is corrected, you can use the no share command to remove the share, then share to re-import it. For metadata shares, the entire volume must be re-imported: use the nsck ... rebuild command for the volume.

Samples

`bstnA# show namespace`
lists all namespaces. See Figure 21.1 for sample output.

`bstnA# show namespace medarcv`
shows all volumes in a CIFS namespace; see Figure 21.2 on page 21-38.

`bstnA# show namespace wwmed`
shows all volumes in an NFS namespace; see Figure 21.3 on page 21-43 for sample output.

`prtlndA# show namespace nemed volume /acctShdw`
shows one volume with manual free-space calculations. See Figure 21.4 on page 21-45 for sample output.

`bstnA# show namespace insur volume /claims`
shows one volume in a multi-protocol (NFS and CIFS) namespace. See Figure 21.5 on page 21-46 for sample output.

`bstnA# show namespace medco volume /vol share generic`
shows one share in a direct volume. See Figure 21.6 on page 21-48 for sample output.

Related Commands

namespace
enable (gbl-ns, gbl-ns-vol)
enable (gbl-ns-vol-shr)
ip address
grep
nsck ... rebuild

Figure 21.1 *Sample Output: show namespace*

```
bstnA# show namespace

Configured Namespaces
---------------------

Namespace                       Description
------------------------------  ------------------------------------------
medco
wwmed                           namespace for World-Wide Medical network
medarcv
insur                           WW Medical insurance claims and records
labResearch
```

Figure 21.2 *Sample Output: show namespace medarcv (CIFS)*

```
bstnA# show namespace medarcv

Namespace "medarcv" Configuration
Description: (none)
Metadata Cache Size: 512 MB
Proxy User: acoProxy2
Filer SMB Signatures: Enabled
```

```
SAM Reference Filer: fs2 (192.168.25.27)

Supported Protocols
-------------------
  cifs

CIFS Authentication
-------------------
Protocols:
  NTLM
  NTLMv2
  Kerberos

Participating Switches
----------------------
  bstnA (Volume Group 2) [Current Switch]


Windows Management Authorization Policies
-----------------------------------------
  readOnly
  fullAccess
  snapViewers

Volumes
-------
  /lab_equipment
    CIFS : compressed files: no; named streams: yes; persistent ACLs: no
           sparse files: no; Unicode on disk: yes; case sensitive: no

             Volume freespace: 249 GB  automatic
           Volume total space: 321 GB
                  CIFS quotas: Not Enabled
              Auto Sync Files: Enabled
                Metadata size: 92 kB
         Metadata free space: 16 GB
              Oplock support: Enabled
          Notify-change mode: Normal
             CIFS path cache: Enabled
      CIFS access based enum: Not Enabled
                   Snapshots: Enabled
            Migration method: Staged
                       State: Enabled

                  Host Switch: bstnA
                     Instance: 3
                 Volume Group: 2
                    Processor: 1.1
                        Files: 68 used (9 dirs), 3.9 M free, 248 M max (automatic)
    Metadata shares:

      Filer          Backend Path           Contains Metadata  Status
      ------------------------------------------------------------------
      nas1           /vol/vol2/meta6        Yes                Online

    Share backlots
      Filer                      fs2 [192.168.25.27]
      CIFS Share                 backlot_records
      Features                   cifs-acls cifs-case-blind
      SID Translation            No
      Ignore SID errors          No
      Status                     Online
```

```
    Import Sync Attributes      Yes
    Import Priority             65535 (Lowest)
    Free space on storage       1.3 GB (1,405,523,968 B)
    Total space on storage      1.9 GB (2,138,540,032 B)
    Policy Maintain Freespace   3 %
    Policy Resume Freespace     5 %
    Transitions                 1
    Last Transition             Fri 16 Nov 2012 02:11:56 AM EST

  Share equip
    Filer                       nas10 [192.168.25.49]
    CIFS Share                  equipment
    Features                    cifs-acls cifs-case-blind
    SID Translation             No
    Ignore SID errors           No
    Status                      Online
    Volume Root Backing         Yes
    Import Sync Attributes      Yes
    Import Skip Managed Check   Yes
    Import Priority             1 (Highest)
    Free space on storage       94 GB (101,073,227,776 B)
    Total space on storage      120 GB (128,849,018,880 B)
    Policy Maintain Freespace   3 %
    Policy Resume Freespace     5 %
    Transitions                 1
    Last Transition             Fri 16 Nov 2012 02:11:56 AM EST

  Share equipSnap               [replica-snap]
    Filer                       nas11 [192.168.25.48]
    CIFS Share                  equipBkup
    Features
    Status                      Online
    Free space on storage       1.5 GB (1,653,608,448 B)
    Total space on storage      1.5 GB (1,717,989,376 B)
    Transitions                 1
    Last Transition             Fri 16 Nov 2012 02:11:57 AM EST

  Share leased
    Filer                       nas10 [192.168.25.49]
    CIFS Share                  for_lease
    Features                    cifs-acls cifs-case-blind
    SID Translation             No
    Ignore SID errors           No
    Status                      Online
    Import Sync Attributes      Yes
    Import Priority             1 (Highest)
    Free space on storage       78 GB (84,064,509,952 B)
    Total space on storage      120 GB (128,849,018,880 B)
    Policy Maintain Freespace   3 %
    Policy Resume Freespace     5 %
    Transitions                 1
    Last Transition             Fri 16 Nov 2012 02:11:56 AM EST

  Share leasedSnap              [replica-snap]
    Filer                       nas11 [192.168.25.48]
    CIFS Share                  leasedBkup
    Features
    Status                      Online
    Free space on storage       1.5 GB (1,653,608,448 B)
    Total space on storage      1.5 GB (1,717,989,376 B)
    Transitions                 1
    Last Transition             Fri 16 Nov 2012 02:11:57 AM EST
```

```
  Share scanners
    Filer                      fs5 [192.168.25.71]
    CIFS Share                 xraysScanners
    Features                   cifs-acls cifs-case-blind
    CIFS Access-based Enum     Exclude
    SID Translation            No
    Ignore SID errors          No
    Status                     Online
    Import Sync Attributes     Yes
    Import Priority            65535 (Lowest)
    Free space on storage      75 GB (81,326,866,432 B)
    Total space on storage     79 GB (85,896,564,736 B)
    Policy Maintain Freespace  3 %
    Policy Resume Freespace    5 %
    Transitions                1
    Last Transition            Fri 16 Nov 2012 02:11:56 AM EST


/rcrds
  CIFS : compressed files: yes; named streams: yes; persistent ACLs: yes
         sparse files: yes; Unicode on disk: yes; case sensitive: no

           Volume freespace: 3.8 GB  automatic
         Volume total space: 5.9 GB
                CIFS quotas: Not Enabled
            Auto Sync Files: Enabled
              Metadata size: 236 kB
        Metadata free space: 16 GB
            Filer Subshares: Enabled
             Oplock support: Enabled
          Notify-change mode: Normal
             CIFS path cache: Not Enabled
     CIFS access based enum: Enabled; Auto-enable
                  Snapshots: Enabled
          Migration method: Staged
                     State: Enabled

                Host Switch: bstnA
                   Instance: 3
               Volume Group: 2
                  Processor: 1.1
                      Files: 247 used (38 dirs), 3.9 M free, 248 M max (automatic)
  Metadata shares:

    Filer          Backend Path           Contains Metadata  Status
    ---------------------------------------------------------------
    nas1           /vol/vol2/meta3        Yes                Online

  Share bulk
    Description                new server to hold big files (such as xrays)
    Filer                      fs2 [192.168.25.27]
    CIFS Share                 bulkstorage
    Features                   cifs-acls cifs-case-blind
    SID Translation            Yes
    Ignore SID errors          No
    Status                     Online
    Import Sync Attributes     Yes
    Import Priority            65535 (Lowest)
    Free space on storage      1.3 GB (1,405,523,968 B)
    Total space on storage     1.9 GB (2,138,540,032 B)
    Policy Maintain Freespace  5 %
```

```
    Policy Resume Freespace     6 %
    Transitions                 1
    Last Transition             Fri 16 Nov 2012 02:11:56 AM EST

  Share charts
    Description                 various medical charts
    Filer                       fs1 [192.168.25.20]
    CIFS Share                  histories
    Features                    cifs-acls cifs-case-blind
    SID Translation             No
    Ignore SID errors           No
    Status                      Online
    Import Sync Attributes      Yes
    Import Priority             65535 (Lowest)
    Free space on storage       1.4 GB (1,508,139,008 B)
    Total space on storage      1.9 GB (2,138,540,032 B)
    Policy Maintain Freespace   5 %
    Policy Resume Freespace     6 %
    Transitions                 1
    Last Transition             Fri 16 Nov 2012 02:11:56 AM EST

  Share rx
    Description                 prescriptions since 2002
    Filer                       fs4 [192.168.25.29]
    CIFS Share                  prescriptions
    Features                    cifs-acls cifs-case-blind
    SID Translation             No
    Ignore SID errors           No
    Status                      Online
    Volume Root Backing         Yes
    Critical Share              Yes
    Import Sync Attributes      Yes
    Import Priority             65535 (Lowest)
    Free space on storage       1.1 GB (1,181,872,128 B)
    Total space on storage      1.9 GB (2,144,333,824 B)
    Policy Maintain Freespace   5 %
    Policy Resume Freespace     6 %
    Transitions                 1
    Last Transition             Fri 16 Nov 2012 02:11:56 AM EST


/test_results
  CIFS : compressed files: yes; named streams: yes; persistent ACLs: yes
         sparse files: yes; Unicode on disk: yes; case sensitive: no

           Volume freespace: 6.5 GB
         Volume total space: 9.9 GB
              Metadata size: 28 kB
             Oplock support: Automatic
         Notify-change mode: Normal
            CIFS path cache: Enabled
    CIFS access based enum: Not Enabled
                 Snapshots: Not Enabled
                     State: Enabled

               Host Switch: bstnA
                  Instance: 3
              Volume Group: 2
                 Processor: 1.1
                     Files: 1 used (1 dir), 31 M free

  Share 2005_charts
```

```
      Filer                       medarcv [Acopia Namespace]
      Volume Path                 /rcrds
      SID Translation             No
      Ignore SID errors           No
      Status                      Online
      Free space on storage       3.8 GB (4,095,533,056 B)
      Total space on storage      5.9 GB (6,421,413,888 B)
      Transitions                 1
      Last Transition             Fri 16 Nov 2012 02:11:56 AM EST

    Share chemistry
      Filer                       fs1 [192.168.25.20]
      CIFS Share                  chem_results
      SID Translation             No
      Ignore SID errors           No
      Status                      Online
      Free space on storage       1.4 GB (1,508,139,008 B)
      Total space on storage      1.9 GB (2,138,540,032 B)
      Transitions                 1
      Last Transition             Fri 16 Nov 2012 02:11:56 AM EST

    Share hematology
      Filer                       fs3 [192.168.25.28]
      CIFS Share                  hematology_results
      SID Translation             No
      Ignore SID errors           No
      Status                      Online
      Free space on storage       1.3 GB (1,458,823,168 B)
      Total space on storage      1.9 GB (2,097,147,904 B)
      Transitions                 1
      Last Transition             Fri 16 Nov 2012 02:11:56 AM EST
```

Figure 21.3 *Sample Output: show namespace wwmed (NFS)*

```
bstnA# show namespace wwmed

Namespace "wwmed" Configuration
Description: namespace for World-Wide Medical network
Metadata Cache Size: 512 MB
NFS Character Encoding: ISO-8859-1

Supported Protocols
-------------------
  nfsv3 nfsv3-tcp

Participating Switches
----------------------
  bstnA (Volume Group 1) [Current Switch]


Volumes
-------
  /acct

            Volume freespace: 263 GB  automatic
          Volume total space: 328 GB
               Metadata size: 1.9 MB
         Metadata free space: 16 GB
                   Snapshots: Not Enabled
            Migration method: Staged
                       State: Enabled
```

```
              Host Switch: bstnA
                 Instance: 2
             Volume Group: 1
                Processor: 1.1
                    Files: 4,433 used (439 dirs), 3.9 M free, 252 M max (automatic)
Metadata shares:

  Filer          Backend Path           Contains Metadata  Status
  ---------------------------------------------------------------------
  nas1           /vol/vol2/meta1        Yes                Online

Share bills
  Filer                      das8 [192.168.25.25]
  NFS Export                 /work1/accting
  Features                   unix-perm
  Status                     Online
  Critical Share             Yes
  Import Sync Attributes     Yes
  Import Priority            65535 (Lowest)
  Free space on storage      23 GB (24,949,542,912 B)
  Total space on storage     70 GB (75,278,499,840 B)
  Policy Maintain Freespace  2 %
  Policy Resume Freespace    3 %
  Free files on storage      17M
  Transitions                1
  Last Transition            Fri 16 Nov 2012 12:52:10 AM EST

Share bills2
  Filer                      das3 [192.168.25.23]
  NFS Export                 /exports/acct2
  Features                   unix-perm
  Status                     Online
  Import Sync Attributes     Yes
  Import Priority            65535 (Lowest)
  Free space on storage      111 GB (119,787,619,328 B)
  Total space on storage     113 GB (121,333,164,032 B)
  Policy Maintain Freespace  2 %
  Policy Resume Freespace    3 %
  Free files on storage      13M
  Transitions                1
  Last Transition            Fri 16 Nov 2012 12:52:11 AM EST

Share budget
  Filer                      das1 [192.168.25.19]
  NFS Export                 /exports/budget
  Features                   unix-perm
  Status                     Online
  Volume Root Backing        Yes
  Import Priority            65535 (Lowest)
  Free space on storage      80 GB (86,870,663,168 B)
  Total space on storage     89 GB (95,607,627,776 B)
  Policy Maintain Freespace  2 %
  Policy Resume Freespace    3 %
  Free files on storage      10M
  Transitions                1
  Last Transition            Fri 16 Nov 2012 12:52:08 AM EST

Share it5
  Filer                      das7 [192.168.25.24]
  NFS Export                 /lhome/it5
  Features                   unix-perm
```

```
      Status                      Online
      Import Sync Attributes      Yes
      Import Priority             65535 (Lowest)
      Free space on storage       48 GB (51,670,810,624 B)
      Total space on storage      56 GB (60,261,593,088 B)
      Policy Maintain Freespace   2 %
      Policy Resume Freespace     3 %
      Free files on storage       14M
      Transitions                 1
      Last Transition             Fri 16 Nov 2012 12:52:13 AM EST
```

Figure 21.4 *Sample Output: show namespace nemed volume /acctShdw*

```
prtlndA# show namespace nemed volume /acctShdw

Namespace "nemed" Configuration
Description: namespace for NorEast Medical network
Metadata Cache Size: 512 MB
NFS Character Encoding: ISO-8859-1

Supported Protocols
-------------------
  nfsv3

Participating Switches
----------------------
  prtlndA (Volume Group 1) [Current Switch]


Volumes
-------
  /acctShdw       (Shadow Volume)

            Volume freespace: 18 GB  manual
           Volume total space: 20 GB
                Metadata size: 2.6 MB
          Metadata free space: 16 GB
                    Snapshots: Not Enabled
             Migration method: Staged
                        State: Enabled

                  Host Switch: prtlndA
                     Instance: 1
                 Volume Group: 1
                    Processor: 1.1
                        Files: 4,866 used (488 dirs), 3.9 M free, 252 M max (automatic)
    Metadata shares:

      Filer          Backend Path           Contains Metadata  Status
      ------------------------------------------------------------------
      nas-p1         /vol/vol2/mdata_A      Yes                Online

    Share back1
      Filer                       das-p1 [192.168.74.81]
      NFS Export                  /exports/BU
      Features                    unix-perm
      Status                      Online
      Volume Root Backing         Yes
      Import Priority             65535 (Lowest)
      Free space on storage       17 GB (18,465,959,936 B)
      Freespace adjustment        1.0 GB (1,073,741,824 B)
      Total space on storage      19 GB (21,129,465,856 B)
```

```
      Policy Maintain Freespace   2 %
      Policy Resume Freespace     4 %
      Free files on storage       2M
      Transitions                 1
      Last Transition             Fri 16 Nov 2012 02:34:57 AM EST

    Share back2
      Filer                       das-p2 [192.168.74.82]
      NFS Export                  /exports/bkup
      Features                    unix-perm
      Status                      Online
      Critical Share              Yes
      Import Priority             65535 (Lowest)
      Free space on storage       110 GB (118,804,558,848 B) (excluded from volume)
      Total space on storage      113 GB (121,333,164,032 B)
      Policy Maintain Freespace   2 %
      Policy Resume Freespace     4 %
      Free files on storage       13M
      Transitions                 1
      Last Transition             Fri 16 Nov 2012 02:34:58 AM EST
```

Figure 21.5 *Sample Output: show namespace insur volume /claims (MPNS)*

```
bstnA# show namespace insur volume /claims

Namespace "insur" Configuration
Description: WW Medical insurance claims and records
Metadata Cache Size: 512 MB
Proxy User: acoProxy2
Filer SMB Signatures: Enabled
NFS Character Encoding: ISO-8859-1

Supported Protocols
-------------------
  cifs nfsv3 nfsv3-tcp

CIFS Authentication
-------------------
Protocols:
  NTLM
  NTLMv2
  Kerberos

Participating Switches
----------------------
  bstnA (Volume Group 10) [Current Switch]


Windows Management Authorization Policies
-----------------------------------------
  readOnly
  fullAccess
  snapViewers

Volumes
-------
  /claims
    CIFS : compressed files: no; named streams: yes; persistent ACLs: yes
           sparse files: no; Unicode on disk: yes; case sensitive: no

             Volume freespace: 19 GB  clients use dir-master-only
```

```
       Volume total space: 19 GB
              CIFS quotas: Not Enabled
         Auto Sync Files: Enabled
           Metadata size: 244 kB
      Metadata free space: 16 GB
          Oplock support: Disabled
       Notify-change mode: Normal
          CIFS path cache: Enabled
  CIFS access based enum: Not Enabled
               Snapshots: Enabled
        Migration method: Staged
                   State: Enabled

             Host Switch: bstnA
                Instance: 4
            Volume Group: 10
               Processor: 1.1
                   Files: 197 used (27 dirs), 63 M free
Metadata shares:

  Filer          Backend Path           Contains Metadata  Status
  ------------------------------------------------------------------
  nas1           /vol/vol2/meta2        Yes                Online

Share shr1-next
  Filer                       nasE1 [192.168.25.51]
  NFS Export                  /patient_records
  CIFS Share                  patient_records
  Features                    cifs-acls unix-perm cifs-case-blind
  SID Translation             Yes
  Ignore SID errors           Yes
  Status                      Online
  Strict Attr Consistency     No
  Import Sync Attributes      Yes
  Import Priority             65535 (Lowest)
  Free space on storage       2.4 GB (2,577,547,264 B)
  Total space on storage      2.4 GB (2,580,463,616 B)
  Policy Maintain Freespace   1.0 GB
  Policy Resume Freespace     2.0 GB
  Free files on storage       314843
  Transitions                 1
  Last Transition             Fri 16 Nov 2012 02:11:57 AM EST

Share shr1-old
  Filer                       nas1 [192.168.25.21]
  NFS Export                  /vol/vol2/insurance
  CIFS Share                  insurance
  Features                    cifs-acls shared-attr cifs-case-blind
  SID Translation             Yes
  Ignore SID errors           No
  Status                      Online
  Volume Root Backing         Yes
  Strict Attr Consistency     No
  Import Sync Attributes      Yes
  Import Rename Non Mappable Directories Yes
  Import Priority             65535 (Lowest)
  Free space on storage       16 GB (17,881,849,856 B)
  Total space on storage      17 GB (18,360,987,648 B)
  Apparent size of storage    2.4 GB (2,579,496,960 B)
  Policy Maintain Freespace   1.0 GB
  Policy Resume Freespace     2.0 GB
  Free files on storage       619659
```

```
      Transitions                1
      Last Transition            Fri 16 Nov 2012 02:11:57 AM EST
```

Figure 21.6 *Sample Output: show namespace medco ... share ... (direct share)*

```
bstnA# show namespace medco volume /vol share generic

Namespace "medco" Configuration
Description: (none)
Metadata Cache Size: 512 MB
NFS Character Encoding: ISO-8859-1

Supported Protocols
-------------------
  nfsv3-tcp

Participating Switches
----------------------
  bstnA (Volume Group 9) [Current Switch]


Volumes
-------
  /vol

            Volume freespace: 116 GB
           Volume total space: 171 GB
                Metadata size: 28 kB
                    Snapshots: Not Enabled
                        State: Enabled

                  Host Switch: bstnA
                     Instance: 1
                 Volume Group: 9
                    Processor: 1.1
                        Files: 1 used (1 dir), 31 M free

    Share generic
      Filer                      nas3 [192.168.25.47]
      NFS Export                 /exports
      Status                     Online
      Free space on storage      52 GB (56,813,814,784 B)
      Total space on storage     57 GB (62,229,050,368 B)
      Free files on storage      6M
      Virtual inodes             16M
      Transitions                1
      Last Transition            Fri 16 Nov 2012 12:51:42 AM EST
```

Status Condition	Description/Action
`Error: A directory was removed from the filer during import.`	A directory was found during an early scan of the filer share, then later it was missing. Either a client or client application must have accessed the filer directly (that is, not through this managed volume) during the import. This ARX cannot support unexpected changes to the back-end shares during (or after) import. After you correct the problem, use nsck ... rebuild volume to reimport all shares in the volume.
`Error: A filename/directory could not be mapped to the namespace character encoding.`	One of the share's directories has a name with Unicode characters that are unsupported by the character-encoding nfs setting. CIFS file names are Unicode and can contain any character, but NFS servers and clients must each configure their character encoding for file names. The volume cannot import a directory with any un-mappable characters in its name. You can use the import rename-directories unmapped-unicode command to allow the volume to rename such directories during import, or you can rename them manually at the filer. Then restart the share import: enter gbl-ns-vol-shr mode and re use the enable (gbl-ns-vol-shr) command.
`Error: a higher priority share failed to initialize import.`	A share with a higher import priority has failed its import, so this share cannot import. If any share import fails, the managed volume cannot import any shares with lower import priorities. Find the import error for the failed share(s), look for the error in this table, and take action as directed. This error is resolved as soon as all higher-priority shares successfully import.
`Error: Access denied by filer.`	This indicates that the ARX does not have proper permissions to connect to the back-end filer, or to perform some operation in the filer share. For NFS exports, check your back-end filer configuration: the back-end share should allow *root* access to all of the ARX's proxy IP addresses. Use the show exports command examine all permission settings at the filer. Use the show ip proxy-addresses command to list all configured proxy IP addresses. For CIFS shares, the switch uses the namespace's *proxy user* (username and password). The proxy-user credentials must belong to the Administrators group at every filer behind the namespace. Use the probe exports command to check this. The proxy-user (gbl-ns) command sets the proxy user credentials for a namespace. This may also indicate that the ARX is examining a read-only directory, such as the "~snapshot" directory in a NetApp filer. Such directories should be ignored during import. The gbl-filer ignore-name command identifies these types of directories. After you find and fix the error, you can restart the share import: enter gbl-ns-vol-shr mode for this share and re use the enable (gbl-ns-vol-shr) command.

Table 21.1 *Share Status Conditions*

Status Condition	Description/Action
`Error: Another share in the volume failed to import.`	All shares in the volume must import successfully for the volume to come online. Address the share that does not have this error, then restart this share import: enter gbl-ns-vol-shr mode for this share and re use the enable (gbl-ns-vol-shr) command.
`Error: Attempted to import subdirectory which is also a managed root`	This share is a parent to an already-imported share. Namespace shares cannot overlap. Use the filer command to change the path or share name. After you correct the problem, use nsck ... rebuild volume to reimport all shares in the volume.
`Error: Attributes of the share root are inconsistent.`	The share's root directory has attributes (such as owner, group, and permission settings) that are inconsistent with those of the already-imported shares. You can access the back-end filer directly to change these attributes, or you can use the import sync-attributes command to allow the volume to change the attributes for you. Then re-enable the share (enable (gbl-ns-vol-shr)) to restart its import.
`Error: Authentication error during import.`	The ARX does not have proper permissions to connect to the back-end filer. For NFS exports, check your back-end filer configuration: the back-end share should allow *root* access to all of the ARX's proxy IP addresses. Use the show exports command examine all permission settings at the filer. Use the show ip proxy-addresses command to list all configured proxy IP addresses. For CIFS shares, the switch uses the namespace's *proxy user* (username and password). The proxy-user credentials must belong to the Administrators group at every filer behind the namespace. The proxy-user (gbl-ns) command sets the proxy user credentials for a namespace. After you find and fix this issue, you can restart the share import: enter gbl-ns-vol-shr mode for this share and re use the enable (gbl-ns-vol-shr) command.
`Error: Bad Export Path/Share Name`	The share cannot be found on the external filer. Use the filer command to change the path or share name for this share, then re-enable the share (enable (gbl-ns-vol-shr)) to retry the import.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: CIFS ABE check failed.`	(CIFS) The volume supports cifs access-based-enum (ABE), and attempted to replicate ABE settings between its back-end shares. This replication process failed. The same process also checks for CIFS subshares (filer-subshares), so you can use sync subshares from-namespace ... tentative to get a full report on this issue. This often occurs because the ARX does not have proper permissions to check for ABE support on this back-end share. The ARX uses the namespace's *proxy user* (username and password) as its identity when it checks for ABE support. The proxy-user credentials must belong to the Administrators group at this back-end filer. You can use the proxy-user (gbl-ns) command to choose new proxy user credentials for the namespace. After you find and fix this issue, use nsck ... rebuild volume to reimport all shares in the volume.
`Error: CIFS error during import scan.`	(CIFS) This is a CIFS error that is not an access or network error, but prevented the import. The syslog shows the specific error. Use show logs syslog to read the syslog, or grep *string* logs syslog to search for a specific string in the syslog. After you correct the error, re-enable the share (enable (gbl-ns-vol-shr)) to retry the import.
`Error: CIFS filer connection limit reached.`	(CIFS) The back-end filer imposes a limit on CIFS connections, and the import requires more connections than the filer allows. Correct this at the back-end filer and then re-enable the share to retry the import.
`Error: CIFS operation attempted on a symbolic link.`	(multi-protocol) The volume software encountered an NFS symbolic link on this back-end share, and the volume has cifs deny-symlinks enabled. You can resolve this issue by using the no cifs deny-symlinks command to allow the volume software to follow these links. Alternatively, you can remove all NFS symbolic links from the back-end share. After you fix this issue, use nsck ... rebuild volume to reimport all shares in the volume.
`Error: CIFS operation failed with an error during import.`	(CIFS) The back-end filer returned an unexpected CIFS error during import. The syslog shows the specific error. Use show logs syslog to read the syslog, or grep *string* logs syslog to search for a specific string in the syslog. You may need to escalate to F5 Support. After you correct the error, re-enable the share (enable (gbl-ns-vol-shr)) to retry the import.
`Error: CIFS operation failed with an error indicating a filer fault.`	(CIFS) The filer returned an unexpected error during the import, and the error indicates a problem at the filer itself. The syslog shows the specific error. (Use show logs syslog to read the syslog, or grep *string* logs syslog to search for a specific string in the syslog.) Check the filer itself and correct the problem there. After you correct the error, re-enable the share (enable (gbl-ns-vol-shr)) to retry the import.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: CIFS operation returned an error indicating an unsupported option.`	(CIFS) The back-end share returned an error indicating that it does not support a CIFS option that the ARX requires. Consult the *F5 Data Solutions Compatibility Matrix* (included in this doc set) to confirm that the filer has been qualified for use behind the ARX. If the share cannot possibly support CIFS behind an ARX, you can use no share to remove the share from the volume. After you fix this issue, use nsck ... rebuild volume to reimport all shares in the volume.
`Error: CIFS proxy-user not privileged.`	(CIFS) The namespace software attempted to write a test file to the share and failed. Go to the filer and check permissions for the namespace's proxy-user (gbl-ns); the proxy user must be part of the Backup Operators and/or Administrators group on the filer. After you correct the error, re-enable the share (enable (gbl-ns-vol-shr)) to retry the import.
`Error: CIFS subshare check failed.`	(CIFS) The volume supports filer-subshares and/or cifs access-based-enum (ABE), and attempted to replicate subshares, subshare ACLs, and/or ABE settings between its back-end shares. This replication process, also known as *subshare synchronization*, failed. As a result, any front-end export of the failed subshare will be degraded. The output of show cifs-service *fqdn* shows all of the degraded subshares in a given *fqdn* service. Use sync subshares from-namespace ... tentative to get a full report on this issue. To repair it, use the sync subshares from-namespace or sync subshares from-service command without the tentative option.
`Error: Cannot find credentials for connection to filer.`	(CIFS) Proxy user credentials have not been properly configured for the CIFS share. These credentials must be a username and password from the Administrators group. Use the proxy-user command to add or edit these credentials, and use the proxy-user (gbl-ns) command to apply them to a namespace. After you correct the error, re-enable the share (enable (gbl-ns-vol-shr)) to retry the import.
`Error: Collision rename failed due to open file on share.`	(CIFS) The share had a file that collided with an already-imported file, and the volume failed to rename it because a CIFS application has it open and locked. You can use show cifs-service open-files to find the open file, close cifs file to close it, and then retry the share import (with enable (gbl-ns-vol-shr)).

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Could not update attributes on a file/directory during import.`	Failure to update file attributes can be caused by loss of connectivity during the import. Use the show exports command and/or ping to check the connection to the filer. For CIFS shares, this may indicate permissions problems. Proxy-user credentials may not have been properly configured for the CIFS share. These credentials must be a username and password from the Administrators group. Special, immutable directories can also cause this. The .snapshot directory (on some systems) is an example of an immutable directory, though .snapshot directories are properly ignored by the import software. To ignore other directories on this filer, use the gbl-ext-filer ignore-name command. After you correct the error, re-enable the share (enable (gbl-ns-vol-shr)) to retry the import.
`Error: Detected collisions/inconsistencies during no modify import.`	Two or more of the volume's shares had common file names that either *collided* or had NFS/CIFS naming *inconsistencies*, and this volume disallows import if it encounters either of these problems. As an example of a collision, suppose share A and share B had the same file in the same path, \docs\index.htm: these files would collide. A naming inconsistency can only occur for a directory in a multi-protocol (NFS and CIFS) namespace; the CIFS-side directory name has unicode characters that are inexpressible on the NFS-side (see the documentation for the character-encoding nfs command). The volume must be allowed to *modify* the directory (or one of the colliding files) for the import to succeed: the directory or file must be renamed. All duplicate files and naming inconsistencies are recorded in the import reports for each share. These reports are named "import.*job-id*.*share-name*.*share-id*.rpt." Use show reports to list all import reports and read their contents. Using the import report for each share, resolve all file collisions and naming inconsistencies before re-importing. Go to the filers and rename the files, move them, and/or resolve that certain file renames are acceptable. Once the issues are cleared, use the gbl-ns-vol reimport-modify and modify commands to allow modification (renames) on import. (If any other shares are still importing, you must wait for their imports to finish before you can use the modify command.) To rename inconsistent NFS/CIFS directories, use the import rename-directories unmapped-unicode command, too. After you correct the problem, use nsck ... rebuild volume to reimport all shares in the volume.
`Error: Device Offline`	The share's back-end filer went offline for an extended time in the middle of the import. You can remove the share from the volume, or resolve the issue at the filer and retry the import. To retry the import, you can use the enable (gbl-ns-vol-shr) command on this share.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: DFS links found on share during import.`	Managed volumes do not support Distributed File System (DFS) links. To find all of the DFS links on all of a volume's shares, you can import the share(s) with no modify. Remove all DFS links from the back-end share. Then use the enable (gbl-ns-vol-shr) command to retry the import.
`Error: Directory attribute collision detected during import.`	A directory on this share has the same name and path as a directory on another share, but the attributes are different. (Directory attributes are permission settings for various users and groups, along with various other flags.) For example, suppose share A had a \docs\img directory and share B had a directory with the same name but different write permissions for the "Everyone" group. To allow the volume to modify directory attributes on import, you can use modify on the volume and import sync-attributes on the share. (If any other shares are still importing, you must wait for their imports to finish before you can use the modify command.) Then use the enable (gbl-ns-vol-shr) command to retry the import.
`Error: Directory collision; Unable to guarantee strict-attribute-consistency.`	(multi-protocol) A directory on this share has the same name and path as a directory on an already-imported share, but some of the CIFS attributes are obscured for at least one of them. Obscured CIFS attributes, due to filer-generated names for the directories, are common on a multi-protocol filer. The volume requires that all attributes (for CIFS and NFS) be synchronized for duplicate directories, and that is impossible under the circumstances. This can only occur for NFS-only directories, with names that are illegal in CIFS. If possible, change the directory name(s) so that they are accessible from CIFS. As an alternative, you can use no strict-attribute-consistency to remove the requirement for strict-attribute consistency; this reduces all undiscovered CIFS attributes to zero. Once the volume has stopped importing any shares, you can do this for all shares in the volume. Then restart this share import with the enable (gbl-ns-vol-shr) command.
`Error: Directory name collision detected during import.`	A directory on this share has the same name and path as a file on an already-imported share. For example, share B has a directory named /var/log but share A was already imported with a file named /var/log. To allow the volume to correct this by changing the directory name on import, you can use modify on the volume and import rename-directories on the share. Alternatively, you can directly access the filer(s) and correct the problem there. Retry this share import (with enable (gbl-ns-vol-shr)) after you address the issue.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Dynamic Import operation failed due to a collision.`	Two or more of the volume's shares had common file/directory names that somehow *collided*, causing one of the shares to fail its import. The following collisions can cause this failure: • A file in share A has the same name and path as a file in share B. • A directory in share A has the same name and path as a file in share B. • A directory in share A has the same name and path as a directory in share B, and the directories have different attributes (such as permissions). If they had the same attributes, they would not collide. • A file or directory in share A has a CIFS *case collision* with a file or directory on share B, and the volume is set for no cifs case-sensitive. For example, "myDir/MYFILE.doc" on share A could collide with "myDir/myFile.doc" on share B. All duplicate files and naming inconsistencies are recorded in the import reports for the share. These reports are named "import.*job-id*.*share-name*.*share-id*.rpt." Use show reports type Imp to list all import reports, and use show reports *report-name* to read a report. Using the import report for this share, resolve all file collisions and naming inconsistencies before re-importing. Go to the filers and rename the files, move them, and/or resolve that certain file renames are acceptable. You can also use some share-import options to have the volume automatically rename files, rename directories, or reset directory attributes in this share during import (import rename-files, import rename-directories, or import sync-attributes). If you use any share-import options, use the gbl-ns-vol reimport-modify and modify commands to allow modification (renames) on import. Then re-import the share with the enable (gbl-ns-vol-shr) command.
`Error: Dynamic Import operation failed.`	The share import failed for an undetermined reason. Run the collect diag-info CLI command to collect diagnostic information, then contact F5 Support.
`Error: Dynamic Import operation timed out.`	An internal import operation timed out, possibly due to a filer-connectivity issue. Use the show exports command and/or ping to troubleshoot the connection to the filer.
`Error: Encountered file(s) with more than 1024 hard links during import.`	The managed volume software supports a maximum of 1024 hard links per file. One or more files on this back-end share exceed this limit. These files are recorded in the share's import report, named "import.*job-id*.*share-name*.*share-id*.rpt." Use show reports type Imp to list all import reports, and use show reports *report-name* to read a particular report. Then access the filer directly to reduce the number of hard links for all of these files. After you correct the problem, use nsck ... rebuild volume to reimport all shares in the volume.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Export In Use By Another Switch`	The filer share contains a hidden directory, .acopia, written by a different ARX. Two ARXes cannot import the same share. If the share is not imported by another switch anymore, remove the directory manually. If this happens to all shares after a switch replacement, the old switch's UUID was not properly applied to the replacement switch. Consult the appropriate *Hardware Installation* manual for switch-replacement instructions. After you correct the problem, use nsck ... rebuild volume to reimport all shares in the volume.
`Error: Export is read-only, cannot import.`	Check the share configuration at the filer: the ARX requires read/write access throughout the share's directory tree. Use the show exports command to examine all permission settings at the filer. Use the following guidelines for permissions problems: • For NFS exports, check your back-end filer configuration: the back-end share should allow *root* access to all of the ARX's proxy IP addresses. Use the show ip proxy-addresses command to list all configured proxy IP addresses. • For CIFS shares, the switch uses the namespace's *proxy user* as its identity. The proxy user, created by the proxy-user command, must belong to the Administrators group. The proxy-user (gbl-ns) command sets the proxy user for a namespace. Once the permissions issue is resolved, you can re-import the share with the enable (gbl-ns-vol-shr) command.
`Error: Failed import initialization due to filer full condition.`	The back-end share is full, so the ARX volume could not create the .acopia directory and/or write some small files to mark the share. Go to the filer and clear out some disk space before you retry the import. Once some space is free, you can use the enable (gbl-ns-vol-shr) command to restart the share import.
`Error: Failed import initialization due to permissions or access error.`	This indicates permissions problems at the back-end filer. Use the show exports command to examine all permission settings at the filer. Use the following guidelines for permissions problems: • For NFS exports, check your back-end filer configuration: the back-end share should allow *root* access to all of the ARX's proxy IP addresses. Use the show ip proxy-addresses command to list all configured proxy IP addresses. • For CIFS shares, the switch uses the namespace's *proxy user* as its identity. The proxy user, created by the proxy-user command, must belong to the Administrators group. The proxy-user (gbl-ns) command sets the proxy user for a namespace. Once the permissions issue is resolved, you can re-import the share with the enable (gbl-ns-vol-shr) command.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Failed to Contact Filer.`	The back-end device could not be located with the IP address configured for the filer. From gbl-ext-filer mode, use the ip address command to reset the filer's address. Use the show exports command, expect traceroute, and/or ping to check the connection to the filer. After the filer connection is re-established, you can use the enable (gbl-ns-vol-shr) command to restart the share import.
`Error: Failed to contact filer for attribute lookup.` `Error: Failed to contact portmap service during import scan.`	(NFS) These errors indicate an NFS-server problem at the filer. Once the filer's NFS service is restored, you can use the enable (gbl-ns-vol-shr) command to restart the share import.
`Error: Failed to create a file/directory on filer.`	This may indicate a full disk on the back-end filer or permissions problems. Use the show exports command to examine all permission settings at the filer. Use the following guidelines for permissions problems: • For NFS exports, check your back-end filer configuration: the back-end share should allow *root* access to all of the ARX's proxy IP addresses. Use the show ip proxy-addresses command to list all configured proxy IP addresses. • For CIFS shares, the switch uses the namespace's *proxy user* as its identity. The proxy user, created by the proxy-user command, must belong to the Administrators group. The proxy-user (gbl-ns) command sets the proxy user for a namespace. Once the permissions issue or space issue is resolved, you can re-import the share with the enable (gbl-ns-vol-shr) command.
`Error: Failed to create directory collision due to open CIFS file in path.`	(multi-protocol) A directory on this back-end share has a case collision with an already-imported directory, and a CIFS client has a file open in the already-imported directory. For example, suppose the volume imports /vol/mydir from one share, a client connects and opens a file in that share (/vol/mydir/myfile.exe), and then the volume attempts to import a second share with /vol/MYDIR. The volume must mark the first share as "NFS-only" for the import to proceed, but cannot as long as the client holds "myfile.exe" open. The second share therefore fails its import. You can use show cifs-service open-files to find the open file, close cifs file to close it, and then retry the share import (with the enable (gbl-ns-vol-shr) command).

`Error: Failed to find valid mount version during import scan.`
`Error: Failed to mount filer during import scan.`
`Error: Failed to open NFS connection to filer during import scan.`
	(NFS) These errors indicate an NFS-server problem at the filer. Once the filer's NFS service is fully restored, you can use the enable (gbl-ns-vol-shr) command to restart the share import.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Failed to resolve IP address for filer.`	The back-end device could not be located with the filer name or IP address configured for the filer. From gbl-ext-filer mode, use the ip address command to reset the filer's address. Use the show exports command and/or ping to troubleshoot the connection to the filer. After the filer connection is re-established, you can use the enable (gbl-ns-vol-shr) command to restart the share import.
`Error: Failed to retrieve Kerberos information.`	(CIFS) The namespace supports Kerberos authentication (see cifs authentication), but the namespace software is unable to confirm that the share is configured to support Kerberos, too. Check the connection to the back-end filer with show exports and/or ping. Restart the import (with enable (gbl-ns-vol-shr)) after you correct the problem.
`Error: Failed to take pathlock due to internal operation in progress.`	The share import failed due to an internal-software conflict, possibly due to a timing issue or a brief filer-connectivity issue. Retry the import by using the enable (gbl-ns-vol-shr) command on this share. If this error returns, contact F5 Support.
`Error: Failed to unmount filer during import scan.`	(NFS) Persistent errors from a back-end share caused the mount to fail. Check the NFS service at the back-end filer. Once the filer's NFS service is fully restored, you can use the enable (gbl-ns-vol-shr) command to restart the share import.
`Error: File collision detected during import.`	A file on this share has the same name and path as a file on an already-imported share. To fix this, you can manually go to the filer and rename the file, or you can set the modify flag on this volume. By setting the modify flag, you allow the volume to rename the file on import. (If any other shares are still importing, you must wait for their imports to finish before you can use the modify command.) You must also have the default settings for import rename-files and import rename-directories on this share. Then restart the import with the enable (gbl-ns-vol-shr) command.
`Error: File exists in CIFS, but not in NFS.`	(multi-protocol) The share is configured for both NFS and CIFS, but a directory that is visible in the back-end-CIFS share is not found in the NFS export. The filer probably has two different names for the directory: one for CIFS clients and one for NFS clients. Check the directory at the back-end share, and rename it so that both versions have the same name. Alternatively, you can set the modify flag on this volume. By setting the modify flag, you allow the volume to rename the directory on import. (If any other shares are still importing, you must wait for their imports to finish before you can use the modify command.) You may also need to set import rename-directories unmapped-unicode for this share; this allows the volume to rename directories whose CIFS names do not map to the character encoding for NFS. Then restart the import with the enable (gbl-ns-vol-shr) command.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: filename collision.`	A file on this share has the same name and path as a file on an already-imported share. To fix this, you can manually go to the filer and rename the file, or you can set the modify flag on this volume. By setting the modify flag, you allow the volume to rename the file on import. (If any other shares are still importing, you must wait for their imports to finish before you can use the modify command.) You must also have the default settings for import rename-files and import rename-directories on this share. Then retry the import.
`Error: Filer does not have enough free space to perform import.`	(NFS) The ARX creates a temporary database on each share during the import process: there is not enough room on this share to create the database. The database requires at least 100 MegaBytes of free space. Clear some free space at the filer share, then restart the import with the enable (gbl-ns-vol-shr) command.
`Error: Filer does not provide an NFS filehandle for a file in its directory listing.`	(multi-protocol) The filer provided a file name in an NFS READDIR call, but did not translate that file name into a filehandle in a subsequent LOOKUP call. Some multi-protocol filers behave this way when a CIFS file name has characters that are not supported by NFS; the filer returns the UTF-8 version of the file name for the READDIR call, which it later rejects in a subsequent LOOKUP call. The ARX can see the file in the directory listing, but cannot get access to the file itself. Rename the file at the filer share, then use nsck ... rebuild volume to reimport all shares in the volume.
`Error: Filer does not support specified protocols.`	The protocol(s) configured for the back-end share are not actually supported at the filer. Use the show exports command to check the protocols supported by the filer. Use the show global-config namespace command to view the required protocols for the namespace. The share must support all of the namespace's protocols. You can remove the share from the volume (with no share), or you can add protocol support at the back-end filer. If you add the protocol support to the filer, you can then restart the share import with the enable (gbl-ns-vol-shr) command.
`Error: Filer error during import.`	This may indicate a disk failure at the external filer. Correct the filer issue and retry the import (with the enable (gbl-ns-vol-shr) command).

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Filer generated name collision is not supported by filer.`	(multi-protocol) The share contains an entry whose filer-generated name (FGN) matches a real name in an already-imported share, or whose real name conflicts with an already-imported FGN. The volume is backed by a mix of Hitachi HNAS (BlueArc) filers and filers from other vendors, and this naming conflict is not supported for this vendor combination. The import report shows the name of the conflicting file or directory. The name appears with an "FC" notation. The report is named "import.*job-id*.*share-name*.*share-id*.rpt." Use show reports to list all import reports and read their contents. Using the import report, resolve all of these name collisions before re-importing. Go to the filer to rename the files and/or directories. Then use enable (gbl-ns-vol-shr) to re-import the share.
`Error: Filer operation returned an unexpected error, see logs for more information.`	The back-end filer behind this share returned an error that the ARX does not recognize. You can use the show logs syslog command to view the system log and learn more about the circumstances around the failure. We recommend that you contact F5 Support if you see this import error. You may be requested to run the collect command, which assembles diagnostic information for F5 Engineering. After the issue is resolved, use nsck ... rebuild volume to reimport all shares in the volume.
`Error: Filer returned invalid link count for a file.`	The import failed because the back-end filer returned a link count of zero or a negative number for a file, which is invalid. A file's link count should be one or more. The ARX syslog contains the file with the invalid link count; use grep "link count" logs syslog to find the path. Use this path to help diagnose and correct the filer issue. Then use the enable (gbl-ns-vol-shr) command to retry the import.
`Error: Hard Links Inconsistent`	(NFS) One or more files in the exported directory have hard links from outside the directory. That is, the switch tried to import a hard link to file X but detected that there was at least one more hard link to file X in some other directory, not included in the NFS export. Remove the hard links from the external directories and retry the share import.
`Error: Hard Links Not Supported`	(CIFS) Hard-links exist in a CIFS share (not supported in CIFS by the ARX).
`Error: I/O error encountered talking to filer.`	Use the show exports command and/or ping to check the connection to the filer. Restart the import (with enable (gbl-ns-vol-shr)) after the connection is restored.
`Error: Import database I/O failure.`	(NFS) The ARX creates a temporary database during the import process. The database was deleted while the import was underway. This may indicate that other clients are bypassing the ARX to access the share, which is an unsupported configuration. You can restart the share import with the enable (gbl-ns-vol-shr) command.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Insufficient files reserved for the volume.`	This share contains more files than the volume can hold. Use the auto reserve files command to automatically increase the number of files that this volume can hold as the volume grows. If you prefer to manually set the maximum files for the volume, use reserve files to manually increase the maximum. Then restart the share import with the enable (gbl-ns-vol-shr) command.
`Error: Insufficient free space to store metadata.`	A volume's metadata share requires gigabytes of free space; the one chosen for this volume has less than one-half gigabyte. Choose a metadata share with more free space: use the metadata share command to configure a new metadata share. After you choose another metadata share, or clear up several gigabytes of free space on the current metadata share, you can use the enable (gbl-ns-vol-shr) command to restart the import.
`Error: Internal IPC communications failure during import.` `Error: Internal thread error during import.`	These errors each indicate an internal software problem. Run the collect diag-info CLI command to collect diagnostic information, then contact F5 Support. Once the issue is understood, you can restart the share import with the enable (gbl-ns-vol-shr) command.
`Error: Maximum number of namehash collisions for a filename reached in a directory.`	A directory in this share has too many entries (files and/or subdirectories) whose names differ only in case (for example, "thisFile.txt," "THISfile.txt," "thisFILE.txt," and so on). The limit on case collisions for a single entry name is over 16,000. A client tried to create a new entry that would exceed the limit for a particular entry name, and the volume blocked the action with an "access denied" error. You can resolve this issue by accessing the volume as a client and moving as many of the file entries as possible. Create a new directory to hold many of them, or rename them to a non-matching name. For example, you could change the earlier examples to "thatFile.txt," "THATfile.txt," and so on to stop them from making a case-insensitive match with "thisfile.txt." The specific directory and file name appears in a syslog message labeled "ERROR_MAX_HASH_COLLISIONS." Use grep ERROR_MAX_HASH_COLLISIONS logs syslog to search for this error in the syslog. Note that the colliding names are usually case collisions, but not always. The collision occurs after a mathematical conversion of the file names, called a *hash*. After you clear enough colliding files from the directory, you can restart the share import with the enable (gbl-ns-vol-shr) command.
`Error: Maximum Path Length Exceeded`	A path to one of the files exceeds the maximum, 1024 characters. You must shorten the path(s) to import this share. Once all of the share's paths are below 1024 characters, you can restart the share import with the enable (gbl-ns-vol-shr) command.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Metadata database I/O failure.`	There was a database I/O failure for the share. This may be caused by a transient network error, or a filer problem. You can use the show exports command to check the filer and connection for common problems. Once the external problem is resolved, use nsck ... rebuild volume to reimport all shares in the volume.
`Error: Metadata filer configuration is invalid. Data written has been lost, corrupting the database.`	This is not an import error: it typically occurs while the volume is providing service. This condition disables the volume. Some of this share's metadata was lost at the volume's metadata share. This indicates that the metadata share did not synchronously write its data; it wrote the data to a cache, told the ARX that the write succeeded, and then lost the data between the cache and the disk. This is only possible on an NFS export where either • the "async" option was set, or • "async" is the default for the NFS implementation (as with some releases of Linux) CIFS shares do not have this problem. At the filer, set the "sync" option for the NFS export. We also recommend that you specify the "no_wdelay" option. Use nsck ... rebuild volume to re-initialize the metadata share and reimport all shares in the volume.
`Error: Metadata share unavailable during import.`	The switch could not contact the metadata share. Use the show exports command and/or ping to check the connection to the metadata share's filer. The show export command also verifies that the share is accessible by *root* (for NFS shares) or the namespace's proxy-user (gbl-ns) (for CIFS shares). After the connection is re-established, restart the share import with the enable (gbl-ns-vol-shr) command.
`Error: Mismatched NFS filename encoding.`	(multi-protocol) The NFS character-encoding setting for the namespace does not match the character encoding supported at the filer. If this share was imported, lost NFS files could result. Reset the namespace character encoding (using the character-encoding nfs command) and retry the import. You can use the enable (gbl-ns-vol-shr) command to restart the share import.
`Error: Mount to filer failed.`	(NFS) The switch was unable to NFS-mount the filer share. Check the NFS service at the back-end filer.
`Error: Multi-protocol share is split.`	(multi-protocol) The filer command specified an NFS export and a CIFS share over two different directory trees. This is unsupported. Retry the command with the correct share and export names, then retry the import with enable (gbl-ns-vol-shr).
`Error: Multi-protocol user mapping error. (nfs read failure)`	(multi-protocol) During import, the volume creates a test file through CIFS and then attempts to read it through NFS. The volume was unable to read the file as *root*. Check the NFS configuration at the back-end share, correct the problem, and retry the import (enable (gbl-ns-vol-shr)).

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Multi-protocol user mapping error. (nfs remove failure)`	(multi-protocol) During import, the volume creates a test file through CIFS and then attempts to delete it through NFS. The volume could read and write the file (as *root*), but was unable to remove it. This may indicate a permissions problem in the top-level directory for the share. Check the NFS configuration at the back-end share, correct the problem, and retry the import (enable (gbl-ns-vol-shr)).
`Error: Multi-protocol user mapping error. (nfs write failure)`	(multi-protocol) During import, the volume creates a test file through CIFS and then attempts to write to it through NFS. The volume could read the file, but was unable to write to it (as *root*). Check the NFS configuration at the back-end share, correct the problem, and retry the import (enable (gbl-ns-vol-shr)).

`Error: NFS connect failed.`
`Error: NFS error during import scan.`
`Error: NFS error occurred during import scan.`
`Error: NFS Error while scanning storage`
	(NFS) Persistent NFS errors from a back-end share caused the import to fail. Check the NFS service at the back-end filer.
	After the filer's NFS service is fully restored, you can restart the share import with the enable (gbl-ns-vol-shr) command.

`Error: NFS Version Mismatch`	(NFS) The back-end share does not support the NFS version(s) configured for the external filer. Use the show exports command to check the protocols supported by the filer. Use the filer command to change the configured NFS version(s) for the share/export.
`Error: NFSv2 error during import scan.` `Error: NFSv3 error during import scan.`	(NFS) Persistent NFS errors from a back-end share caused the import to fail. Check the NFS service at the back-end filer. After the filer's NFS service is fully restored, you can restart the share import with the enable (gbl-ns-vol-shr) command.
`Error: Network communications failure during import scan.`	Use the show exports command, expect traceroute, and/or ping to troubleshoot the connection to the filer. After the filer connection is re-established, you can restart the share import with the enable (gbl-ns-vol-shr) command.

`Error: Network error or timeout during CIFS attributes check.`
`Error: Network error or timeout during CIFS privilege check.`
`Error: Network error or timeout during CIFS write check.`
	(CIFS) A connection error occurred in the middle of a CIFS-permissions test. Use the show exports command, expect traceroute, and/or ping to troubleshoot the connection to the filer. Once the connection is fully restored, you can restart the share import with the enable (gbl-ns-vol-shr) command.

`Error: No CIFS write access.`	The namespace's proxy-user does not have adequate privileges to write to this CIFS share, so the import failed. The proxy user must belong to the Administrators group on this filer. You can choose new, more-privileged credentials for your proxy user, or you can go to the filer and add the current proxy user to a more-privileged group. The probe exports command can verify that the new proxy-user credentials pass this write test. Then restart the import with the enable (gbl-ns-vol-shr) command.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action

`Error: No space left on share, or share unreachable`
 	
`Error: No space on back-end device.`
	The ARX got an error writing a database that is created as part of the import process. This database is created on the filer share itself, so the error indicates a problem with the filer.
	
	Clear some free space at the filer share, then restart the import with the enable (gbl-ns-vol-shr) command.

`Error: Operation aborted at user request.`	An administrator issued the cancel import command to stop this share import. You can restart the import with the enable (gbl-ns-vol-shr) command.
`Error: Permission Denied`	This indicates that the ARX did not have proper permissions to connect to the back-end filer. For NFS exports, check your back-end filer configuration: the back-end share should allow access to all of the ARX's proxy IP addresses. Use the show exports command to check the filer's permissions and configuration. Use the show ip proxy-addresses command to list all configured proxy IP addresses. For CIFS shares, the switch uses the proxy user for the namespace; the proxy-user (gbl-ns) command sets these credentials. The proxy user must belong to the Administrators group. After the permissions problem is corrected, you can restart the import with the enable (gbl-ns-vol-shr) command.
`Error: Protocol specific error, see logs for more information.`	(CIFS) The back-end filer returned an unexpected CIFS error during import. The syslog shows the specific error. Use show logs syslog to read the syslog, or grep *string* logs syslog to search for a specific string in the syslog. You may need to escalate to F5 Support. After you correct the error, use nsck ... rebuild volume to reimport all shares in the volume.
`Error: Proxy User does not map to root user via NFS.`	(multi-protocol) The *proxy user* is a Windows username and password that the volume can use as its identity for share import and for policy operations. In a multi-protocol (CIFS and NFS) namespace, the proxy user on the Windows side must map to the *root* user on the UNIX side. You can select a new proxy user for the namespace with the command. If necessary, map the proxy user to root at the filer itself; the *ARX Site Planning Guide* has instructions for creating this mapping on common multi-protocol filers.
`Error: Remove incomplete: found files on share.`	An administrator failed to remove the share with no share because client-visible files are still present on the share. Use the remove-file-entries option to remove all of the file entries from the volume; this produces a client-visible effect, so do this with caution. Alternatively, you can use remove-share nomigrate.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Root Squash is Enabled`	(NFS) This indicates that the ARX did not have proper permissions to create files at the back-end filer. Check your back-end filer configuration: the back-end share should have *no-root-squash* set for all of the ARX's proxy IP addresses. (On some filers, you accomplish this by mapping the *anonymous* user to UID 0 (zero).) Use the show exports command to check the filer's permission settings. Use the show ip proxy-addresses command to list all configured proxy IP addresses. Once the filer's root-squash option is disabled, you can restart the import with the enable (gbl-ns-vol-shr) command.
`Error: RPC error during import scan.`	(NFS) Persistent NFS errors from a back-end share caused the import to fail. Check the NFS service at the back-end filer. After the filer's NFS service is fully restored, you can restart the share import with the enable (gbl-ns-vol-shr) command.
`Error: Share attributes incompatible with volume attributes.`	(CIFS) The CIFS attributes set for the volume (with compressed-files, named-streams, persistent-acls, sparse-files, and/or unicode-on-disk) are not all supported at the back-end share. Use the show exports command to check the supported CIFS attributes for the share. You can remove the share from the volume (with no share) or disable the conflicting CIFS attribute(s) in the managed volume. If you elect to keep the share in the volume, use the enable (gbl-ns-vol-shr) command to restart the share import.
`Error: Share does not support Kerberos authentication.`	(CIFS) The namespace supports Kerberos authentication (see cifs authentication), but this share does not. The share must support *extended security negotiations* for the import to succeed. Also, the ARX needs the correct service-principal name (*SPN*) for the filer; you can use show exports ... capabilities to verify that the ARX has discovered the correct SPN for the filer, or you can use the spn command to set it manually. After you ensure that the filer support Kerberos and the ARX has the filer's SPN, you can restart the import (enable (gbl-ns-vol-shr)). Alternatively, you can remove the share from the volume with no share.
`Error: Share type does not match namespace supported protocols.`	Each back-end share must support *all* of the namespace's configured protocols (any combination of NFSv2, NFSv3(/UDP), NFSv3/TCP, and CIFS). Use the show global-config namespace command to view the namespace's protocols. You can remove the share from the volume (with no share) or enable the missing service(s) at the filer. If you elect to keep the share in the volume, use the enable (gbl-ns-vol-shr) command to restart the share import.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action

`Error: Specified export is not a directory.`
 	
`Error: Specified export is too long.`
 	
`Error: Specified export not found on filer.`
	The export name is incorrect in the external-filer configuration. Use the filer command to change the configured name for the share/export.

`Error: Stat of a file/directory failed during storage init.`	(NFS) The ARX was unable to access the attributes of a file or directory. Check the NFS service at the back-end filer. To check the ARX's connectivity to the filer and perceived permissions at the filer, use show exports. You can restart the share import with enable (gbl-ns-vol-shr) after you resolve the filer issue.
`Error: Storage is already in use by this switch pair.`	The share is already imported into another managed volume on this switch (or its redundant peer). Evidently, the back-end share has at least one alias and it was previously imported under another name. A share can only be managed by one volume at a time, under a single name. From gbl-ns-vol-shr mode in the CLI, use no filer to detach from the back-end share. Then choose another back-end path with the filer command, or use no share to remove the share from the volume. To restart the import (with or without this share), use nsck ... rebuild volume to reimport all shares in the volume.
`Error: Storage remove failed because pending directory operations could not be flushed.`	Someone attempted to remove a share (with no filer, no share, remove-share migrate, remove-share nomigrate, or remove service), and an internal error caused the removal to fail. Contact F5 Support if you see this message.
`Error: Unable To Connect to CIFS Share.`	The connection to the back-end CIFS share failed due to possible configuration errors or a broken connection to the back-end filer. Use the show exports command, expect traceroute, and/or ping to troubleshoot the connection to the filer. Once the connection is fully established, you can restart the share import with the enable (gbl-ns-vol-shr) command.
`Error: Unable to read file due to stale filehandle.`	(NFS) The filer provided a filehandle in response to an earlier NFS call, and is now pronouncing it "stale." This is inconsistent filer behavior; it may indicate that NFS clients are bypassing the ARX to access the share. Check the filer share and correct the problem there. Once the filer issue is corrected, use nsck ... rebuild volume to reimport all shares in the volume.

`Error: Unexpected CIFS privilege check failure.`
 	
`Error: Unexpected CIFS write check failure.`
	A CIFS-permission test failed at the filer for an undetermined reason. This may be a filer issue or a connectivity issue. Check the filer and the connection, and retry the share import (with enable (gbl-ns-vol-shr)). If this error stops the import a second time, run the collect diag-info CLI command to collect diagnostic information and contact F5 Support.

Table 21.1 *Share Status Conditions (Continued)*

Status Condition	Description/Action
`Error: Unresolvable case-blind name collision.`	(CIFS) A file on this share has the same name and path as a file on an already-imported share, based on a case-blind comparison, and the volume is configured with no cifs case-sensitive. That is, some of the characters have differing cases, but the characters match (for example, "index.htm" matches "index.HTM"). If the volume is not case-sensitive, it cannot see the difference between the two names. You have three options to address this error: • Manually go to the filer and rename the file. • Set the modify flag on this volume. By setting the modify flag, you allow the volume to rename the file on import. You must wait for all the volume's shares to finish importing before you can use this command. • Use cifs case-sensitive to make the volume case-sensitive. Then retry the import with the enable (gbl-ns-vol-shr) command.
`Import Interrupted`	An administrator stopped the import with the cancel import command. You may be able to restart the import with no enable (gbl-ns-vol-shr) and then enable. If the import was stopped too far in the process, you must first use nsck ... destage to shut down the volume, remove and re-add the share, then enable the volume again.
`Internal Error (`*`number`*`)`	Internal problem; contact F5 personnel.
`Metadata Only`	The share is designated to store namespace metadata only, so there are no client-accessible files to import. This serves as an explanation; it is not an import error.
`Offline: Volume Disabled`	No imports are possible unless the volume is enabled. Use the enable (gbl-ns, gbl-ns-vol) command to enable the volume.
`Offline: Volume Failed`	The metadata share for the volume failed to import. Use metadata share to designate a new dedicated share for metadata.
`Remove Interrupted`	An administrator stopped the share removal with the cancel remove command. To restart the removal process, use remove-share migrate, remove-share nomigrate, no share, or no filer.
`Storage job aborted.`	The import process was interrupted by an nsck ... rebuild force. The rebuild operation will re-import the share.
`Storage job aborted: DNAS is shutting down.`	DNAS, the internal name for volume software, is shutting down. This may be the result of administrative action, such as a remove service, during the import. The volume has stopped running, so the import is canceled.
`Uninitialized (No filer assigned)`	Use the filer command to assign a filer to the share, then retry the import.

Table 21.1 *Share Status Conditions (Continued)*

show namespace mapping

Purpose Use the show namespace mapping command to view the back-end shares behind a namespace.

Mode (any)

Security Role(s) crypto-officer or storage-engineer

Syntax

```
show namespace mapping
show namespace mapping namespace
show namespace mapping namespace volume vol-path
```

namespace (1-30 characters) focuses on one namespace. Without this, the command displays the shares behind all namespaces.

vol-path (1-1024 characters) focuses on one volume.

Guidelines The output is a two-column table with the namespace and volume in the left column and the physical filer shares in the right column. For direct volumes, this shows one line per attach point. The word, [replica-snap], appears next to any replica-snap shares.

Use show server-mapping to show all of the filer shares behind client-side share, from a front-end service such as cifs or nfs.

Sample(s)

`bstnA# show namespace mapping`

shows filer shares behind all namespaces. For sample output, see Figure 21.7 on page 21-68.

`bstnA# show namespace mapping wwmed volume /acct`

shows the back-end shares behind a particular volume. See Figure 21.8 on page 21-70.

Related Commands show namespace
show server-mapping
show global-config namespace

Figure 21.7 *Sample Output: show namespace mapping*

```
bstnA# show namespace mapping

Namespace              Physical Server
--------------------   --------------------
insur:/claims
                       nas1:/vol/vol2/meta2*
                       \\nas1\insurance
                       \\nasE1\patient_records


Namespace              Physical Server
--------------------   --------------------
labResearch:/labMice
                       \\fs3\mice_lab_results
```

```
                        \\fs4\white_mice
                        nas1:/vol/vol2/meta8*

labResearch:/labRats
                        \\fs6\labRats
                        \\fs7\lrTests
                        nas1:/vol/vol2/meta9*



Namespace               Physical Server
--------------------    --------------------
medarcv:/lab_equipment
                        \\fs2\backlot_records
                        \\fs5\xraysScanners
                        nas1:/vol/vol2/meta6*
                        \\nas10\equipment
                        \\nas10\for_lease
                        \\nas11\equipBkup [replica-snap]
                        \\nas11\leasedBkup [replica-snap]

medarcv:/rcrds
                        \\fs1\histories
                        \\fs2\bulkstorage
                        \\fs4\prescriptions
                        nas1:/vol/vol2/meta3*

medarcv:/test_results
  2005charts            medarcv:/rcrds/2005
  chemLab               \\fs1\chem_results/.
  hematologyLab         \\fs3\hematology_results/.



Namespace               Physical Server
--------------------    --------------------
medco:/vol
  vol1/corp             nas1:/vol/vol2/shr
  vol1/notes            nas1:/vol/vol2/notes
  vol2                  nas3:/exports/data
  vol3/mtgMinutes       nas2:/vol/datavol1/direct/mtgs
  vol3/sales            nas2:/vol/datavol1/direct/export



Namespace               Physical Server
--------------------    --------------------
wwmed:/acct
                        das1:/exports/budget
                        das3:/exports/acct2
                        das7:/lhome/it5
                        das8:/work1/accting
                        nas1:/vol/vol2/meta1*



Where * denotes metadata only physical server.
```

Figure 21.8 *Sample Output: show namespace mapping ... volume ...*

```
bstnA# show namespace mapping wwmed volume /acct

Namespace               Physical Server
--------------------    --------------------
wwmed:/acct
                        das1:/exports/budget
                        das3:/exports/acct2
                        das7:/lhome/it5
                        das8:/work1/accting
                        nas1:/vol/vol2/meta1*


Where * denotes metadata only physical server.
```

show namespace status

Purpose Use the show namespace status command to view the import status of a namespace.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show namespace status namespace
show namespace status namespace volume vol-path
show namespace status namespace volume vol-path share
share-name
show namespace status all
```

namespace (1-30 characters) identifies a namespace.

vol-path (1-1024 characters) focuses on one volume.

share-name (1-64 characters) narrows the focus to a single share.

all displays status for all namespaces.

Guidelines A managed-volume share imports files from its back-end filer when it is enabled and its parent volume is also enabled. A direct-volume share does not import; it only connects to its attach points (see attach). Use enable (gbl-ns-vol-shr) to enable a share, and use enable (gbl-ns, gbl-ns-vol) to enable a volume.

This shows one table for each namespace and a sub-table for each volume. The top row of each volume table contains the name of the volume and its status. The Status is one of these values:

- Enabled - Available for client access.
- Starting indicates that the volume software is starting. This can occur after enabling the volume, after an ARX reload, or after an nsck ... rebuild.
- Stopping means that the volume software is shutting down. An nsck ... destage or nsck ... rebuild can cause this.
- Migrating shows that the volume is migrating its metadata from one share to another. The nsck ... migrate-metadata command starts a metadata migration.
- Disabled indicates that the volume is administratively disabled.

Shares are grouped under their volumes. For each share, this shows the Share name (or “metadata-share” for a metadata-only share), the name of the external Filer, the NFS Export or CIFS Share behind this namespace share, and the Status of the share. An [rs] appears before a replica-snap share, which holds snapshots of a standard share in the same volume.

Guidelines: Share Status

In the Status column, the following messages show the progress of a successful import:

- Pending - Scheduled for import.
- Importing - Actively inventorying the back-end storage.
- Online - Import completed successfully, back-end storage is ready to be used. A direct share goes directly to this state.

The following status conditions may also appear:

- Disabled means that someone issued a no enable for the share or its volume.
- Offline means that the volume lost contact with the back-end filer, or that someone disabled the share with no enable. Use show exports to check the connection with the filer.
- Interrupted indicates that the ARX lost contact with the filer during the import.
- Error shows that the import failed; use show namespace for a more-detailed error message.
- Removing metadata appears when the share is being removed from the namespace. (Use no share, remove-share migrate, or remove-share nomigrate to remove a share.)
- Updating Probes occurs after an upgrade to a new software release (see boot system). The initial import of the share used a now-outdated probe to prove that the share is viable for import, so the volume is now running the latest probe test. This status should change to Online very quickly; examine the share with show exports and/or probe exports if this status persists for a long period of time.
- Uninitialized means that no one has assigned a filer to the share.

Use the show namespace command for more-detailed progress and error messages.

Sample(s)

```
bstnA# show namespace status all
```

shows volume shares and filer import status for all configured namespaces. For sample output, see Figure 21.9 on page 21-73.

```
bstnA# show namespace status wwmed
```

shows volume shares and filer import status for the ‘wwmed’ namespace. See Figure 21.10 on page 21-75.

```
bstnA# show namespace status wwmed volume /acct
```

shows volume shares and filer import status for the ‘wwmed~/acct’ volume. See Figure 21.11 on page 21-75.

Related Commands

show namespace
show global-config **namespace**
enable (gbl-ns-vol-shr)
enable (gbl-ns, gbl-ns-vol)

Figure 21.9 *Sample Output: show namespace status all*

```
bstnA# show namespace status all

Namespace: medco
Description:

     Share                     Filer                                Status
         NFS Export
    ------------------------ ----------------------------------- -----------

  Volume: /vol                                                   Enabled
     corporate                 nas1                                 Online
         NFS: /vol/vol2

     sales                     nas2                                 Online
         NFS: /vol/datavol1/direct

     generic                   nas3                                 Online
         NFS: /exports

Namespace: wwmed
Description:

     Share                     Filer                                Status
         NFS Export
    ------------------------ ----------------------------------- -----------

  Volume: /acct                                                  Enabled
     budget                    das1                                 Online
         NFS: /exports/budget

     bills                     das8                                 Online
         NFS: /work1/accting

     bills2                    das3                                 Online
         NFS: /exports/acct2

     it5                       das7                                 Online
         NFS: /lhome/it5

     metadata-share            nas1                                 Online
         NFS: /vol/vol2/meta1

Namespace: medarcv
Description:

     Share                     Filer                                Status
        CIFS Share
    ------------------------ ----------------------------------- -----------

  Volume: /lab_equipment                                         Enabled
     equip                     nas10                                Online
        CIFS: equipment

     leased                    nas10                                Online
        CIFS: for_lease

     backlots                  fs2                                  Online
        CIFS: backlot_records

     scanners                  fs5                                  Online
        CIFS: xraysScanners
```

```
     metadata-share            nas1                                  Online
         NFS: /vol/vol2/meta6

     [rs] equipSnap            nas11                                 Online
        CIFS: equipBkup

     [rs] leasedSnap           nas11                                 Online
        CIFS: leasedBkup

[rs] : replica-snap share

  Volume: /rcrds                                                   Enabled
     rx                        fs4                                   Online
        CIFS: prescriptions

     charts                    fs1                                   Online
        CIFS: histories

     bulk                      fs2                                   Online
        CIFS: bulkstorage

     metadata-share            nas1                                  Online
         NFS: /vol/vol2/meta3

  Volume: /test_results                                            Enabled
     chemistry                 fs1                                   Online
        CIFS: chem_results

     hematology                fs3                                   Online
        CIFS: hematology_results

     2005_charts               medarcv                               Online
        Path: /rcrds

Namespace: insur
Description:

     Share                     Filer                                 Status
        CIFS Share
         NFS Export
    ------------------------ ----------------------------------- -----------

  Volume: /claims                                                  Enabled
     shr1-old                  nas1                                  Online
        CIFS: insurance
         NFS: /vol/vol2/insurance

     shr1-next                 nasE1                                 Online
        CIFS: patient_records
         NFS: /patient_records

     metadata-share            nas1                                  Online
         NFS: /vol/vol2/meta2

Namespace: labResearch
Description:

     Share                     Filer                                 Status
        CIFS Share
    ------------------------ ----------------------------------- -----------
```

```
  Volume: /labMice                                                Enabled
     Mresults                 fs3                                 Online
        CIFS: mice_lab_results

     metadata-share           nas1                                Online
         NFS: /vol/vol2/meta8

     Mrecords                 fs4                                 Online
        CIFS: white_mice

  Volume: /labRats                                                Enabled
     Rresults                 fs6                                 Online
        CIFS: labRats

     Rrecords                 fs7                                 Online
        CIFS: lrTests

     metadata-share           nas1                                Online
         NFS: /vol/vol2/meta9
```

Figure 21.10 *Sample Output: show namespace status wwmed*

```
bstnA# show namespace status wwmed

Namespace: wwmed
Description: namespace for World-Wide Medical network

     Share                    Filer                               Status
         NFS Export
    ------------------------ ----------------------------------- -----------

  Volume: /acct                                                   Enabled
     budget                   das1                                Online
         NFS: /exports/budget

     bills                    das8                                Online
         NFS: /work1/accting

     bills2                   das3                                Online
         NFS: /exports/acct2

     it5                      das7                                Online
         NFS: /lhome/it5

     metadata-share           nas1                                Online
         NFS: /vol/vol2/meta1
```

Figure 21.11 *Sample Output: show namespace status wwmed volume /acct*

```
bstnA# show namespace status wwmed volume /acct

Namespace: wwmed
Description: namespace for World-Wide Medical network

     Share                    Filer                               Status
         NFS Export
    ------------------------ ----------------------------------- -----------

  Volume: /acct                                                   Enabled
```

```
budget                    das1                              Online
    NFS: /exports/budget

bills                     das8                              Online
    NFS: /work1/accting

bills2                    das3                              Online
    NFS: /exports/acct2

it5                       das7                              Online
    NFS: /lhome/it5

metadata-share            nas1                              Online
    NFS: /vol/vol2/meta1
```

smb2-allowed

Purpose Use the smb2-allowed command to enable SMB2 on this ARX. SMB2 is a newer, widely-used, version of the original CIFS (SMB) protocol.

Use no smb2-allowed to disable SMB2.

Mode gbl

Security Role(s) storage-engineer or crypto-officer

Syntax

```
smb2-allowed
no smb2-allowed
```

Default(s)

```
no smb2-allowed
```

Guidelines SMB2 must be allowed for the current ARX before clients can use SMB2 to connect to any CIFS service or access any namespace storage through SMB2.

For a given CIFS service to support SMB2, it must also meet the following requirements:

- SMB2 must be enabled at the CIFS service. By default, SMB2 is enabled for every CIFS service; you can use the no smb2-allowed (gbl-cifs) command to disable SMB2 on a service-by-service basis.
- All filers behind the CIFS service must support SMB2. You can use the show cifs-service filer-capabilities *fqdn* command to find all filers behind the *fqdn* service and probe them for SMB2 support.

This command affects all future CIFS connections, but does not affect currently-connected clients. If a client is already connected with SMB when you allow SMB2, the client continues with SMB until the connection is cut, and future connections use SMB2 where appropriate. If a client has been using SMB2, you later use this command to disable it, and the client fails to drop down to SMB, you need to reboot your client to force it to drop down.

We recommend allowing SMB2 service with this command while creating your SMB2 namespaces and volumes, before you create the front-end CIFS service.

Sample

```
bstnA(gbl)# smb2-allowed
```

allows SMB2 on the "bstnA" ARX.

Related Commands smb2-allowed (gbl-cifs)
show cifs-service filer-capabilities

windows-mgmt-auth (gbl-ns)

Purpose

If a CIFS service has MMC browsing enabled, only authorized Windows clients can manage the service. You can use the windows-mgmt-auth command to create a Windows-management-authorization (*WMA*) group, a list of Windows clients with MMC permissions, and then you can use this command to apply one or more such groups to the current namespace.

You can also use this command to select privileged CIFS clients who can access snapshots. (This command has no effect on snapshot access by NFS clients.)

Use no windows-mgmt-auth to remove a WMA group from the current namespace.

Mode

gbl-ns

Security Role(s)

storage-engineer or crypto-officer

Syntax

`windows-mgmt-auth` ***`name`***
`no windows-mgmt-auth` ***`name`***

name (1-64 characters) identifies the WMA group.

Default(s)

None

Guidelines

Each WMA group has special MMC access to the namespace. The Windows clients in the group share the same MMC permissions. All cifs services backed by this namespace (if they have browsing enabled) use the WMA group(s) that you identify with this command.

You can also use WMA groups to manage CIFS-client access to snapshots. Use the permit snapshot monitor (see permit (gbl-mgmt-auth)) command to allow group members to view snapshots, and use the snapshot privileged-access command in any volume where the group(s) should access snapshots.

Use this command multiple times to associate multiple WMA groups with the namespace. The show windows-mgmt-auth command shows all available groups and their configurations.

If a single front-end cifs service has exports from more than one namespace, this set of WMA groups must be the same for all of the exported namespaces. (The export (gbl-cifs) command exports a namespace volume through CIFS.) If this command makes the current namespace inconsistent with the other namespace(s) behind the same CIFS service, the CLI prompts with an opportunity to make the same change in the other namespace(s). Enter **yes** to allow the CLI to propagate the change to the other namespace(s).

Samples

```
bstnA(gbl-ns[medarcv])# windows-mgmt-auth testers
bstnA(gbl-ns[medarcv])# windows-mgmt-auth fullAccess
bstnA(gbl-ns[medarcv])# windows-mgmt-auth readOnly
```

associates three WMA groups to the “medarcv” namespace. A Windows client in one of these groups has the MMC permissions defined in the group. No other Windows client has MMC access to the namespace.

```
bstnA(gbl-ns[medarcv])# no windows-mgmt-auth testers
```

removes a WMA group, “testers,” from the “medarcv” namespace.

Related Commands windows-mgmt-auth
show windows-mgmt-auth

22

Volume

attach

Purpose

Use the attach command to attach (or map) a back-end directory to the current share. The ARX volume's clients can then see this back-end directory as though it was one of the volume's directories.

Use the no form of the command to remove the attachment.

Mode

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

`attach` ***`attach-point-directory`*** **`[to`** ***`physical-directory`*** **`]`**
`[access-list` ***`map-name`*** **`]`**

attach-point-directory (1-256 bytes of UTF-8-encoded characters) is the relative pathname within the direct volume. This creates a new virtual directory, visible to clients.

to ***physical-directory*** (optional, 1-256 bytes of UTF-8-encoded characters) is the physical directory on the share relative to the share export point. If omitted, use the same name as ***attach-point-directory***. To attach to the root of the back-end share, use a period (.).

map-name (optional, 1-64 characters) is the NFS access list to associate with the share.

`no attach` ***`attach-point-directory`***

attach-point-directory (1-4096 bytes of UTF-8-encoded characters) is the relative pathname within the volume. This creates a new virtual directory, visible to clients.

Default(s)

None.

Guidelines

This command only applies to a share in a direct volume. It has no effect in a managed volume.

Use the show server-mapping command to view the current mappings.

Clients can mount any level of the attach-path directory. For example, suppose you create a /etc volume with a /bin/local attach point; clients can mount /etc, /etc/bin, or /etc/bin/local. Mounting a directory below that point would defeat the purpose of federating the attach points into a direct volume, so it is not supported.

Sample(s)

```
bstnA(gbl-ns-vol-shr[wwclim~/temp~zones])# attach mid-atl
```

attaches a direct-mapped directory named mid-atl to the "wwclim~/temp~zones" share.

Related Commands

namespace -> volume -> share
direct
nfs-access-list
show server-mapping

auto reserve files

Purpose

A managed volume has a limited number of *file credits*, where one credit is required for each of its files and directories. By default, the volume automatically increases its allocation of file credits as needed. This is a desirable situation unless the volume is used as follows:

- A direct volume uses this volume as it would use an external filer (with the filer or managed-volume commands), and
- the direct volume supports NFS.

The direct NFS volume cannot function properly if both of the above conditions are true; the direct volume requires an unchanging number of file credits in all of its "filers." You can use no auto reserve files for a managed volume behind an NFS direct volume. This is especially important if the direct volume is on a remote ARX, which cannot detect the configuration issue.

Use the affirmative form, auto reserve files, to allow the managed volume to automatically reserve its own file credits.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no auto reserve files
auto reserve files
```

Default(s)

```
auto reserve files
```

Guidelines

The auto-reserve-files feature is recommended for most managed volumes. You only need to disable it if the managed volume is being used as an NFS filer by a direct volume.

This command has no effect on a direct volume itself.

Sample

```
bstnA(gbl-ns-vol[ns1~/etc])# no auto reserve files
```

disables automatic file reservations in the "ns1~/etc" volume. This makes it safe to use this volume as a "filer" in an NFS-supporting direct volume.

Related Commands

namespace -> volume
direct
managed-volume
show namespace
show global-config namespace

auto sync files

Purpose

A managed volume's metadata contains information about names and locations of all files on all of its back-end filers. A filer application, such as anti-virus software, could possibly move, delete, or rename a file without the volume's knowledge; this obsoletes the metadata about the file. If a CIFS user receives an error that indicates that a file's metadata is incorrect (the file is missing), the managed volume can automatically launch a *sync* operation to synchronize the metadata with the filer contents. Use the auto sync files command to allow the current volume to automatically synchronize its metadata.

Use no auto sync files to disable automatic synchronizations.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
auto sync files [rename-files]
no auto sync files [rename-files]
```

rename-files (optional) allows the volume to rename a newly-discovered file that collides with another file in the volume. (Two files are said to *collide* when they share the same path and name on their respective back-end shares.)

Default(s)

disabled

Guidelines

For NFS-only volumes, or for volumes where automatic synchronization is disabled, you can use the sync files command to manually launch a sync-files operation. You must also use the manual command for files that are created on the back-end filer; the auto-sync operation only works for missing files. Directory syncing is not automated by this command; use sync directories to synchronize the volume's metadata with new directories that may have been created at the filer.

Each auto-sync operation has a unique job ID and generates a report of its progress. The reports follow this naming convention: "auto-sync.*sync-job-id*.*volume*.rpt." Use show reports to list all reports, including auto-sync reports. To follow the progress of the auto-sync operation, you can use tail reports report-name follow.

The show sync command shows the current status of one or more sync operations. You can use the wait-for sync command to wait for the operation to complete. To cancel the operation, use cancel sync.

If the rename option is off, a newly-discovered file is not synchronized in the metadata unless its path is unique in the volume. If the option is enabled, a conflicting file is renamed as follows: "myfile.txt" on the "bills" share becomes "myfile_bills.32.txt," assuming the auto-sync-job ID is 32.

Guidelines: Access-Based Enumeration (ABE)

CIFS shares on some back-end filers use a security feature called *Access-Based Enumeration* (ABE) that can complicate auto-sync operations. An ABE-enabled share provides customized directory listings for every client, where the directory listing only shows files and directories where the client has read access. A managed volume examines a client's directory listing before forwarding it to the client; if a file is missing from that listing due to ABE, the volume logs that the file is "missing" and performs extra processing to find that it is still in the directory. To dampen this logging and help the volume software to work around ABE-based file listings, use the cifs access-based-enum command.

Samples

```
bstnA(gbl-ns-vol[medarcv~/usr])# auto sync files rename-files
```

allows automatic metadata synchronization in the "medarcv~/usr" volume. If any newly-discovered files would collide with other files in the volume, the volume is allowed to rename them.

```
bstnA(gbl-ns-vol[medarcv~/rcrds])# no auto sync files
```

disables auto sync in the "/rcrds" volume.

Related Commands

namespace **->** volume
sync files
show sync
wait-for sync
cancel sync
cifs access-based-enum
show namespace
show global-config **namespace**

cancel import

Purpose The cancel import command stops a managed volume from importing a share.

Modes priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`cancel import namespace`** ***`ns`*** **`volume`** ***`vol-path`*** **`share`** ***`share-name`***

ns (1-30 characters) is the namespace where the import is occurring.

vol-path (1-1024 characters) specifies the namespace volume.

share-name (1-64 characters) is the share that is being imported. Use show namespace status for a list of shares that are in the process of importing.

Default(s) None

Guidelines After canceling a share import, you must remove the share from the managed volume. You have three options:

1. Use remove-share migrate to migrate all imported files to another share in the volume.
2. Use remove-share nomigrate to avoid migrating any imported files, thus causing them to disappear from the view of the volume's clients. This command copies all imported directories to another share in the volume, and if the share is later re-imported, the directories use the new share as their home.
3. Use no enable (gbl-ns-vol-shr) and nsck ... rebuild to remove all imported files and directories from the volume. This causes files and directories to disappear from client view. It also forces all of the volume's NFS clients to re-mount the share (CIFS clients are briefly disconnected, which may go unnoticed).

Sample

```
bstnA# cancel import namespace ns volume /vol share testrun
Storage job cancelled successfully.
bstnA#
```

cancels the import in progress for the "ns~/vol~testrun" share.

Related Commands

enable (gbl-ns-vol-shr)
show namespace status
remove-share migrate
remove-share nomigrate
nsck ... rebuild

cifs access-based-enum

Purpose

A back-end CIFS share with *Access-Based Enumeration* (ABE) provides customized directory listings to its clients; a directory listing only contains files and folders where the client has read access. Managed-volume software mistakenly assumes that the missing (read-only) files and directories are metadata inconsistencies, and fills them back into the directory listing. Clients of the managed volume therefore see all files and directories, whether or not they have permission to read them. This defeats the purpose of ABE. Use the cifs access-based-enum command to inform the volume that its back-end shares have ABE enabled, and prevent the volume from revealing inaccessible files and subdirectories to its CIFS clients. (This command has a lesser effect in direct volumes, explained in the *Guidelines* below.)

Use no cifs access-based-enum in a volume backed by shares with ABE disabled.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
cifs access-based-enum [auto-enable]
no cifs access-based-enum
```

auto-enable (optional) automatically enables ABE on the back-end share when you first invoke enable (gbl-ns-vol-shr) and import the share. (Each of the volume's shares must have the import sync-attributes setting for this operation to succeed.) This ensures ABE consistency for all of the volume shares that you import from now on, and is therefore recommended.

Default(s)

```
no cifs access-based-enum
```

Guidelines

This command does not apply to a volume in an NFS-only namespace.

An ARX volume with multiple shares should have a consistent ABE setting for all of them. ABE is enforced at the back-end share, not the ARX. If one back-end share has ABE enabled and another share does not, restricted files may seem to appear and disappear if they migrate between those two shares.

The CLI prompts for confirmation if the managed volume already contains any shares with conflicting ABE settings; enter **yes** to proceed. If you use the auto-enable option, the volume software replicates all ABE settings between the volume's back-end shares at share-import time. (Share-import time is the first time each share is enabled, or the first time they are enabled after someone runs nsck ... destage or nsck ... rebuild on the volume.) If the volume was enabled without ABE's auto-enable flag, or if back-end shares otherwise have inconsistent ABE settings, you can use the [no] cifs access-based-enum (priv-exec) command. This priv-exec command enables (or disables) ABE on all of the filer shares behind a volume.

This command has no effect on any shares that have already been imported. The command sets a flag that the volume software uses during share import. The volume software otherwise ignores the flag.

In a direct volume, the back-end settings for ABE have complete control over the volume's ABE behavior. This command only controls whether or not a direct volume advertises that it supports ABE (as described below).

Guidelines: ABE Restrictions on Shares

After you enable ABE in the volume, the volume enforces ABE consistency on any imported shares or subshares. The enable (gbl-ns-vol-shr) command cannot succeed for a share if its ABE setting differs from that of the volume. The auto-enable flag resolves this by enabling ABE on the filer at share-import time. Each new share must have the import sync-attributes setting for this to succeed. For shares on filers that cannot support ABE, you can use the cifs access-based-enum exclude command on the share.

This last command is not generally recommended, as it creates ABE inconsistencies that are confusing to clients, but it may be necessary to support your set of CIFS filers. ABE is not supported on filers before Windows 2003 R1, NetApp 7.2.3, or EMC 5.5. ABE is also not supported on any current release of Samba.

Some filers cannot consistently support CIFS ACLs at all, and therefore preclude any ABE support in the volume. If the volume is backed with any of these filers, you must set no persistent-acls at the volume before it can import from them. The volume therefore ignores ACLs entirely. A volume with this setting cannot reliably support ABE, because ABE depends on the permissions settings in file ACLs.

Guidelines: Volume Responses to the abecmd.exe Utility

The *abecmd.exe* utility, freely available from Microsoft, can query remote shares to determine whether or not they have ABE enabled. If a front-end export (gbl-cifs) for this volume receives such a query, this command determines how to answer it.

The *abecmd.exe* utility cannot change the ABE setting for an ARX CIFS service, because this command controls ABE at the ARX-volume level. A change to ABE in one ARX-CIFS share would change ABE in all of the shares that export the same ARX volume.

A direct volume does not probe its shares for their ABE settings, so the cifs access-based enum command has a lesser effect in direct volumes. It only determines how to answer the ABE query from the *abecmd.exe* utility. The back-end shares have complete control over whether or not ABE is enabled in a direct volume.

Samples

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# cifs access-based-enum
```

sets the “medarcv~/lab_equipment” volume for ABE processing. This is a managed volume, so it no longer allows anyone to enable (gbl-ns-vol-shr) new shares (and import them) with ABE disabled.

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# no cifs access-based-enum

This volume still contains at least one share with access-based enumeration
(ABE) enabled.  Disabling ABE on the volume without also disabling ABE on
all the volume's shares may result in user access problems.

Are you sure you want to disable ABE behavior for this volume? [yes/no] yes
```

sets the “medarcv~/lab_equipment” volume to discontinue all ABE processing. The volume presumes that file listings from all clients should be consistent for a given directory, and corrects any inconsistencies it finds. This would defeat the purpose of ABE, so ABE should be disabled for all back-end shares behind the volume. From priv-exec mode, you can use no cifs access-based-enum (priv-exec) to disable ABE at all of the volume’s back-end shares, too.

Related Commands

namespace -> volume
cifs access-based-enum (priv-exec)
cifs access-based-enum exclude
import sync-attributes

cifs access-based-enum (priv-exec)

Purpose

A back-end CIFS share with *Access-Based Enumeration* (ABE) provides customized directory listings to its clients; a client's directory listing only contains files and folders where he or she has read access. The filers behind the volume create these customized listings and the managed volume passes them back to clients. Each of the volume's shares has a separate setting to determine whether or not ABE is enabled. If ABE is enabled at the managed volume, it should also be enabled at all of the volume's backing shares. From priv-exec mode, you can use the cifs access-based-enum command to enable ABE on all of a volume's back-end shares and subshares at once.

Use no cifs access-based-enum to disable ABE on all of the back-end shares and subshares behind the volume. This is appropriate for a situation where ABE is disabled on the volume.

Modes

priv-exec

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
cifs access-based-enum ns vol-path [force]
no cifs access-based-enum ns vol-path [force]
```

ns (1-30 characters) identifies the namespace that contains the ABE-enabled volume.

vol-path (1-1024 characters) specifies the managed volume where you want to fully enable ABE.

force (optional) tells the operation to enable ABE even on filer shares that are configured but not yet enabled.

Default(s)

None.

Guidelines

This command does not apply to a direct volume or a volume in an NFS-only namespace.

The ABE-enable operation produces a report to show results from each back-end share. The report name appears after you enter the command, along with a summary of the overall results.

This command is an alternative for using Microsoft's *abecmd.exe* utility on each filer. For any share with cifs access-based-enum exclude, this attempts to disable ABE.

The no form of this command disables ABE on all of the volume's back-end shares. Directly before or after disabling ABE at the volume's filers, use the no cifs access-based-enum command to disable ABE on the volume itself. The CLI prompts for confirmation before disabling ABE on any back-end filers; enter **yes** to proceed. If the volume has ABE disabled while a backing share has ABE enabled, clients may see the following error in their directory listings:

```
02/24/2009  04:17 PM               9 myFile.doc
The parameter is incorrect.
```

This error indicates that the file or directory ("myFile.doc" in the above example) resides on a back-end share with ABE enabled.

CIFS clients only see the results of this command if they connect *after* you invoke it.

Guidelines: Windows Server Clusters

Some Windows filers are nodes in a larger Windows Server Cluster. A *server cluster* is a Windows redundancy feature. The ABE-enable or ABE-disable operation only occurs on the currently-active node for any server cluster behind the volume. If the cluster fails over, ABE support is likely to be inconsistent on the newly-active node. After you run this command, go to the cluster's administrative interface and manually enable ABE; this applies the ABE setting to all nodes in the cluster. If the cluster's interface does not offer an option for ABE, use the abecmd.exe utility to individually set it at each node. As an alternative, you can trigger a failover at the cluster and then re-run this command.

Samples

```
bstnA# cifs access-based-enum medarcv /rcrds

% INFO: Changed access-based enumeration settings for 13 of 13 shares.
There were 0 errors and 0 shares were left unchanged. See report
'cifsAbeChange_20100227014459.rpt'.
```

enables ABE on every filer share behind the "medarcv~/rcrds" volume. Refer to Figure 22.1 for a sample report.

```
stkbrgA# cifs access-based-enum bgh /naumkeag_wing force
...
```

enables ABE on every filer share behind the "bgh~/naumkeag_wing" volume. This also enables ABE on shares that are not yet enabled in the /naumkeag_wing volume.

```
bstnA# no cifs access-based-enum medarcv /lab_equipment

Volume medarcv~/lab_equipment is not configured for access-based enumeration (ABE).
Disabling ABE on this volume's shares will bring them into agreement
with the volume configuration.

Proceed to disable ABE for this volume's shares? [yes/no] yes

% INFO: Changed access-based enumeration settings for 3 of 3 shares. There were 0 errors and 0
shares were left unchanged. See report 'cifsAbeChange_200903031058.rpt'.
```

disables ABE on every filer share behind the "medarcv~/lab_equipment" volume. The "medarcv~/lab_equipment" volume is now prepared to have ABE disabled.

Related Commands

namespace -> volume
cifs access-based-enum
cifs access-based-enum exclude

Figure 22.1 *Sample Report: cifsAbeChange*

```
bstnA# show reports cifsAbeChange_20100227014459.rpt
**** CIFS Access-based Enumeration Change Report: Started at Sat Feb 27 01:44:59 2010 ****
**** Software Version: 5.02.000.12541 (Feb 23 2010 20:12:44) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:


   Command:  cifs-access-based-enum medarcv /rcrds
       ABE:  Enabled on this volume.
```

```
Access-based enumeration (ABE) is now ENABLED for the following shares:

    Share:  CELEBS$
    Filer:  fs1 (192.168.25.20)
    (Subshare of histories at relative path "VIP_wing")

    Share:  histories
    Filer:  fs1 (192.168.25.20)

    Share:  Y2004
    Filer:  fs1 (192.168.25.20)
    (Subshare of histories at relative path "2004")

    Share:  Y2005
    Filer:  fs1 (192.168.25.20)
    (Subshare of histories at relative path "2005")

    Share:  bulkstorage
    Filer:  fs2 (192.168.25.27)

    Share:  CELEBS$
    Filer:  fs2 (192.168.25.27)
    (Subshare of bulkstorage at relative path "VIP_wing")

    Share:  Y2004
    Filer:  fs2 (192.168.25.27)
    (Subshare of bulkstorage at relative path "2004")

    Share:  Y2005
    Filer:  fs2 (192.168.25.27)
    (Subshare of bulkstorage at relative path "2005")

    Share:  CELEBS$
    Filer:  fs4 (192.168.25.29)
    (Subshare of prescriptions at relative path "VIP_wing")

    Share:  prescriptions
    Filer:  fs4 (192.168.25.29)

    Share:  Y2004
    Filer:  fs4 (192.168.25.29)
    (Subshare of prescriptions at relative path "2004")

    Share:  Y2005
    Filer:  fs4 (192.168.25.29)
    (Subshare of prescriptions at relative path "2005")

    Share:  Y2006
    Filer:  fs4 (192.168.25.29)
    (Subshare of prescriptions at relative path "2006")

**** Total processed:              13
**** Elapsed time:           00:00:01
**** CIFS Access-based Enumeration Change Report: DONE at Sat Feb 27 01:45:00 2010 ****
```

cifs access-based-enum exclude

Purpose

A back-end CIFS share with *Access-Based Enumeration* (ABE) provides customized directory listings to its clients; a directory listing only contains files and folders where the client has read access. A managed volume with ABE enabled should have ABE enabled at all of its back-end shares. You can use this command to exclude a share whose backing filer cannot support ABE (such as a Samba filer). This option is designed for a tiered volume, where a lower tier of storage may be on an older or less-expensive filer that does not support ABE.

Use the no form to declare that the current share's back-end filer can now support ABE.

Modes

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
cifs access-based-enum exclude
no cifs access-based-enum exclude
```

Default(s)

```
no cifs access-based-enum exclude
```

Guidelines

This command does not apply to a share in a direct volume or an NFS-only namespace. It only applies to a managed-volume share where cifs access-based-enum is enabled.

This command is not recommended for all sites: a managed volume should typically have consistent ABE settings with all of its back-end filers.

WARNING

If ABE is enabled on share A but disabled on share B, a client without permission to read the "xyz.doc" file sees the file disappear if it migrates to share A. The same client sees the "xyz.doc" file reappear if it migrates back to share B where ABE is disabled. This inconsistency can be confusing.

The command is designed for volumes where the above confusion is unavoidable. To mitigate the confusion, place all ABE-excluded shares on the lowest tier of storage in the volume. Do not mix any ABE-enabled shares on the same tier. Also, do not use any combination of place-rules or share-farms that mix files between shares with different ABE settings, other than the place-rule(s) that enforce the tiering policy.

If the managed volume has ABE enabled but the current share's filer cannot support ABE, you require this command to allow an enable (gbl-ns-vol-shr) for the share. A managed volume stops a share from importing if the volume has ABE enabled and the back-end share does not.

The no form of this command informs the volume that ABE is now supportable at the back-end share. This can occur after a filer upgrade. If the volume has ABE enabled, you can use this command to remove the current share's exclusion from ABE. If necessary, you can then use the cifs access-based-enum (priv-exec) command to enable ABE on all the volume's back-end shares at once.

Sample

```
bstnA(gbl-ns-vol-shr[ns2~/vol1~sh5])# cifs access-based-enum exclude
```

excludes the “sh5” share from ABE support in the “ns2~/vol1” volume.

Related Commands

namespace -> volume
cifs access-based-enum
enable (gbl-ns-vol-shr)
cifs access-based-enum (priv-exec)

cifs case-sensitive

Purpose

Some back-end CIFS shares support case-sensitive file names, where “file.txt” and “FILE.txt” are stored as two different files. Other CIFS filers and file servers use both names to refer to the same file. (The same is also true of many CIFS-client applications.) A volume that supports CIFS cannot support case-sensitive names unless *all* of its back-end shares also support them. If that is the case, and if your CIFS clients also support them, you can enable case-sensitive naming support in this volume with cifs case-sensitive.

Use no cifs case-sensitive to stop the volume from supporting case-sensitive names.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
cifs case-sensitive
no cifs case-sensitive
```

Default(s)

```
no cifs case-sensitive
```

Guidelines

This command can only function in a CIFS-only namespace. It does not apply to an NFS-only namespace, and is not allowed in a multi-protocol (CIFS and NFS) namespace.

You can only enable CIFS-case sensitivity if all the back-end shares behind the volume support it.

Guidelines: Disabling Case Sensitivity

You may want to add a share that cannot support case-sensitivity to a CIFS volume that previously supported it. For those situations, you can disable case sensitivity with the no cifs case-sensitive command. The volume must be offline to disable case sensitivity: first use nsck ... destage to take the volume offline, then run the no cifs case-sensitive command, and then re-enable all of the volume’s shares.

In a CIFS volume with case sensitivity disabled, two files or directories whose names differ only in case are said to have a *case collision*. For example, there would be a case collision between the “\dir” directory and “\Dir.” If you disable case sensitivity in a running volume, your pre-existing files and directories may have one or more case collisions. After bringing the volume’s shares back online, the volume reacts to case collisions based in the settings of certain CLI commands: if no modify is set, any share with a case collision fails its import; if modify is set, the volume uses the renaming rules set by the CLI commands below.

- import rename-files
- import sync-attributes
- import rename-directories

Check each share’s import report for a list of all renamed or altered files or directories in the share. Use show reports type lmp for a list of all import reports; there is one for each imported share. Look in the import report for files and/or directories labeled “CC” (Case Collision).

Samples

```
bstnA(gbl-ns-vol[win-ns~/vol])# cifs case-sensitive
```

sets the current CIFS volume to support case-sensitive naming. A file named "myFile" can co-exist with a different file named "MyFile."

```
bstnA(gbl-ns-vol[medarcv~/rcrds])# no cifs case-sensitive
```

resets the "medarcv~/rcrds" volume so that it treats "myFile" and "MyFile" as the same file.

Related Commands

namespace -> volume
nsck ... report metadata-only

cifs deny-symlinks

Purpose

In a multi-protocol (CIFS and NFS) volume, CIFS clients can use symlinks created by NFS clients. For example, an NFS client can create a symlink named “alink” that points to a directory named “adir/subdir.” Any client can then access the “adir/subdir” directory through the “alink” symlink, as though “alink” was an actual directory. You can use the cifs deny-symlinks command to deny symlink access to the current volume’s CIFS clients. This command makes it impossible for CIFS clients to traverse a symlink.

Use no cifs deny-symlinks to allow CIFS clients to follow NFS symlinks.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
cifs deny-symlinks
no cifs deny-symlinks
```

Default(s)

```
no cifs deny-symlinks
```

Guidelines

This command is only meaningful in a multi-protocol namespace. It has no effect on a volume in an NFS-only or CIFS-only namespace.

ARX CIFS clients can never see a symlink that points to a file or directory outside their front-end export. For example, consider a CIFS share of the “\mydir\subdir” directory, where the share is named CIFSVIEW. ARX clients see the contents of “\subdir” in the root of any drive mapped to CIFSVIEW. If a symlink named “lowersub” points to “\mydir\anotherDir,” the clients of the CIFSVIEW share never see the “lowersub” symlink. The symlink would lead to a directory next to the client’s CIFSVIEW share, which may not exist on the client machine, so the managed volume does not display the link to its CIFS clients.

You can use the nsck ... report symlinks command to get a list of all the symlinks in a namespace or volume. You can also use the find command to find the target file or directory for a particular symlink. To see how many CIFS clients are using NFS symlinks, along with some related statistics, you can use the show statistics cifs symlinks command.

Samples

```
bstnA(gbl-ns-vol[mp-ns~/vol])# cifs deny-symlinks
```

prevents all CIFS clients from viewing or using any NFS symlinks in the “mp-ns~/vol” volume.

```
bstnA(gbl-ns-vol[insur~/claims])# no cifs deny-symlinks
```

re-instates CIFS symlinks in the “insur~/claims” volume.

Related Commands

namespace -> volume
nsck ... report metadata-only
nsck ... report symlinks
find
show statistics cifs symlinks

cifs file-system-name

Purpose CIFS client applications can query a managed volume for the name of its file system. By default, the volume chooses a name based on the CIFS capabilities of its back-end shares; for example, it advertises "NTFS" if and only if all of its back-end shares can support NTFS-related features. Some client applications require specific settings in this string in order to function properly. On the advice of F5 Support, you can use the cifs file-system-name command to manually set the file system name for this volume.

Use no cifs file-system-name to return to the automatic default.

Mode gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
cifs file-system-name name
no cifs file-system-name
```

name (1-16 characters) is the file-system name that this volume advertises to client applications. Use a name that is well-known, such as "NTFS" (see the *Guidelines*, below).

Default(s) the volume calculates an appropriate setting for the string (no cifs file-system-name)

Guidelines Use this command *only under the guidance of F5 personnel*. It has the potential to confuse client applications and cause them to malfunction.

By default, the volume advertises one of three file-system names, depending on the CIFS options supported by its back-end shares:

- NTFS - if the volume's shares all support named-streams, persistent-acls, unicode-on-disk, and they all preserve case in their file/directory names ("MyFile.txt" is *not* converted to "MYFILE.TXT").
- FAT32 - if the volume's shares support unicode-on-disk and case preservation, but not the others.
- FAT - if the volume supports none of the above options.

Some applications require "NTFS" to function properly, and do not require all of the above NTFS capabilities. For example, if one of the volume's shares does not support named streams, the volume advertises FAT32. The Robocopy application reads the "FAT32" file-system name and assumes that the managed volume cannot support large (4GB+) files. Support for named streams is irrelevant to the application. For Robocopy to create large files in this volume, you must set "NTFS" for the volume.

In any case, use only one of the names above, in uppercase: "NTFS," "FAT32," or "FAT"

Sample

```
bstnA(gbl-ns-vol[insur~/claims])# cifs file-system-name NTFS
```

advertises an "NTFS" file system for the "/claims" volume in the "insur" namespace.

Related Commands

namespace -> volume
named-streams
persistent-acls
unicode-on-disk

cifs notify-change-mode

Purpose A CIFS filer or service can send change notifications to its clients on request. A *change notification* is an event from the CIFS service indicating some change in the file system, such as a renamed file or a new directory. This is a standard offering from CIFS implementations; file-managing applications use this to provide regularly-updated views of remote file systems. A CIFS volume on an ARX may create performance issues with this feature, since it collects change notifications from multiple back-end shares and forwards their aggregate to its clients. The volume and/or its clients can be overwhelmed with change notifications. By default, the volume reduces this traffic by ignoring any changes to a directory's subtree; it only sends notifications of changes in a directory's root. On the advice of F5 Support, you can use the cifs notify-change-mode command to either remove all change-notification traffic or increase it dramatically.

Use no cifs notify-change-mode to return to the default, which is sufficient for most volumes.

Modes gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
cifs notify-change-mode {use-subtree-flag | no-changes-sent}
no cifs notify-change-mode
```

use-subtree-flag and **no-changes-sent** are a required choice:

use-subtree-flag causes the CIFS volume to honor a client's request for subtree notifications. By default, the volume silently ignores these client requests and only sends notifications for changes in a directory's root. You should do this only on the advice of F5 Support, as it dramatically increases the change notifications from the back-end filers to the ARX, and from the ARX to its clients.

no-changes-sent shuts down change notifications from the volume.

Default(s) `no cifs notify-change-mode`

Guidelines The default should be sufficient for most (if not all) installations. Use this command only on the advice of F5 Support.

Samples

```
bstnA(gbl-ns-vol[win-ns~/vol])# cifs notify-change-mode no-changes-sent
```

disables CIFS-change notifications from the "win-ns~/vol" volume.

```
bstnA(gbl-ns-vol[medarcv~/rcrds])# no cifs notify-change-mode
```

causes the "medarcv~/rcrds" volume to support limited change notifications. Any client can request change notifications for any given directory, but the volume only sends notifications for changes in the directory's root. If the client requests subtree notifications, too, the volume silently ignores that part of the request.

Related Commands namespace -> volume

cifs oplocks-disable

Purpose The ARX supports CIFS opportunistic locks (*oplocks*) by default. You can use the cifs oplocks-disable command to deny oplocks to this volume's CIFS clients.

Use no cifs oplocks-disable to reinstate oplock support.

Mode gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
cifs oplocks-disable [auto]
no cifs oplocks-disable
```

auto (optional) disables oplocks on a per-client basis, in response to client timeouts (see the *Guidelines*, below).

Default(s) the volume supports oplocks (no cifs oplocks-disable)

Guidelines This command only applies to a namespace that supports CIFS.

The default is sufficient for most installations. Use this command only on the advice of F5 Support.

With oplocks enabled, CIFS-client A can reserve an oplock on file X. While the client holds the oplock, no other client application can write to the file; Client A can safely write to a local cache while it holds the oplock. When another client tries to write to file X, the CIFS volume holds the request, tells Client A that its oplock is revoked, takes all writes from Client A, then allows the other client to proceed.

When oplocks are disabled, the volume does not grant any oplocks to any CIFS clients.

The auto keyword causes the volume to enable or disable oplocks individually for each CIFS client. The volume disables oplocks for a CIFS client that fails to an "oplock break" command in 10 seconds. (The "oplock break" command tells the client to finish its writes and release the oplock to another client.) For 10 minutes, oplocks are disabled for the CIFS client. After the 10 minutes expire, the volume re-enables oplocks for the client.

Samples

```
bstnA(gbl-ns-vol[insur~/claims])# cifs oplocks-disable
```

shuts off oplock support for the "/claims" volume in the "insur" namespace.

```
bstnA(gbl-ns-vol[medarcv~/rcrds])# cifs oplocks-disable auto
```

makes the "medarcv~/rcrds" volume manage oplocks on a client-by-client basis, automatically.

```
bstnA(gbl-ns-vol[ns1~/])# no cifs oplocks-disable
```

reinstates the oplock feature in the "ns1~/" volume.

Related Commands namespace -> volume

cifs path-cache

Purpose

Each NSM processor learns about a managed volume's file and directory paths from the volume software. The volume software runs on an ACM processor. The NSM processor caches all file/directory paths as it learns them, to avoid repetitive queries to the ACM. You can use no cifs path-cache to stop caching these paths at the NSM processors.

The affirmative command, cifs path-cache, resumes path caching at all NSM processors.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no cifs path-cache
cifs path-cache
```

Default(s)

```
cifs path-cache
```

Guidelines

Use this command only on the advice of F5 Support.

You can use show cifs-service path-cache to view the current state of the CIFS-path cache. For path-cache statistics since the last reboot, use show statistics cifs path-cache. The clear statistics cifs path-cache command clears all path-cache statistics.

SMB2 does not use the path cache under any circumstances, so this command has no effect on SMB2 clients (see smb2-allowed and smb2-allowed).

Additionally, this setting has no effect on a direct volume.

Samples

```
bstnA(gbl-ns-vol[win-ns~/vol])# no cifs path-cache
```

disables all path caching for the "win-ns~/vol" volume.

```
bstnA(gbl-ns-vol[win-ns~/vol])# cifs path-cache
```

enables path caching for the same volume.

Related Commands

namespace -> volume
show cifs-service path-cache
show statistics cifs path-cache
clear statistics cifs path-cache

compressed-files

Purpose

A volume that supports *compressed files* allows its clients to compress its files and preserves the file compression for policy migrations and shadow copies. If any back-end-CIFS filer does not support compressed files, you must disable the feature for its namespace volume. This applies to volumes in namespaces that support CIFS, not in NFS-only namespaces. Use the no compressed-files command to stop the volume from using compressed files.

Use the affirmative form, compressed-files, to reinstate compressed files.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

`no compressed-files`
`compressed-files`

Default(s)

The compressed-files setting at the first-enabled share.

Guidelines

A Windows client can compress any seldom-used file to conserve on disk space (on Windows XP, right-click the file -> Properties -> General tab -> Advanced... button -> Compress ... checkbox). A volume without compressed-file support does not preserve this compression if the file migrates between back-end shares (with a place-rule), or if the file is copied to a shadow volume (with a shadow-copy-rule).

You cannot enable a share if it supports compressed files and its back-end-CIFS filer does not. The enable operation fails with an error that lists all CIFS features that must be disabled, possibly including this one. Use the enable (gbl-ns, gbl-ns-vol) command to enable all shares in a new namespace, or use the enable (gbl-ns-vol-shr) command to enable a new share in an already-enabled namespace.

If you remove the share(s) or upgrade the back-end filer(s), you can reinstate this feature for the volume. Use remove-share migrate or remove-share nomigrate to remove a share from a namespace.

You can use the show exports command to see all CIFS options for the share.

Samples

`bstnA(gbl-ns-vol[medarcv~/lab_equipment])# no compressed-files`

shuts off compressed files for the "/lab_equipment" volume in the "`medarcv`" namespace.

`bstnA(gbl-ns-vol[medarcv~/rcrds])# compressed-files`

reinstates compressed files for the "/rcrds" volume in the "medarcv" namespace.

Related Commands

namespace -> volume
enable (gbl-ns, gbl-ns-vol)
enable (gbl-ns-vol-shr)
show exports

critical

Purpose

Use this command to designate the current share as "critical," so that a failover may occur if the share connection fails.

Use the no form of the command to make the share non-critical.

Mode

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
critical
no critical
```

Default(s)

```
no critical
```

Guidelines

This command is only meaningful on a switch with a redundant peer.

If a critical share fails on the current peer *and* the redundant peer has no failures, control fails over to the other peer. A peer with "no failures" has no service-affecting software or hardware faults, and has access to all critical shares and critical routes. For example, if the current switch loses contact with a critical share but its peer has lost contact with the quorum disk, no failover occurs. (This prevents a critical-share failure from causing an unnecessary failover.)

The ARX tests for failure with a regular ICMP ping. Every 30 seconds, the ARX sends a ping to the share's filer. If the filer fails to respond, the ARX waits an additional 30 seconds before asking the peer if it is possible to fail over. The pings continue indefinitely at 30-second intervals. If the filer responds before the failover is initiated, the failover does not occur.

Use the critical route command to establish a critical route. To make a dedicated metadata share into a critical resource, use the metadata critical command.

To show all critical shares and routes, use the show redundancy critical-services command.

Samples

```
bstnA(gbl-ns-vol-shr[archives~/etc~s1])# critical
```

designates the current share, archives~/etc~s1, as a critical share. If the ARX loses contact with the back-end share *and* its redundant peer has no serious issues, a failover occurs.

```
bstnA(gbl-ns-vol-shr[ns~/vol1~shareA])# no critical
```

removes the current share from the list of critical resources.

Related Commands

namespace -> volume -> share
critical route
metadata critical
show redundancy critical-services

direct

Purpose

Use the direct command to indicate that the current volume contains direct-mapped shares. This establishes the volume as a direct volume rather than a managed volume.

Use the no form of the command to indicate the volume contains no direct-mapped shares.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
direct
no direct
```

Default(s)

no direct

Guidelines

This command can be executed only when the volume has no filers assigned in any of its shares.

You cannot create a direct volume in a namespace that supports NFSv2 (see the documentation for protocol).

Use the attach command in the volume's shares to directly attach the volume to various mount points on back-end filers.

A direct volume can support up to 4,278,190,080 total inodes, or files.

To use an ARX managed volume as the back-end for one of the direct volume's shares, use the managed-volume command.

Sample

```
bstnA(gbl-ns-vol[wwclim~/temp])# direct
```

indicates the wwclim~/temp volume contains direct-mapped shares.

Related Commands

namespace -> volume
attach
managed-volume

enable (gbl-ns-vol-shr)

Purpose Use the enable command to activate the current share.

Use no enable to disable the current share.

Mode gbl-ns-vol-shr

Security Role(s) storage-engineer or crypto-officer

Syntax

```
enable [take-ownership]
no enable
```

take-ownership (optional) applies only to a share in a managed volume (as opposed to a direct volume). This causes the managed volume to take ownership of the back-end share. Use this option only if you are sure that the share is not in active use by a managed volume on another ARX. For example, some sites use filer applications to replicate all data from one site to another; if an ARX had managed volumes at the primary site, the ARX's ownership marker (a file) would be copied to the second site. An ARX at the second site could only import the replicated share if you use the take-ownership option.

Important

The take-ownership *option could possibly remove a share from another managed volume that is in service. Use this option only for cases where the share is spuriously marked by another ARX. The CLI prompts for confirmation if you use this option; enter* ***yes*** *to proceed.*

Default(s) Disabled

Guidelines You must enable a share for the volume to import it or otherwise use it. This is true for both managed volumes and direct volumes. We recommend that you enable all of the volume's shares first, then enable the volume with the enable (gbl-ns, gbl-ns-vol) command. This allows the volume to import all of the shares simultaneously, and gives the volume's clients immediate access to all of the shares at once.

Important

For shares backed by NetApp or EMC, you may need to access the filer directly and pre-create some qtrees or EMC quota trees. This rare configuration issue only occurs if:
- this is a managed volume,
*- you want to support both free-space quotas (*freespace cifs-quota*), and*
- you also want to support filer-subshares *in this volume.*
In this case, a NetApp share requires one qtree per subshare, and an EMC share must be an EMC File System with one tree quota per subshare. Pre-create the NetApp qtrees and/or EMC quota trees before you enable the share. See the Guidelines: Subshare Replication with Free-Space Quotas section of the filer-subshares *documentation.*

Guidelines (Cont.)

When a share is first imported into a managed volume, the volume generates an import report. The import reports are accessible through show reports. Their names have the following formats:

`import.`*`share-id`*`.`*`share-name`*`.`*`job-id`*`.rpt`

Use show reports *report-name* to see its contents. See Figure 22.2 on page 22-29 for a sample import report.

The no enable command stops using the current share in the volume. All files on the share become inaccessible to clients. This is not generally recommended in a managed volume, where it also shuts down all policy rules in the volume.

Use remove-share migrate or remove-share nomigrate to remove the share.

Guidelines: Enabling a New Share in an Exported CIFS Volume

This section only applies to a share that supports CIFS storage.

A new filer behind a running CIFS service often requires some configuration at a local Domain Controller (DC). Best practices dictate that a CIFS service is trusted to delegate CIFS connections to all of the filers behind it, and it is not trusted to delegate CIFS to any other filer. This is called *constrained delegation*. The DCs manage the CIFS service's ability to delegate, as well as the back-end filers to whom it can delegate. An authorized Windows administrator must therefore go to a DC and add this new filer to its CIFS service's "delegate to" list. If the filer's storage is exported by more than one CIFS service, the administrator must add it to the "delegate to" list for all of them.

After you enable the share, you can use show cifs-service *fqdn* to list the volumes exported by the *fqdn* CIFS service. For every service that exports the current volume, you can use probe delegate-to *fqdn* to check the DC(s) and determine whether or not delegation is properly configured for the new filer. If not, a properly authorized administrator must access the DC and add the new filer to the CIFS service's "delegate to" list.

The filer must be joined to the same Windows Domain as the above CIFS service(s), and it must support Kerberos authentication. The show cifs-service *fqdn* command shows the domain to which the *fqdn* service is joined.

Guidelines: Changing the Maximum Files in a Direct Share

This guideline only applies to a share in a direct volume.

When you enable a share in a direct volume, the volume software records the maximum number of files in the share. To increase that maximum, you must disable the share (with no enable), change the maximum at the back-end filer (through direct access), and then re-enable the direct share (enable).

Samples

```
bstnA(gbl-ns-vol-shr[archives~/multimedia~lun77])# enable
```

enables the current share.

```
prtlndA(gbl-ns-vol-shr[insur_bkup~/insurShdw~backInsur])# enable
take-ownership
This command allows the switch to virtualize shares that are used by
other Acopia switches.
Allow switch to take ownership of share? [yes/no] yes
```

enables the current share and takes ownership of it.

```
bstnA(gbl-ns-vol-shr[archives~/radio~xyz])# no enable
```

disables the archives`~/radio~xyz` share.

Related Commands namespace -> volume -> share
show namespace

Figure 22.2 *Sample Report: import*

```
bstnA# show reports import.6.charts.12.rpt
**** Share Import Report: Started at 06/17/2011 00:48:36 -0400 ****
**** Software Version: 6.01.000.14028 (Jun 13 2011 20:09:24) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

**** Namespace: medarcv
**** Volume:    /rcrds
**** Share:     charts
**** IP Addr:   192.168.25.20
**** Export:    histories
**** Options:
****            modify:                          yes
****            rename-directories:              yes
****            rename-non-mappable-directories: no
****            rename-files:                    yes
****            sync-attributes:                 yes
****            skip-managed-check:              no
****            import protection:               yes
****            import treewalk-threads:         8
****            import priority:                 65535 (Lowest)

**** NOTE: Since both sync-attributes and rename-directories were specified,
****       only directories that collide with existing filenames will be
****       renamed.  Directories with colliding attributes will have their
****       attributes synchronized to the namespace.

**** Ignore List:
****            .snapshot


Share                Physical Filer
-------------------  ------------------------------------------------------------
[charts           ]  192.168.25.20:histories

**** LEGEND
****
****  Actions
****   R   : Entry renamed.
****   S   : Directory attributes synchronized.
****   ?   : Indicates additional settings or operations required for full
****         operation.  See specific issues for SD and IN below.
****   !   : Unable to resolve given specified options or configuration.
****
****  Entry Type
****   F   : Entry is a file.
****   D   : Entry is a directory.
****
****  Issue
****   NC  : Name collision.
****   AC  : Attribute collision.
****   RA  : Attributes of share root are inconsistent.
****   MG  : Subdirectory of this share is already imported as managed share.
****   CC  : Case-blind collision (MPNS and CIFS-only).
****   MC  : Maximum name collisions for this name in a directory
```

```
****   ER  : Entry removed directly from filer during import.
****   AE  : Error accessing entry.
****   RV  : Reserved name on filer not imported.
****   DF  : DFS link found during import.
****   TC  : Trailing character: space and period are not supported by all filer vendors.
****   HL  : Exceeded limit of 1024 hard links to a file.
****   PF  : Directory promote failed and master directory is mis-located
****   DA  : Directory attributes sync failed and are inconsistent between shares
****   NF  : Specified Path was not found.
****   IG  : Directory ignored during import.
****   @@  : Other error.


Import Scan:
================================================================================

------------------------------------------------------------
Import Scan Start Time:   06/17/2011 00:48:39 -0400
------------------------------------------------------------

Type      Path
--------  ------------------------------------------------------------------------
[  D TC]  /2007/late_test_results.
[  D TC]  /2004/planR.
[  F TC]  /2007/hlo_medical_record.dat.

Directories found:                      22
Files found:                           143
Directories Scanned/Second:             22
Files Scanned/Second:                  143
Total Entries Scanned/Second:          165

------------------------------------------------------------
Import Scan Stop Time:    06/17/2011 00:48:39 -0400
Import Scan Elapsed Time: 00:00:01
------------------------------------------------------------

Directories imported by scan:                                    22
Directories imported dynamically:                                 0
Files imported by scan:                                         143
Files imported dynamically:                                       0
Directories renamed due to name/attribute conflict:               0
Files renamed due to name conflict:                               0
Directory attributes synchronized:                                0


Sync Subshare Report Name:  syncSshrNewStorageReport_201106170448

**** Elapsed time:           00:00:04
**** Share Import Report: DONE at 06/17/2011 00:48:40 -0400 ****
```

filer

Purpose Use the filer command to bind the current share to a back-end-filer share.

Use the no form of this command to remove the binding to any source export/share.

Mode gbl-ns-vol-shr

Security Role(s) storage-engineer or crypto-officer

Syntax: filer

```
filer filer {nfs share-name | cifs share-name}+
  [access-list list-name] [cluster cluster-name]
```

filer (1-64 characters) is the name of the filer. Use show external-filer for a list of configured filers to choose from.

You can enter either **nfs *share-name*** or **cifs *share-name*** to identify a share at the filer. You can enter both (in any order) to specify a multi-protocol share on the filer:

nfs *share-name* (1-900 characters) is an NFS export at the filer.

cifs *share-name* (1-1024 characters) is a CIFS share at the filer.

list-name (optional, 1-64 characters) applies an NFS-access list to the NFS export. Use show nfs-access-list for a list of all configured NFS access lists.

cluster-name (optional, 1-64 characters) is only relevant if the ARX is part of a disaster-recovery (DR) configuration. In a DR configuration, there is an active ARX cluster with one set of filers and a backup cluster with a mirrored set of filers. This determines which cluster uses this *filer*. Run the filer command twice per share if you use DR: once to designate the filer for the active cluster, and again to determine the filer at the backup cluster. Use show cluster for a list of configured clusters. If you omit this, the command assumes that this is the local cluster.

Syntax: no filer

```
no filer [cluster cluster-name]
no filer relocate-dirs target-share [remove-file-entries
  [verbose] [offline]]
```

cluster-name (optional) is only relevant if the ARX is part of a disaster-recovery (DR) configuration. This specifies the cluster where the filer binding is removed. If you specify a remote cluster, the filer binding is simply removed from the remote cluster's configuration and none of the remaining options are relevant. If you omit this option, the CLI removes the binding from the local cluster.

relocate-dirs is required if you are disconnecting from an already-imported share in a managed volume. This is never required for a direct volume or a replica-snap share.

> ***target-share*** is another share in the same managed volume. The volume migrates all master directories in the current share to this target.

You can enter the optional flags (**remove-file-entries**, **offline**, and **verbose**) in any order. As above, these options only apply to managed volumes and do not apply to replica-snap shares:

> **remove-file-entries** removes all files from volume metadata that still reside on this back-end share.
>
> **verbose** (optional, if you choose **remove-file-entries**) causes the operation to list all removed files in its "removeShare" report.
>
> **offline** (optional, if you choose **remove-file-entries**) is only for back-end shares that are offline or otherwise unreachable. This forces the disconnect without scanning the back-end for its directory attributes. Relocated master directories therefore have all of their file attributes set to 0 (zero).

Default(s)

None

Guidelines

The filer's export/share must support all of the namespace's protocols (for example, a namespace that supports NFSv2 and NFSv3 cannot import a filer that supports only NFSv3). The show exports command shows the supported protocol(s) for each export/share, and the show global-config namespace command shows the protocol(s) supported by a namespace. Use show external-filer to list all filers.

Never import more than one instance of a given back-end share. If the filer uses aliases for its shares, the ARX assumes that each alias is a different share.

In a direct volume, you can use a managed volume as a filer instead. Use the managed-volume command to do this.

Guidelines: Importing a NetApp Share with Multiple Qtrees

If you choose a NetApp share with multiple qtrees, all the qtrees must have the same security style. Also, the qtrees must all use either NTFS or Unix style; managed volumes do not support qtrees with a "Mixed" security style.

A volume with such a NetApp share cannot support a shadow-copy-rule, the shadow command, or the migrate retain-files option in a place-rule.

Guidelines: Removing the Filer Binding

In a managed volume, you can remove the filer binding with the simple no filer command before the share is first imported. The share import is triggered by a combination of enable (gbl-ns-vol-shr) for the share and enable (gbl-ns, gbl-ns-vol) for either the volume or the namespace. A direct volume does not import its shares, so you can always detach its filers with the simple no filer syntax. The same applies to replica-snap shares in a managed volume.

After the managed volume imports the share, you can use a place-rule to remove all files from the share. Then you use the relocate-dirs argument in no filer to relocate the share's master directories to some other share in the same volume. (A volume often has multiple copies of its directories in each share, so that it can migrate files between them; a *master directory* is the copy where the volume puts all new files.) The volume scans the current back-end share for the file attributes of these directories, to be duplicated at the *target-share*. If this is impossible because the back-end share is offline, use the offline flag (along with remove-file-entries) to create new instances of the directories with zeroed-out file attributes.

The CLI creates a "removeShare" report to show the progress of the no filer command. Each report is named "removeShare.*share-name*.rpt," where *share-name* is the ARX-share name (not the name of the share at the filer).

The no filer command fails if any of the client-visible files in the managed volume are on the share; the remove-file-entries flag removes the files from the volume metadata and allows the no filer operation to succeed. The number of files removed appears in the removeShare report. The additional verbose flag adds the names of each removed file into the report. The additional offline flag, described above, applies to unreachable filers.

In a direct volume, you can use no filer, without any options, before or after you enable the share. This removes all of the share's files and directories from client view.

To remove the share configuration, you can use no share.

As an alternative to no filer or no share in a managed volume, you can use remove-share migrate to remove an imported managed-volume share, remove-share nomigrate to remove a managed-volume share that failed to import, or remove-share offline to remove a share that is unreachable.

Samples

```
bstnA(gbl-ns-vol-shr[wwmed~/acct~bills2])# filer das3 nfs /data/acct2
```

binds the "bills2" share to an NFS filer share.

```
bstnA(gbl-ns-vol-shr[mpns~/vol~stor4])# filer nas6 nfs /vol/vol1 cifs
VOL1
```

binds the "stor4" share to a multi-protocol share.

```
bstnA(gbl-ns-vol-shr[ns1~/vol~test])# no filer relocate-dirs share4
```

removes the filer binding for "test," migrating all of its master directories to "share4" in the "ns1~/vol" volume.

Related Commands

namespace -> volume -> share
show external-filer
show exports
show global-config namespace
no share
managed-volume

filer-subshares

Purpose A managed volume that supports CIFS can optionally support subshares and their share-level ACLs. A *subshare* is a CIFS share inside an imported CIFS share. Through a volume that supports subshares, a properly-configured cifs service can pass its clients from a front-end subshare to the corresponding back-end subshare. The back-end filer can then apply the subshare's ACL to the client's actions. Use the filer-subshares command to support subshares at this volume.

The no form of this command removes subshare support from this managed volume. A volume without subshare support always connects to the root of the back-end share, thereby using the ACL defined there, whether or not the client connects to a front-end subshare.

Mode gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
filer-subshares [native-names-only]
no filer-subshares
```

native-names-only (optional) causes a subshare to be "degraded" if it encounters any name collisions. A *subshare name collision* occurs when the volume attempts to replicate the subshare on a filer that already supports the same share name for a different directory. For example, suppose filer A has these two shares:

- e:\exports as "ESHARE" and
- e:\exports\dirA\books as "BOOKS." This is a subshare of "ESHARE."

Further suppose that filer B has these shares:

- d:\exports as "MYSTUFF" and
- c:\Users\juser\Documents as "BOOKS." This is *not* a subshare of "MYSTUFF."

If a managed volume imports both "ESHARE" and "MYSTUFF," it finds the "BOOKS" subshare under "ESHARE." It must replicate "BOOKS" on filer B at the same relative path, d:\exports\dirA\books. However, "BOOKS" is already defined on filer B for a directory on the C drive. The default solution to this issue is to generate a special subshare name, such as "_acopia_BOOKS_2$." Use the **native-names-only** option to forbid any such generated subshare names. The volume generates a full report on subshare replication (described below) during import; if this replication fails due to a subshare-name collision like this, the report shows the subshare name(s) that collided.

Default(s) `no filer-subshares`

Guidelines

This command only applies to a managed volume that supports CIFS.

The ARX is a proxy between front-end clients and back-end shares; it passes clients through to the back end for authentication. Share-level ACLs are created and enforced at the back-end filers. You can use show exports ... paths to probe for all shares on a given filer as well as the directory path behind each share. This shows you all of the share-to-subshare relationships on the filer.

Subshares require special configuration and processing on the ARX. After you use the filer-subshares command on a volume, you can use export (gbl-cifs) ... filer-subshare to export each of the subshares to your CIFS clients. When a client connects to one of these subshares, the volume passes the client connection directly to an equivalent back-end-filer subshare. The filer then uses its subshare ACL to determine the client's access privileges. Without the subshare configuration, the volume connects the client to the root of the share, and the filer uses the root share's ACL.

Guidelines: Subshare Replication

When you enable (gbl-ns, gbl-ns-vol) a volume with filer-subshares, the volume probes all its filer shares for subshares, and then it replicates the subshares and ACLs between those shares. The volume must ensure that all instances of a subshare have identical ACLs, so that a client has the same access privileges no matter which back-end subshare holds the desired file. For the replication to succeed, the volume must be set to modify and each of its shares must have the import sync-attributes flag raised.

The goal of subshare replication is to make all subshare definitions match as follows:

- share name,
- directory path, relative to the root of the share, and
- ACL.

The volume requires admin-level permissions to read or replicate the directory paths and subshares on the back-end filers. This is a higher level of access than you typically need for proxy-user credentials. You may need to change to more-privileged credentials behind the namespace's proxy-user, or configure the namespace with a new proxy-user (gbl-ns) altogether.

If synchronization is needed for a subshare but impossible due to a configuration issue (no modify, no import sync-attributes, and/or insufficient proxy-user permissions), the corresponding front-end subshare will be considered "degraded." Clients cannot connect to a degraded subshare.

For a replica-snap share, which is typically read-only, the subshare replication process cannot create directories on the share. Therefore, it cannot occur when the replica-snap share is enabled. Instead, the process waits for the snapshot replica-snap-rule to create the subshare directories, and then it defines the share name and ACL for each of them.

Guidelines: Subshare Replication with Free-Space Quotas

The volume-level freespace cifs-quota command makes a CIFS volume advertise free space in terms of each client's back-end storage quota. That is, if three back-end subshares each have a quota of 1G, clients with that quota see up to 3G of free space (1G for each of the three shares) instead of the actual free space. You set these free-space quotas at the back-end filers. The quotas are not replicated with other subshare data; if you use storage quotas on your back-end filers, replicate the storage quotas manually to all the back-end filers behind this managed volume.

For NetApp filers and EMC servers, manually replicate the subshares as qtrees and/or EMC tree quotas *before* you enable the volume. Do this at every NetApp and EMC share behind this managed volume. Use the same directory name, relative path, and share name for all of them. No special directories are required to set quotas on Windows servers, so you can add quotas after subshare replication on those devices.

Guidelines: Subshare-Name Collisions

Replicated back-end subshares use the same name as the source subshare, called the *native* subshare name, whenever possible.

The volume cannot use a native subshare name if any other CIFS share on the back-end filer is already using that name. This is called a subshare-name *collision*. For subshare-name collisions, the managed volume names the subshare with the following syntax:

_acopia_*subshare*_*id*$

where

- *subshare* is the native name of the original subshare.
- *id* is a unique integer to identify the subshare. This is unique for every subshare behind the ARX and its redundant peer.
- $, at the end of the share name, hides the subshare name from most network clients.

You can use the **native-names-only** option to prevent any such generated names. This means that a subshare-name collision puts the front-end subshare in a "degraded" state. CIFS clients cannot connect to a degraded front-end subshare.

Guidelines: Finding and Correcting Subshare Collisions

A managed volume with filer subshares generates a "sync subshares" report when it and its shares are enabled. This report describes the results of the subshare-replication process. The report name has the prefix, "syncSshrNewStorageReport." Use show reports *report-name* to see its contents. A sample report appears below; see Figure 22.3 on page 22-38. Inconsistent subshares are flagged in this report, if any are found; you can access the filer directly to change the subshare definition or ACL, then use sync subshares from-namespace to retry the subshare-replication process.

Guidelines: Subshare-ACL Collisions

When two import shares have matching subshares with different share-level ACLs, they collide. For example, two shares could contain directories named "myDir\yourDir" at their roots, both shared as YOURDIR but with slightly different ACLs. The ACLs must match at all instances of the subshare, so the volume uses the directory that was chosen as "master" to choose the ACL. The *master directory* is the one that is first imported. Before you import the shares, you can use the import priority command to choose a particular share (typically the Tier-1 share) to win directory mastership wherever there is such a collision. The volume copies the ACL from the master-directory subshare to all of the matching subshares.

An ACL collision does not prevent import, or require further action.

Guidelines: Windows Server Clusters

A Windows *Server Cluster* is a redundant configuration of Windows servers, called *nodes*. Subshare replication has a limitation when used on a Windows-server cluster: it can only replicate the subshares on the cluster's currently-active node. The replicated subshares and their ACLs only exist on the active node, and cannot be used after a cluster failover.

If you activate filer subshares in a volume backed by a server cluster, manually duplicate the subshares and ACLs on all nodes in the cluster. Subshare replication occurs when you enable (gbl-ns-vol-shr) one of the cluster's shares in the current ARX volume. After the replication is complete, duplicate the subshares and ACLs at the cluster's administrative interface; this ensures that they are on all nodes. Use the same subshare names created by the replication process, such as "Y2005" (a native subshare name) or "_acopia_Y2004_4$" (a generated subshare name).

Guidelines: Exporting the Subshare(s)

To find and export all the subshares from the volume to your front-end CIFS service, use sync subshares from-namespace.

While any of a volume's subshares are shared through a front-end CIFS service, you cannot use no filer-subshares in the volume.

Guidelines: Adding New Subshares

To add a new subshare in a running volume that already supports them:

- Choose an existing directory to be exported as a subshare. If none exists, create one now:
  - If you want to support CIFS quotas (freespace cifs-quota) and the volume is backed by a NetApp filer or EMC server, go directly to the back-end filer(s) and manually create qtree(s) or quota tree(s) for the subshare. Use the same directory name and relative path for all of them. Then run sync directories to incorporate the new directory into the volume metadata.
  - If you are not supporting quotas from NetApp or EMC, connect to the front-end export (gbl-cifs) as a client and create its directory.
- Use find to get the master directory for the new subshare path. This tells you which filer share contains the master directory.
- Access that filer share directly and create the share-level ACL, share definition, and/or directory attributes for the subshare path.
- Use the sync subshares from-namespace command to export the directory as a subshare, and to replicate the subshare on all of the volume's filers that do not already have it. This also copies the ACL and attributes from the master directory to all stripe directories in the volume.
- If the volume is backed by any Windows Servers and it supports freespace cifs-quota, access the Server(s) directly and set a path-based quota on the subshare(s).
- If the volume is backed by any Windows Server Clusters, manually duplicate the subshares and ACLs from the cluster's active node to the cluster's remaining node(s).

You cannot enable filer-subshares in a volume with already-enabled shares.

Samples

`bstnA(gbl-ns-vol[medarcv~/rcrds])# filer-subshares native-names-only`

configures the "/rcrds" volume to support CIFS subshares and share-level ACLs. The native-names-only option tells the volume not to generate subshare names (such as "_acopia_mysubshare_1") if subshare names collide; instead, let the replication fail and put the front-end subshare into a "degraded" state. This command raises a flag only, to be enforced when you enable the volume and its shares.

`bstnA(gbl-ns-vol[ns3~/vol1])# no filer-subshares`

removes CIFS-subshare support from the "/vol1" volume.

Related Commands

namespace **->** volume
import priority
show exports **... paths**
export (gbl-cifs) **... filer-subshare**
sync subshares from-namespace
sync subshares from-service
freespace cifs-quota

Figure 22.3 *Sample Report: syncSshrNewStorageReport_...*

```
bstnA# show reports syncSshrNewStorageReport_201007060501.rpt
**** Synchronize New Storage Report: Started at Tue Jul  6 01:01:27 2010 ****
**** Software Version: 5.02.000.12591 (Jul  3 2010 13:41:14) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

Operation Parameters
--------------------------------------------------------------------------------
Service:
Namespace: medarcv
Volume: /rcrds (modifiable, no access-based-enum)
Options:
    None.


ABE - access-based-enumeration.
ACL - security descriptor access control list.

Operation Execution Detail
--------------------------------------------------------------------------------

Action: Missing share added on filer.
    Share Name: CELEBS$
    Relative Path: VIP_wing
    Import Share: histories
    Filer Name: fs1
    Filer IP: 192.168.25.20

Action: Missing share added on filer.
    Share Name: CELEBS$
    Relative Path: VIP_wing
    Import Share: bulkstorage
    Filer Name: fs2
    Filer IP: 192.168.25.27

Action: Missing share added on filer.
```

```
    Share Name: MP3S
    Relative Path: 2011\mp3downloads
    Import Share: prescriptions
    Filer Name: fs4
    Filer IP: 192.168.25.29

Action: Missing share added on filer.
    Share Name: MP3S
    Relative Path: 2011\mp3downloads
    Import Share: histories
    Filer Name: fs1
    Filer IP: 192.168.25.20

Action: Missing share added on filer.
    Share Name: Y2004
    Relative Path: 2004
    Import Share: histories
    Filer Name: fs1
    Filer IP: 192.168.25.20

Action: Missing share added on filer.
    Share Name: Y2004
    Relative Path: 2004
    Import Share: bulkstorage
    Filer Name: fs2
    Filer IP: 192.168.25.27

Action: Missing share added on filer.
    Share Name: Y2005
    Relative Path: 2005
    Import Share: histories
    Filer Name: fs1
    Filer IP: 192.168.25.20

Action: Missing share added on filer.
    Share Name: Y2005
    Relative Path: 2005
    Import Share: bulkstorage
    Filer Name: fs2
    Filer IP: 192.168.25.27

Action: Missing share added on filer.
    Share Name: Y2010
    Relative Path: 2010
    Import Share: prescriptions
    Filer Name: fs4
    Filer IP: 192.168.25.29

Action: Missing share added on filer.
    Share Name: Y2010
    Relative Path: 2010
    Import Share: bulkstorage
    Filer Name: fs2
    Filer IP: 192.168.25.27



Summary:
--------------------------------------------------------------------------------
10 shares were added.

**** Total processed:              10
```

```
**** Elapsed time:           00:00:00
**** Sync New Storage Report: DONE at Tue Jul  6 01:01:27 2010 ****
```

freespace adjust

Purpose

If you used freespace calculate manual to manually calculate free space for this volume, you can use the freespace adjust command to adjust the free space that is advertised for the current share.

Use no freespace adjust to remove any free-space adjustment.

Modes

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

`freespace adjust [-]`*`adjustment`*`[k|M|G|T]`
`no freespace adjust`

- (optional) makes the adjustment negative.

adjustment is the change in advertised free space.

k|M|G|T chooses the unit of measure: **k**ilobytes, **M**egabytes, **G**igabytes, or **T**erabytes. A kilobyte is 1,024 bytes, a megabyte is 1,024 kilobytes (1,048,576 bytes), and so on.

Default(s)

`no freespace adjust`

k|M|G|T - if none of these is used, the adjustment is measured in bytes.

Guidelines

By default, the volume calculates its free space automatically: it takes the sum of all free space on all shares, but it only counts one of the shares from a given back-end storage volume. Use the freespace calculation manual command to allow manual adjustments in the current volume, then use this command to adjust the free space for each chosen share.

You can also use freespace ignore to ignore the share in the free-space calculation. If you use this command together with freespace ignore on a given share, this command sets the total free space advertised by the share.

Some back-end CIFS filers (such as Windows Servers) can support free-space quotas, and these can have an effect on the free-space adjustment. You can use the freespace cifs-quota command to discover the free space available to a client at each back-end share, add up the client's available free space from all those shares, and pass the sum back to the CIFS client. This command, freespace adjust, then adjusts the current share's numbers based on the client's available space, not based on the total size of the share.

Use the show global-config namespace command to see if this command is set for the share.

CIFS clients only see the results of this command if they connect *after* you invoke it.

Samples

```
bstnA(gbl-ns-vol-shr[ns1~/vol~shr4])# freespace adjust -1G
```

changes the free space that is counted for shr4; this advertises 1 Gigabyte less than the actual free space.

```
bstnA(gbl-ns-vol-shr[wwmed~/acct~bills])# no freespace adjust
```

uses the actual free space for "bills" share.

Related Commands

namespace `->` volume `->` share
freespace calculation manual
freespace ignore
freespace cifs-quota
show global-config **`namespace`**

freespace apparent-size

Purpose

Some sites restrict each of their clients to a single back-end share, using the ARX for occasional migrations only (that is, to migrate their entire directory tree to another share all at once). This creates a situation where the client can only use the free space on their current back-end share, not on all the shares behind the managed volume. You can use freespace calculate dir-master-only for this situation, which shows clients only the free space on their currently-assigned back-end share. If you use that command, you can also use freespace apparent-size to change the free space that is advertised for the current share.

Use no freespace apparent-size to report the full size of the current share to the volume's clients.

Modes

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
freespace apparent-size size[k|M|G|T]
no freespace apparent-size
```

size is the full size to advertise for this share. After the share is imported, you can use the show namespace command to find its actual space.

k|M|G|T chooses the unit of measure: **k**ilobytes, **M**egabytes, **G**igabytes, or **T**erabytes. A kilobyte is 1,024 bytes, a megabyte is 1,024 kilobytes (1,048,576 bytes), and so on.

Default(s)

```
no freespace apparent-size
```

k|M|G|T - if none of these is used, the adjustment is measured in bytes.

Guidelines

By default, the volume calculates its free space automatically: it takes the sum of all free space on all shares. Use the freespace calculation dir-master-only command to advertise only the free space on a single back-end share to any client. The single back-end share is the one with the *master* for the client's front-end-share path; all back-end *stripe* directories are excluded. Then you can optionally use this command to set the advertised free space for each share.

If the back-end share uses up the free space you advertise with this command, but is not actually full, clients see 0 (zero) free space.

Use the show global-config namespace command to see if this command is set for the share.

Samples

```
bstnA(gbl-ns-vol-shr[insur~/claims~shr1-next])# freespace apparent-size
2460M
```

changes the free space that is advertised for shr1-next; this advertises free space as though the share's total size is 2460 Megabytes (2.4 Gigabytes).

```
stoweA(gbl-ns-vol-shr[lodges~/skiPatrol~stoweMtn])# no freespace
apparent-size
```

uses the actual free space size for "stoweMtn" share.

Related Commands namespace `->` volume `->` share
freespace calculation dir-master-only
show global-config **`namespace`**

freespace calculation dir-master-only

Purpose

The volume's free space, as seen by clients, is the sum of all free space in its back-end shares. This is misleading in installations where clients only access a single back-end share behind the volume; such clients are restricted to that single share, and should therefore only see the free space from that share. The single back-end share is said to hold the front-end share's *master* instance; replica shares on other back-end filers are called *stripes*. To show clients the free space in only the master back-end share, use the freespace calculation dir-master-only command.

Use the no form of the command to show clients the sum of all free space on all back-end shares behind this volume.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
freespace calculation dir-master-only
no freespace calculation dir-master-only
```

Default(s)

```
no freespace calculation dir-master-only
```

Guidelines

This command is useful in a managed volume with home directories, or a managed volume that is migrating files once (from one share to another). This only affects the client view of the volume's free space. It has no effect on the free space perceived by the policy engine.

A direct volume does not support this type of free-space calculation.

The master share is the back-end share that holds the master instance of the front-end share's directory. This is the root of the export (gbl-cifs) to which the client connected. This could be the root of the entire volume or it could be a CIFS subshare.

You may want to advertise a fixed size for each share, perhaps to ensure that the free-space seen by a client is consistent if his or her directory tree migrates from one share to another. For each share, you can set an "apparent" space size with the freespace apparent-size command.

Use the show global-config namespace command to see the volume's configuration settings for free space.

Guidelines: Alternative Free-Space Reporting

The freespace calculation manual command is an alternative form of free-space calculation where you can manually adjust the free-space numbers from each share. If you set the volume for manual free-space calculation, you can exclude shares with the freespace ignore command. You can also adjust the free-space number for a share (whether or not its actual free space is ignored) with the freespace adjust command.

Samples

```
bstnA(gbl-ns-vol[insur~/claims])# freespace calculation dir-master-only
```

changes the /claims volume so that it only shows the free space from one back-end share rather than all of them.

```
bstnA(gbl-ns-vol[wwmed~/acct])# no freespace calculation
dir-master-only
```

returns the /acct volume to the default. This volume reports the total free space to its clients, summing up the free space in all of its back-end shares.

Related Commands namespace -> volume
freespace apparent-size

freespace calculation manual

Purpose

The volume's free space, as seen by clients, is the sum of all free space in its shares. By default, the volume finds the back-end storage volume for each of its shares and, if two or more of its shares draw from the same back-end storage volume, it counts the free space from only one of them. Use the freespace calculation manual command to remove this automation, permitting you to manually decide which shares to ignore.

Use the no form of the command to automatically detect multiple shares from the same back-end storage volume.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
freespace calculation manual
no freespace calculation manual
```

Default(s)

```
no freespace calculation manual
```

Guidelines

An NFS volume uses a share's file-system ID (FSID) to identify the storage volume behind the share. A CIFS volume uses the share's Volume Serial Number. By default, only one share is counted if two or more shares have matching IDs.

If you set the volume for manual free-space calculation, you can exclude shares with the freespace ignore command. You can also adjust the free-space number for a share (whether or not its actual free space is ignored) with the freespace adjust command.

Some back-end CIFS filers (such as Windows Servers) can support free-space quotas, and these can affect the free-space calculation. You can use the freespace cifs-quota command to discover the free space available to a client at each back-end share, add up the client's available free space from all those shares, and pass the sum back to the CIFS client. The freespace adjust command then adjusts a share's free-space numbers based on the client's quota, not based on the total size of the share.

Use the show global-config namespace command to see the volume's configuration settings for free space.

Guidelines: Alternative Free-Space Reporting

The freespace calculation dir-master-only command is an alternative form of free-space calculation where the volume only shows the free space in one back-end share behind the volume. It chooses the share that holds the master instance of the client's current directory. This is useful for volumes where a client's directory tree is exclusively on one back-end share at any given time. If you set the volume for this free-space calculation, you can also set the client-visible space on each of the volume's shares with the freespace apparent-size command. You cannot use freespace adjust or freespace ignore with the freespace calculation dir-master-only command.

Samples

```
bstnA(gbl-ns-vol[ns1~/])# freespace calculation manual
```

includes all shares, allowing you to manually ignore some shares.

```
bstnA(gbl-ns-vol[wwmed~/acct])# no freespace calculation manual
```

reverts to automatic free-space calculations. The automatic calculation notes the back-end shares that draw from the same storage pool, and each back-end storage pool only once.

Related Commands

namespace -> volume
freespace ignore
freespace adjust
freespace cifs-quota

freespace cifs-quota

Purpose

A managed volume's free space, as seen by clients, is the sum of all free space in its back-end shares. By default, this does not take into account any storage quotas on any back-end CIFS servers; if a 100G back-end share has a quota of 1G, the client sees 100G of free space when the share is empty. You can use the freespace cifs-quota command to take the back-end quotas into account, so that CIFS clients with a 1G quota only see 1G of space and connections to a CIFS subshare with a 5G quota see only 5G of space. With this option enabled, clients see the sum of the quota-based free space on the volume's back-end shares, not the sum of the full free space.

Use the no form of the command to show all clients the full free space, ignoring all back-end quotas.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
freespace cifs-quota
no freespace cifs-quota
```

Default(s)

```
no freespace cifs-quota
```

Guidelines

This command is only useful in a managed volume that supports CIFS. It is not supported in a direct volume, or one that supports only NFS (see protocol).

Back-end filers can support one or two forms of storage quota:

- *path-based* quotas, which affect a particular directory path.
- *user-based* quotas, which allow each client to use only a limited amount of storage in a particular storage pool.

This command, freespace cifs-quota, makes both of these quota types visible to the ARX volume's clients. If a given client is allotted a total of 4G on the volume's back-end filers, due to path-based quotas and/or the client's user-based quotas at those filers, the client only sees 4G of space.

Use the show global-config namespace command to see the volume's configuration settings for free space.

CIFS clients only see the results of this command if they connect *after* you invoke it.

Guidelines: Relationship of Front-End Shares to Back-End Quotas

Path-based quotas are applied to back-end paths, so the free space that the client sees depends on the front-end CIFS share to which they connect. There are two major forms of front-end CIFS share:

- An export (gbl-cifs) is a full share of the entire volume, and
- An export (gbl-cifs) ... filer-subshare is a subshare.

If a CIFS client connects to a full share (or any path within a full share), the client sees the free-space quotas set at the imported path (or root) of the volume's back-end shares. If the client connects to a subshare, the client sees the free-space quotas set at its back-end subshare paths.

When each subshare is dedicated to a single client (as with a home-directory application), you can use path-based quotas at each back-end subshare. Then the freespace cifs-quota command causes the ARX volume to advertise those quotas to its clients.

◆ Note

In all cases, we recommend setting specific path-based quotas at each back-end share or subshare to which your CIFS clients connect. For NetApp, this means that each share or subshare must be its own qtree. For EMC, the share or subshares must reside in an EMC File System with quotas set at the import share or at each subshare. For Windows, set a path-based quota at the import share or at each subshare.

Guidelines: NetApp and EMC with Filer Subshares

Free-space quotas (this command) and filer-subshares are difficult to use together in a volume backed by either NetApp or EMC. A NetApp only supports path-based quotas on its qtrees, and an EMC server only supports them in a quota tree inside a File System. Each qtree or quota tree gets a single quota. As mentioned above, each subshare requires a path-based quota for proper free space reporting. Therefore, each subshare path must be backed by one qtree or quota tree. This is required for CIFS clients to see the space that they are allowed to use.

The ARX volume automatically replicates its subshare directories to all of its back-end filers. The replication process cannot safely presume that you want a full qtree or quota tree behind these replicated directories, so it can only create the directory itself. Therefore, the subshare-replication process creates directories on NetApp and EMC that cannot support path-based quotas.

For NetApp or EMC, pre-create the qtrees and/or quota trees instead of relying on the ARX volume's subshare-replication process. Do this before you import any storage from the back-end filer(s), as described in the *Guidelines: Subshare Replication with Free-Space Quotas* section of the filer-subshares documentation.

If subshares are already configured as plain directories on your back-end NetApp or EMC, and you want to convert them to qtrees or EMC quota trees, contact F5 for guidance.

Guidelines: Windows Servers with Filer Subshares

You must also set path-based quotas on all Windows servers where subshare-replication occurs. The subshare-replication process does not replicate Windows-quota settings along with other subshare settings. Unlike NetApp and EMC, you can set the quotas on Windows servers after subshare replication is complete; you do not need to pre-create any special directories.

Guidelines: Manually Adjusting the Reported Free Space

You can also manually ignore the free space on some of the volume's back-end shares, or adjust the advertised free space for one or more shares. Use the freespace calculation manual command in the volume, then use freespace ignore and/or freespace adjust in each share. The adjustment from freespace adjust is applied to the back-end quota, not to the total size of the back-end share.

Guidelines: Back-End Shares That Draw from the Same Storage Pool

An ARX volume may import two or more shares that are affected by the same free-space number, and this creates errors in free-space reporting. Whenever the freespace cifs-quota feature is enabled, the ARX volume presumes that a path-based quota is set on all of its back-end shares; it therefore counts the space from each share separately. For example, suppose shares HUEY, DEWEY, and LOUIE all draw from drive E on the same Windows server, and an ARX volume imports all three of them. The volume also adds together all of their free space. This creates an incorrect sum if

- there is no quota (all three shares report the full space of the E drive),
- there is a user-based quota (all shares report the full space allowed to the client), or
- all three shares inherit the same path-based quota from a common ancestor, such as E:\.

Specific path-based quotas for each share, recommended earlier, helps to solve this problem. Additionally, avoid any user-based quotas or common path-based quotas altogether.

If this is not possible, use freespace ignore to ignore each imported share from a common storage pool except one. For example, ignore the freespace on any two of the three shares (such as HUEY and DEWEY).

Guidelines: Alternative Free-Space Reporting

The freespace calculation dir-master-only command is an alternative form of free-space calculation where the volume only shows the free space in one back-end share behind the volume. It chooses the master back-end share for the front-end share to which the client connected. This is useful for volumes where a client's directory tree is exclusively on one back-end share at any given time. If you use this free-space calculation technique, you can also set the client-visible space on each of the volume's shares with the freespace apparent-size command.

Samples

```
stoweA(gbl-ns-vol[lodges~/skiPatrol])# freespace cifs-quota
```

changes the /skiPatrol volume so that it detects path-based and/or user-based quotas on its back-end shares, and reports only the free space that is allotted to each client.

```
provA(gbl-ns-vol[provMed~/rns])# no freespace cifs-quota
```

returns the /rns volume to the default. This volume reports the total free space to its clients, without requesting quota-based free space from its back-end filers. If the back-end filers do not support any CIFS quotas, this setting reduces traffic between the managed volume and its filers.

Related Commands

namespace -> volume
filer-subshares
freespace calculation manual
freespace ignore
freespace adjust

freespace ignore

Purpose

If you used freespace calculate manual to manually calculate free space for this volume, you can use the freespace ignore command to exclude the current share from the free-space calculation.

Use no freespace ignore to include the free space from the current share.

Modes

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
freespace ignore
no freespace ignore
```

Default(s)

```
no freespace ignore
```

Guidelines

By default, the volume calculates its free space automatically: it takes the sum of all free space on all shares, and it excludes all but one share from any given back-end storage volume. Use the freespace calculation manual command to allow manual exclusions in the current volume, then use this command to ignore each chosen share.

You can also use freespace adjust to adjust the amount of free-space that is counted for this share. This can be an alternative to ignoring the share's free space (for example, adjust the space by -100G instead of ignoring it altogether), or you can use it to set your own free-space value for the share (ignore the share's actual space, then adjust the share's free space to +50G).

Use the show global-config namespace command to see if this command is set for the share.

Samples

```
bstnA(gbl-ns-vol-shr[ns1~/vol~shr4])# freespace ignore
```

ignores the free space from shr4.

```
bstnA(gbl-ns-vol-shr[wwmed~/acct~bills])# no freespace ignore
```

includes the free space from the "bills" share.

Related Commands

namespace -> volume -> share
freespace calculation manual
freespace adjust
show global-config namespace

ignore-sid-errors

Purpose In some cases, a back-end-CIFS filer does not recognize a Security ID (SID) on a migrated file. CIFS filers typically return an error and reject the file with the unknown SID; the volume therefore assumes that the migration failed and keeps the file on its original share. Some file servers, however, can be configured to return an error and *accept* the file anyway. EMC Celerra servers have shown this behavior in lab testing. You can use the ignore-sid-errors command to ignore these errors from such a file server.

To react to SID errors by canceling the file migration, use no ignore-sid-errors.

Mode gbl-ns-vol-shr

Security Role(s) storage-engineer or crypto-officer

Syntax

```
ignore-sid-errors
no ignore-sid-errors
```

Default(s)

```
no ignore-sid-errors
```

Guidelines This command does not apply to an NFS-only namespace.

SID errors occur for a SID that is unknown at the destination file server. The SID may be unknown because it came from another file server with Local Groups or Local Users. You can use sid-translation to translate all local SIDs (after some preparation at the filer servers). However, this translation could fail if the filer servers are not fully prepared to support Local Groups.

The command causes the share to ignore certain SID errors from its back-end file server and assume that the migration to the file server succeeded. The SID errors, in response to the CIFS 'set security descriptor' request, are STATUS_INVALID_SID, STATUS_INVALID_OWNER, and STATUS_INVALID_PRIMARY_GROUP. When most filers return these errors, they also reject the file. Use this only for file servers that accept the file in spite of the error.

Important

If you use the ignore-sid-errors *command, be sure that the filer is configured to accept the file (or directory) itself in spite of any SID errors. If the filer returns a SID error for a file and then drops it, the file is lost if the ARX share is set to* ignore-sid-errors.

Samples

```
bstnA(gbl-ns-vol-shr[medarcv~/rcrds~bulk])# ignore-sid-errors
```

causes the volume to ignore any SID errors from the file server behind the "bulk" share.

```
bstnA(gbl-ns-vol-shr[insur~/claims~shr1-old])# no ignore-sid-errors
```

causes the "shr1-old" share to acknowledge any SID errors from its filer. If the policy engine gets a SID error when it tries to migrate a file or directory, the policy engine cancels the migration.

Related Commands namespace -> volume -> share
sid-translation

import priority

Purpose

A managed volume imports files and directories as its clients randomly access them, making file and directory mastership also random. For example, suppose a client accesses "/bigDir/latest.txt" from a volume with two shares, A and B. Both shares contain a "/bigDir" directory and share B contains the "/bigDir/latest.txt" file. This causes the volume to import the share B instance of "/bigDir" first, so that share B wins the directory conflict. The directory on share B becomes the *master directory*. The master directory keeps its name and attributes, but the *stripe directory* in share A may need to have its attributes or name changed for a successful import. You can use the import priority command to set a priority for the current share, so that you determine which shares win these import conflicts. If two shares have different priorities, the higher-priority share wins the conflict.

The volume imports its shares in their priority order; the volume waits for all shares of priority A to complete their imports before it begins importing the share(s) of priority B.

Use the no import priority command to reduce the share to the lowest priority.

Mode

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
import priority number
no import priority
```

number (1-65535 characters) is the priority for the current share. The highest priority is 1 (one); a share with a priority of 1 wins any import conflict with any lower-priority share.

Default(s)

`65535`

Guidelines

This command is strongly recommended for any volume with tiered shares, where one share or set of shares is considered faster and/or more-reliable than the rest. The master instances of all directories should reside on the volume's reliable Tier 1 shares, so that clients can reliably access them. An import priority of 1 on the volume's Tier 1 shares ensures that, wherever there is a directory collision, the Tier 1 share gets the master directory.

If a single share has the highest priority in the volume, it wins every import conflict and does not require any other import commands (such as import rename-directories, import rename-files, or import sync-attributes). For any given collision, those commands are only relevant for the lower-priority share.

Tiering is implemented with two or more file-placement rules (place-rule) that move higher-priority files to the Tier 1 shares and lower-priority files to other shares. Best practices dictate that you also create a place-rule that puts *all* the volume's master directories onto its Tier 1 shares, not just the directories that collided on import. To do this, create a filename-fileset that recursively matches everything in the volume (path exact / and recurse), and use the from (gbl-ns-vol-plc) command to have that fileset match directories and promote them.

The volume imports any shares with the same priority in parallel.

Guidelines: Import Priority Only Applies to the Current Volume

If multiple volumes are importing their shares at the same time, the import priorities in one volume have no effect on the import priorities in the other volume. For example, if volume /usrs has one or more shares at priority 1 and volume /var has shares at priority 100 (the highest priority in the volume), the priority-100 shares in /var may import at the same time as the priority-1 shares in /usrs.

Guidelines: Share Roots

The first-configured share is assigned the master for the volume's root directory, and the import priority does not change this. To control directory mastership at the volume's root directory, use a place-rule that moves all master directories to your chosen share or share-farm, as described above. Such a place rule ensures that all master directories always reside on the share or share farm where you want them.

Samples

```
bstnA(gbl-ns-vol-shr[medarcv~/lab_equipment~equip])# import priority 1
```

sets the highest priority for the "equip" share. If any files or directories in this share collide with those of a lower-priority share, this share wins the conflict and gets the master instance of the file or directory.

```
bstnA(gbl-ns-vol-shr[ns1~/vol~shr4])# no import priority
```

resets the 'shr4' share to the lowest possible priority.

Related Commands

namespace -> volume -> share
place-rule
import rename-directories
import rename-files
import sync-attributes

import rename-directories

Purpose

There are some cases where it is necessary for a managed volume to rename a directory on import:

- It *collides* with an already-imported directory. That is, it has the same name as a directory from an already-imported share, but different file attributes.
- It collides with an already-imported *file* with the same name.

If the managed volume is allowed to modify files and directories during import, it renames these directories by default. Use the no import rename-directories command to prevent these directory renames in the current share, potentially causing this share's import to fail for directories that collide.

Use import rename-directories to permit the managed volume to rename directories in this share.

Mode

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no import rename-directories [unmapped-unicode]
import rename-directories
```

unmapped-unicode (optional) causes the import to rename directories with Unicode-only characters in their names. This occurs at multi-protocol (CIFS and NFS) sites where character-encoding nfs was set at a non-Unicode standard (such as UTF-8), but CIFS users named their directories with Unicode characters. This is off by default to protect against unexpected renames in a large multi-protocol volume.

Default(s)

```
import rename-directories
```

Guidelines

In a managed volume, two same-named directories from different shares are aggregated by default: the volume presents a single directory to its clients, with all the files from both directories. The directories collide, however, if their file attributes (such as permission settings and group ownership) are different, or if the directory's name is the same as an already-imported file. The gbl-ns-vol modify command configures the volume to rename collided directories. The no import rename-directories command disallows directory renames in this particular share, and instead causes one of two reactions to a directory collision:

- for a collision with a file: fail the import for this share.
- for a collision with a directory: if import sync-attributes is set, synchronize the attributes of this share's directory with the already-imported directory. Otherwise, fail the import for this share.

To change the rules for modifying files in this share, you can use the import rename-files command.

This option is only relevant in the second-imported share and any subsequently-imported shares. If two shares are imported together with different import priorities (see import priority), the share with the higher priority always imports first.

Guidelines: Unmapped Unicode

If a multi-protocol filer uses single-byte character encoding for its NFS names, it typically permits its CIFS clients to use multi-byte characters in their directory names. The multi-byte characters are not mappable from CIFS to the NFS-character encoding; the filer typically creates a filer-generated name (FGN) that NFS clients can use to access the directory. On some filers, the NFS FGN bears no apparent resemblance to the CIFS-side name.

A multi-protocol volume does not allow CIFS names with non-mappable characters, so this naming problem causes an import failure by default. If you set the unmapped-unicode option, the volume can rename any such directories on import.

Shares with this problem typically fail their initial import because of the default. You can use show reports import-report to see which directories have non-mappable characters, then decide to either manually rename the directories at the filer or allow the volume to do it on the next import. Then retry the import.

Guidelines: Renaming Format

For cases where directory renames are allowed (import rename-directories, the default), the managed volume renames directories using the following syntax: *dirName_shareName-importId*[*-index*][*.ext*].

- *dirName* is the base name of the directory, not including any extension it may have. If it has any characters that are non-mappable between CIFS and NFS, it contains "(U+*nnnn*)" in place of each non-mappable character. The *nnnn* is the Unicode number for the character, shown in hexadecimal format. The name is truncated if it exceeds 256 characters. For example, "dir(U+30D2)(U+30AA)_myshare-2."
- *shareName* identifies the namespace share.
- *importId* is an internal identifier for the import job.
- *-index* only appears if multiple files or directories map to this name. It is unique for each such file or directory (for example, "renamedDir_1," "renamedDir_1-1," "renamedDir_1-2," and so forth).
- *.ext* is the original extension for the directory, if there was one.

Guidelines: A "no modify" Import

You can run a mock import of each share to review all directory-naming collisions before actually importing the share. Before you enable a share and start its import, you can disable all volume modifications with no modify. This prevents the volume from changing anything in its shares, and also prevents client access to the volume. Then enable the share to invoke the "no modify" import (with enable (gbl-ns-vol-shr)). After the "no modify" import finishes, review the share's import report for naming collisions (as well as any other import issues) and manually correct each issue at the back-end filer. Use show reports type Imp to list all import reports, and show reports *report-name* to read one.

If there were no issues, you can run the modify command to activate the volume and the import is finished. Otherwise, remove the share from the volume (with remove-share nomigrate), then add the share back into the volume to retry the import.

Guidelines: Recovery from Import Failure

If an actual import fails because a directory rename was disallowed, you can use remove-share nomigrate to remove the share from the volume. After the share is removed, you can fix the collisions directly on the filer and then re-import the share into the volume.

Samples

```
bstnA(gbl-ns-vol-shr[wwmed~/acct~bills])# no import rename-directories
```

disables the renaming of directories in the 'bills' share during an import. If a directory in this share collides with a file in an already-imported share, this share import fails. If it collides with a directory, the share import only succeeds if the import sync-attributes command is set to allow the volume to synchronize the directory attributes. If it has any character that is unsupported by the character-encoding nfs setting, the import fails.

```
bstnA(gbl-ns-vol-shr[ns1~/vol~shr4])# import rename-directories
```

enables the renaming of directories in the 'shr4' share during an import.

```
bstnA(gbl-ns-vol-shr[insur~/claims~shr1-old])# import
rename-directories unmapped-unicode
```

permits the 'insur~/claims~shr1-old' share to rename all directories with non-mappable characters in their names.

Related Commands

enable (gbl-ns-vol-shr)
import priority
import rename-files
import sync-attributes
modify

import rename-files

Purpose

This command applies only to a managed volume. When a file from the current share has the same name as a file or directory from an already-imported share, the file is said to *collide*. If the managed volume is allowed to modify files and directories during import, it renames collided files by default. Use the no import rename-files command to prevent these file renames in the current share, causing this share's import to fail for files that collide.

Use import rename-files to permit the volume to rename files in this share.

Mode

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no import rename-files
import rename-files
```

Default(s)

```
import rename-files
```

Guidelines

If the managed volume is allowed to modify files that collide, it renames those files by default. You can use the no import rename-files command to change this default for the current share; this causes the current share to fail its import if any files collide.

If the import fails, you can use remove-share nomigrate to remove the share from the volume. After the share is removed, you can fix the collisions directly on the filer and then re-import the share into the volume.

You can also use the import rename-directories and import sync-attributes commands to set the rules for directory collisions in this share.

This option is only relevant in the second-imported share and any subsequently-imported shares. If two shares are imported together with different import priorities (see import priority), the share with the higher priority always imports first.

Guidelines: Renaming Format

For cases where file renames are allowed (import rename-files, the default), the managed volume renames directories using the following syntax: *fileName_shareName-importId*[-*index*][.*ext*].

- *fileName* is the base name of the file, not including any extension it may have.
- *shareName* identifies the namespace share.
- *importId* is an internal identifier for the import job.
- -*index* only appears if multiple files or directories map to this name. It is unique for each such file or directory (for example, "renamedFile_1.doc," "renamedFile_1-1.doc," "renamedFile_1-2.doc," and so forth).
- .*ext* is the original extension for the file, if there was one.

Guidelines: A "no modify" Import

You can run a mock import of each share to review all file-naming collisions before actually importing the share. Before you enable a share and start its import, you can disable all volume modifications with no modify. This prevents the volume from changing anything in its shares, and also prevents client access to the volume. Then enable the share to invoke the "no modify" import (with enable (gbl-ns-vol-shr)). After the "no modify" import finishes, review the share's import report for naming collisions (as well as any other import issues) and manually correct each issue at the back-end filer. Use show reports type lmp to list all import reports, and show reports *report-name* to read one.

If there were no issues, you can run the modify command to activate the volume and the import is finished. Otherwise, remove the share from the volume (with remove-share nomigrate), then add the share back into the volume to retry the import.

Samples

```
bstnA(gbl-ns-vol-shr[wwmed~/acct~bills])# no import rename-files
```

disables the renaming of files in the 'bills' share during an import. If a file in the 'bills' share collides with an already-imported file or directory, the volume does not import the 'bills' share.

```
bstnA(gbl-ns-vol-shr[ns1~/vol~shr4])# import rename-files
```

enables the renaming of files in the 'shr4' share during an import.

Related Commands

enable (gbl-ns-vol-shr)
import priority
import rename-directories
import sync-attributes
modify

import skip-managed-check

Purpose

A managed volume, by default, performs a strict check on each directory before importing it: the share import fails if there is any evidence that the directory is already imported by another managed volume. This is an important precaution; no two managed volumes should ever control the same back-end directory. However, you can skip the directory check to increase the speed of the import. Use import skip-managed-check to prevent the volume for checking this share's directories.

Use the no import skip-managed-check command to revert to the safer import option.

Mode

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
import skip-managed-check
no import skip-managed-check
```

Default(s)

```
no import skip-managed-check
```

Guidelines

You should only skip the managed-directory check if you are sure that the back-end share is not managed by another ARX volume. For example, you can skip this check for a new filer, the initial cut-in of the ARX, or before an nsck ... rebuild of an existing volume.

Samples

```
bstnA(gbl-ns-vol-shr[medarcv~/lab_equipment~equip])# import
skip-managed-check
```

skips the directory check for a future import of the 'equip' share. This makes imports faster.

```
bstnA(gbl-ns-vol-shr[ns1~/vol~shr4])# no import skip-managed-check
```

authorizes the slower directory checks for the 'shr4' share.

Related Commands

namespace -> volume -> share
enable (gbl-ns-vol-shr)
nsck ... rebuild

import sync-attributes

Purpose

This only applies to a managed volume. When a directory from the current share has the same name as a directory from an already-imported share, but different file attributes, the directory is said to *collide*. If the volume is allowed to modify files and directories during import, it renames collided directories by default. Use the import sync-attributes command to allow the managed volume to instead synchronize the directories' attributes, rather than renaming them.

Use the no form of the command to disallow attribute synchronization for directories.

Modes

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
import sync-attributes
no import sync-attributes
```

Default(s)

```
no import sync-attributes
```

Guidelines

In a managed volume, two same-named directories from different shares are aggregated by default: the volume presents a single directory to its clients, with all the files from both directories. The directories collide, however, if their file attributes (such as permission settings and group ownership) are different. The gbl-ns-vol modify command, along with the default for import rename-directories, configures the volume to rename collided directories. This command provides an alternative for the current share: synchronize the directory attributes instead of renaming.

For the root directory of each share, renaming is not an option. You can use this option to synchronize all root-share attributes automatically, or you can synchronize the attributes manually at the filers (before the import). As an alternative, you can disable all volume modifications with no modify, review each share's import report for attribute collisions (as well as any other import issues), manually correct these issues at the filer, then retry the import with no enable (gbl-ns-vol-shr) followed by enable. (Use show reports type Imp to list all import reports.)

If a directory collides with a file, attribute synchronization is not enough to resolve the conflict. The managed volume must rename the directory (see import rename-directories) or the share import fails.

The no form of the command causes one of two reactions to a directory collision:

- rename the directory (if import rename-directories is set for this share) or
- fail the import (if no import rename-directories is set).

◆ Note

For heterogeneous multi-protocol namespaces, always enable synchronization with the import sync-attributes *command. Both CIFS and NFS attributes are compared in multi-protocol namespaces, greatly increasing the likelihood of directory collisions.*

Guidelines (Cont.)

If the import fails, you can use remove-share nomigrate to remove the share from the managed volume. After the share is removed, you can fix the collisions directly on the filer and then re-import the share into the volume.

This option is particularly important for CIFS-subshare support, and is very helpful for Access-Based Enumeration (ABE). These synchronize directory-attribute settings, such as subshare ACLs and ABE settings, between the volume's filers. You can use the filer-subshares command to set up subshares for the current volume. Use the cifs access-based-enum command to set up ABE.

To change the rules for modifying files in this share, you can use the import rename-files command.

Samples

```
bstnA(gbl-ns-vol-shr[ns1~/vol~shr4])# import sync-attributes
```

allows the volume to synchronize directory attributes in the 'shr4' share.

```
bstnA(gbl-ns-vol-shr[wwmed~/acct~bills])# no import sync-attributes
```

prevents the '/acct' volume from synchronizing directory attributes in the 'bills' share.

Related Commands

import priority
import rename-directories
import rename-files

managed-volume

Purpose Use the managed-volume command to assign a local, managed volume as the "back-end filer" for the current share. This only applies to shares in direct volumes.

Use the no form of the command to remove the managed volume from the share.

Mode gbl-ns-vol-shr

Security Role(s) storage-engineer or crypto-officer

Syntax

```
managed-volume namespace volume [access-list list-name]
no managed-volume
```

name (1-30 characters) is the managed volume's namespace.

volume (1-1024 characters) identifies the managed volume.

list-name (optional, 1-64 characters) is the NFS-access list to associate with the share.

Default(s) None.

Guidelines A direct volume has the option to use a managed volume to stand in as a *virtual filer* for one of its shares. The managed-volume command is analogous to the filer command, which assigns an external-filer share to the current namespace share.

Before you use this command in an NFS volume, you must disable the managed-volume's auto reserve files feature. A direct volume requires that any NFS share has a static, unchanging limit on its number of files.

If the direct volume supports CIFS, the volume that you use as a virtual filer must be in the same namespace.

Sample

```
bstnA(gbl-ns-vol-shr[medco~/vol~sales])# managed-volume wwmed22
/04accts
```

assigns the 'wwmed22~/04accts' managed volume as the filer for the 'sales' share.

Related Commands

namespace -> volume -> share
direct
filer
no auto reserve files

max-volume-groups

Purpose

Each ARX volume shares its memory, CPU time, and other resources with the other volumes in its *volume group*. On the ARX-500, memory and CPU cycles are less plentiful; the maximum volume groups are therefore set at a lower value on that platform. Raising the number of volume groups could create a resource-contention issue on an ARX-500. On the advice of F5 Support, you can use the max-volume-groups command to raise the maximum to the upper limit.

Use the no form of the command to return to the default maximum for volume groups.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
max-volume-groups
no max-volume-groups
```

Default(s)

2 volume groups

Valid Platforms

ARX-500

Guidelines

This command doubles the maximum volume groups. The default is sufficient for most installations. Use this command only under the guidance of F5 personnel.

You can only change this maximum while all volumes are disabled. Use no enable in gbl-ns mode to disable all volumes in a namespace (see the documentation for enable (gbl-ns, gbl-ns-vol)). You cannot reduce the number to a point where current volume-group assignment would be impossible; for example, if a volume is already assigned to volume-group 3, you cannot use no max-volume-groups because it would reduce the maximum to 2 groups.

Each volume group uses the memory set by the metadata cache-size command. The ACM processors require at least that much memory for each volume group. The show processors command shows the memory resident to each processor on your system.

Use the volume-group command to assign a volume to a group. The show volume-group command shows all volume-to-group assignments.

Samples

```
canbyA(gbl)# max-volume-groups
```

allows the maximum possible volume groups on the current switch, an ARX-500. This should only be done on the advice of F5 personnel.

```
canbyA(gbl)# no max-volume-groups
```

resets the maximum number of volume groups back to the recommended default.

Related Commands

volume-group
show volume-group
metadata cache-size
show processors

metadata critical

Purpose

If this switch has a redundant peer, you can use the metadata critical command to declare the metadata for the current managed volume as a critical resource. If an active switch loses access to this volume's metadata share, a failover may occur.

Use the no form of the command to remove the "critical-resource" status from the volume's metadata share.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
metadata critical
no metadata critical
```

Default(s)

None

Guidelines

This command marks the volume's current metadata share as a critical resource. If the volume loses contact with its metadata share, the switch where the volume is running may fail over to the peer switch. (The failover occurs only if the peer switch has access to *all* of its critical resources.) Use the show redundancy critical-services command to see a list of all critical resources on this switch.

The ARX tests for failure with a regular ICMP ping. Every 30 seconds, the ARX sends a ping to the metadata share's filer. If the filer fails to respond, the ARX waits an additional 30 seconds before asking the peer if it is possible to fail over. The pings continue indefinitely at 30-second intervals. If the filer responds before the failover is initiated, the failover does not occur.

If this switch has a redundant peer, this is recommended for all volumes. This command is irrelevant unless redundancy is configured. See redundancy.

Sample

```
bstnA(gbl-ns-vol[wwmed~/acct])# metadata critical
```

declares that the metadata share for wwmed~/acct is a critical resource.

Related Commands

namespace -> volume
metadata share
redundancy
show redundancy critical-services

metadata share

Purpose Use the metadata share command to make a dedicated metadata share available for the current managed volume, or for all managed volumes in the current namespace.

Use the no form of the command to remove the metadata share from the current managed volume or namespace.

Modes gbl-ns

gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
metadata share filer {nfs3|nfs3tcp|cifs} path [cluster cl-name]
no metadata share filer {nfs3 | nfs3tcp | cifs} path
  [cluster cl-name]
```

filer (1-64) is the name of an external filer; use show external-filer to list all of them.

nfs3 | nfs3tcp | cifs chooses the protocol for file-access. This can be outside the set of protocols for the namespace.

path (1-900 for NFS, 1-1024 for CIFS) is the share path on the filer (for example, /arx_meta).

cl-name (optional, 1-64 characters) is only relevant if the ARX is part of a disaster-recovery (DR) configuration. In a DR configuration, there is an active ARX cluster with one set of filers and a backup cluster with a mirrored set of filers. This determines which cluster uses the *filer*. Run the metadata share command twice per volume if you use DR: once to designate the metadata share's host at the active cluster, and again to determine the metadata host at the backup cluster. Use show cluster for a list of configured clusters. If you omit this option, the CLI applies the change to the local cluster.

Default(s) None

Guidelines

A CIFS namespace can use a CIFS or NFS metadata-only share, but an NFS namespace is limited to metadata-only shares that also support NFS. An NFS namespace does not have the proxy-user (gbl-ns) credentials that it needs to access a CIFS metadata-only share.

If this switch has a redundant peer, use the gbl-ns-vol metadata critical command to declare the volume's metadata-only share as a critical resource.

If possible, use one metadata share per managed volume. Multiple managed volumes and/or namespaces can use the same metadata-only share, but this is not recommended.

A metadata share should reside on an extremely fast and reliable filer. Ideally, it should reside in its own file system on the same filer where you store the ARX volume's Tier 1 shares. The show statistics metadata command shows the latency between the ARX and its metadata shares. You can use the nsck ... migrate-metadata command to migrate a volume's metadata to a new filer, even after the volume has imported all of its back-end shares. This takes the volume offline during the migration. It is a safe operation in that it reverts the volume back to its original metadata share if it is interrupted.

For the best volume performance, configure all NFS-metadata shares with the 'sync' and 'no_wdelay' options. These options cause the back-end NFS service to synchronously write to disk without any caching, to improve both reliability and speed. You set these options at the back-end filer, in the configuration for the NFS export. This is not an issue with CIFS-metadata shares, which perform synchronous writes to disk on request.

Do not store your metadata on /vol/vol0 on a NetApp filer. This is where the NetApp filer keeps its operating system; it is not intended for fast access.

Sample

```
bstnA(gbl-ns[wwmed])# metadata share nas1 nfs3 /vol/vol1/meta1
```

sets a metadata-only share to hold all managed-volume metadata for the "wwmed" namespace.

Related Commands

metadata critical
show external-filer
show global-config **namespace**
nsck ... migrate-metadata
show statistics metadata

modify

Purpose

When a managed volume imports files, the import fails if there are any file collisions (that is, if the same file name and path exist on two or more shares). Further, clients cannot modify any files in the volume; the volume is read-only after the import. To allow managed-volume modifications, use the modify command.

Use the no form of the command to prevent all modifications to imported files.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
modify
no modify
```

Default(s)

```
no modify
```

Guidelines

This command does not apply to direct volumes (see direct); only metadata-based managed volumes.

The CLI prompts you with a question about nsck rebuild. The nsck ... rebuild command takes the volume offline and re-imports with a single command. Answer **yes** to turn modify on after a rebuild, or answer **no** to leave it off. A **no** answer makes the volume read-only after any re-import. If you answer **no**, you can use the reimport-modify command to achieve the same effect.

If the modify command is in effect when the volume is enabled, it renames collided files according to the following syntax:

pathname_share-jobid.ext

where

pathname is the original pathname minus any extension that the file may have had,

share is the name of the namespace share that imported the file,

jobid is a unique ID number for this import job, and

ext is the file's original extension.

Directories are only renamed if they have the same name as a previously-imported file, or if they have different file attributes (such as permissions settings) than an already-imported directory. A directory is renamed according to the same conventions as a file.

Whether or not you enabled modifications before import, each share's import report shows any and all duplicate files and directories in the share. Use show reports to get a list of import reports. Import reports are named "import.*share-id*.*share-name*.*job-id*.rpt." Use show reports *import-report* to view the contents of an import report. Duplicate files and directories are each called out in a separate line starting with "Duplicate."

Guidelines (Cont.)

If you enable a volume without modify enabled, clients cannot write to the volume. This creates an opportunity to check the import reports and resolve any file conflicts at your back-end filers. If there where no collisions, use the modify command to enable the volume for client writes. If there were collisions, correct them at the back-end and/or accept them. Then use the nsck ... destage command to take the volume offline, allow modifications with modify, then use enable (gbl-ns-vol-shr) on each of the volume's shares.

Guidelines: A "no modify" Import

You can opt to test an import before committing it with a "no modify" import. This means importing to the managed volume with the modify flag down, then checking all of the import reports for collisions or other issues. Use show reports type Imp to list all import reports, and show reports *report-name* to read one.

If no issues occurred, you can use the modify command after the import to allow write access to clients.

Guidelines: Per-Share Exceptions

You can protect certain shares in the managed volume from file and/or directory renaming. Use the no import rename-files command to disallow file renames in a particular share. Use no import rename-directories to disallow directory renames. To allow the volume to synchronize the attributes of matching directories instead of renaming them, use import sync-attributes; if the volume is allowed to synchronize attributes or rename directories, it synchronizes the directory attributes.

The above preference only applies to directories that collide with other directories. For a directory that collides with an already-imported file, the conflict cannot be resolved with an attribute change.

Shares in multi-protocol volumes have an additional import option for directories: import rename-directories unmapped-unicode. This allows the volume to rename directories whose names contain non-mappable characters; that is, multi-byte characters that are supported by CIFS but not by the setting for character-encoding nfs.

Samples

```
bstnA(gbl-ns-vol[wwmed~/acct])# modify
Automatically re-enable volume modify mode during nsck rebuild?
[yes/no] yes
```

allows the wwmed~/acct volume to modify its back-end shares.

```
bstnA(gbl-ns-vol[ns1~/])# no modify
```

disallows modifications for the "/" volume in the "ns1" namespace.

Related Commands

namespace -> volume
enable (gbl-ns-vol-shr)
show reports
nsck ... destage
nsck ... rebuild
reimport-modify
import rename-files
import rename-directories
import sync-attributes

named-streams

Purpose

A *named stream* (or *Alternate Data Stream*) is a hidden file with meta-information about the main file, such as a summary description or a thumb-nail graphic. If any back-end-CIFS filer does not support named streams, you must disable the feature for its namespace volume. This applies to volumes in namespaces that support CIFS, not in NFS-only namespaces. Use the no named-streams command to stop the volume from using named streams.

Use the affirmative form, named-streams, to reinstate named streams.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no named-streams
named-streams
```

Default(s)

The volume probes the filer behind its first-enabled share and uses the named-streams setting from there.

Guidelines

The Windows Explorer application uses named streams for the Properties -> Summary information; a volume that does not support named streams may not be able to provide any information for the Summary tab. Similarly, a volume without named streams may not support thumb-nail views of its graphics files.

You cannot enable a share if its volume supports named streams and its back-end-CIFS filer does not. The enable operation fails with an error that lists all CIFS features that must be disabled, possibly including this one. Use the enable (gbl-ns, gbl-ns-vol) command to enable all shares in a new namespace, or use the enable (gbl-ns-vol-shr) command to enable a new share in an already-enabled namespace. Use the enable (gbl-ns-vol-shr) command to enable a share.

If you remove the share(s) or upgrade the back-end filer(s), you can reinstate this feature for the volume. See the ARX CLI Maintenance Guide for instructions on removing a share from a namespace.

You can use the show exports command to see all CIFS options for the share.

CIFS clients only see the results of this command if they connect *after* you invoke it.

Guidelines: Advertised File-System Type

This CIFS option, together with three others, determines the File-System type (NTFS or a FAT file system) that the managed volume advertises to its CIFS clients. Client applications may query the file-system type and use it to customize the way it presents files; for example, if the managed volume advertises that it is a FAT file system, a client application may show old-style file-access times for the volume's files.

The volume advertises one of three file-system types, depending on the CIFS options supported by its back-end shares:

- NTFS - if the volume supports named-streams, persistent-acls, unicode-on-disk, and all of the volume's back-end shares preserve case in their file/directory names ("MyFile.txt" is *not* converted to "MYFILE.TXT").
- FAT32 - if the volume supports unicode-on-disk and case preservation, but not the others.
- FAT - if the volume supports none of the above options.

On the advice of F5 Support, you can use the cifs file-system-name command to manually set this file-system name.

Samples

```
bstnA(gbl-ns-vol[medarcv~/usr])# no named-streams
```

shuts off named streams for the "/usr" volume in the "medarcv" namespace.

```
bstnA(gbl-ns-vol[ns1~/])# named-streams
```

reinstates named streams for the "/" volume in the "ns1" namespace.

Related Commands

namespace -> volume
enable (gbl-ns, gbl-ns-vol)
enable (gbl-ns-vol-shr)
show exports

nfs-param

Purpose Use the nfs-param command to change the NFS read size or write size between the current volume and its back-end shares.

Use the no form of the command to revert to the default.

Mode gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
nfs-param {rsize | wsize} {1024 | 2048 | 4096 | 8192 | 16384 |
32768 | 65536}
no nfs-param {rsize | wsize}
```

rsize | wsize determines if this is the NFS-read size or write size.

1024 | ... 65536 is the size of reads or writes, in bytes.

Default(s) 8192

Guidelines The nfs-param command sets the size of NFS reads or writes from a volume to its back-end shares.

This parameter applies only to volumes that support NFS. Confirm that all filers behind the volume can support the read and/or write size that you choose.

Samples

```
bstnA(gbl-ns-vol[wwmed~/acct])# nfs-param rsize 16384
bstnA(gbl-ns-vol[wwmed~/acct])# nfs-param wsize 16384
```

sets the same read and write sizes for the "wwmed~/acct" volume.

```
bstnA(gbl-ns-vol[ns1~/])# no nfs-param rsize
```

reverts to the default NFS-read size in "ns1~/."

Related Commands namespace -> volume

persistent-acls

Purpose

A volume with persistent Access-Control Lists (ACLs) can display the ACLs of its files and directories to its clients. If any of the volume's back-end-CIFS filers do not also support persistent ACLs, you must disable this feature for the volume. This applies to volumes in namespaces that support CIFS, not volumes in NFS-only namespaces. Use the no persistent-acls command to stop the volume from supporting persistent ACLs.

Use the affirmative form, persistent-acls, to reinstate persistent ACLs.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no persistent-acls
persistent-acls
```

Default(s)

The volume probes the filer behind its first-enabled share and uses the persistent-acls setting from there.

Guidelines

This is only recommended for volumes whose back-end filers all share this limitation. We do not recommend mixing shares that support ACLs with shares that cannot.

This controls the CIFS-client view of ACLs in the volume. If a Windows client accesses the "Properties" for any of the volume's files or directories, the "Security" tab does not appear unless the volume has the persistent-acls setting.

You cannot enable a share if its volume supports persistent ACLs and its back-end-CIFS filer does not. The enable operation fails with an error that lists all CIFS features that must be disabled, possibly including this one. Use the enable (gbl-ns-vol-shr) command to enable a share.

A volume with no persistent-acls does not copy ACLs from one back-end share to another when it performs migrations. This applies to all migrations, including migrations between two filers that support persistent ACLs. This may be surprising in a volume backed by filers that cannot support ACLs and other filers that can support them. As stated above, we do not recommend mixing ACL-supporting filers with non-ACL-supporting filers behind the same managed volume.

If you remove the share(s) or upgrade the back-end filer(s), you can reinstate this feature for the volume. See the ARX CLI Maintenance Guide for instructions on removing a share from a namespace.

You can use the show exports command to see all CIFS options for the share, including persistent ACLs.

Guidelines: Advertised File-System Type

This CIFS option, together with three others, determines the File-System type (NTFS or a FAT file system) that the managed volume advertises to its CIFS clients. Client applications may query the file-system type and use it to customize the way it presents files; for example, if the managed volume advertises that it is a FAT file system, a client application may show old-style file-access times for the volume's files.

The volume advertises one of three file-system types, depending on the CIFS options supported by its back-end shares:

- NTFS - if the volume supports named-streams, persistent-acls, unicode-on-disk, and all of the volume's back-end shares preserve case in their file/directory names ("MyFile.txt" is *not* converted to "MYFILE.TXT").
- FAT32 - if the volume supports unicode-on-disk and case preservation, but not the others.
- FAT - if the volume supports none of the above options.

On the advice of F5 Support, you can use the cifs file-system-name command to manually set this file-system name.

Samples

```
bstnA(gbl-ns-vol[medarcv~/usr])# no persistent-acls
```

shuts off persistent ACLs for the "/usr" volume in the "medarcv" namespace.

```
bstnA(gbl-ns-vol[ns1~/])# persistent-acls
```

reinstates persistent ACLs for the "/" volume in the "ns1" namespace.

Related Commands

namespace -> volume
enable (gbl-ns-vol-shr)
show exports

reimport-modify

Purpose

An nsck ... destage or nsck ... rebuild shuts off the volume's modify setting as it takes the volume offline. The modify flag stays down when the shares reimport, so the volume does not modify any re-imported files *and* clients cannot modify the volume. To keep the modify flag up during a re-import, use the reimport-modify command.

Use the no form of the command to prevent all modifications to re-imported shares.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
reimport-modify
no reimport-modify
```

Default(s)

no reimport-modify

Guidelines

This command does not apply to direct volumes (see direct); only metadata-based managed volumes.

Raise this flag *before* you use nsck to rebuild or otherwise take the volume offline. The flag is only effective for future rebuilds.

The CLI prompts for confirmation before raising the flag; answer **yes** to continue.

The modify command determines whether or not modifications occur; this command determines whether or not to re-instate the volume's modify status after the volume is taken offline. This is an extra security measure to ensure that the system avoids any unexpected file renames.

When you invoke the modify command, a prompt requests whether or not you want to set this flag at the same time.

Whether or not you enabled modifications before import, each share's import report shows any and all duplicate files and directories in the share. Use show reports to get a list of import reports. Import reports are named "import.*share-id*.*share-name*.*job-id*.rpt." Use show reports *import-report* to view the contents of an import report. Duplicate files and directories are each called out in a separate line starting with "Duplicate."

Samples

```
bstnA(gbl-ns-vol[wwmed~/acct])# reimport-modify

Automatically re-enable volume modify mode during NSCK rebuild?
[yes/no] yes
```

allows the wwmed~/acct volume to modify its back-end shares after nsck brings it back online.

```
bstnA(gbl-ns-vol[ns1~/])# no reimport-modify
```

causes nsck to turn off the modify flag if it takes the "/" volume offline.

Related Commands

namespace **->** volume
enable (gbl-ns-vol-shr)
show reports
nsck ... destage
nsck ... rebuild

replica-snap

Purpose

You can designate a managed-volume share as a repository for snapshots, where the back-end filer behind the share holds snapshots from another share in the volume. You can use filer-replication methods to duplicate the files and directories from one filer to another, and you can take snapshots at the second filer without using disk space on the first. Use the replica-snap command to indicate that the current share is dedicated to snapshot storage in this way. This indicates that the managed volume should only present the back-end share's snapshots, and ignore any non-snapshot data. The managed volume's clients can then access these snapshots and restore previous versions of their files as needed.

Use the no replica-snap command to change the current share to a standard client-data share.

Mode

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
replica-snap
no replica-snap
```

Default(s)

```
no replica-snap
```

Guidelines

This command has no effect in a direct volume.

You can only run this command when the share is disabled (no enable (gbl-ns-vol-shr)). If the managed volume has already imported the share's primary storage, you must destage the volume to use the share (see nsck ... destage). Typically, the command is applied to new shares.

You cannot use no replica-snap after the share is enabled, but you can use no share to remove a replica-snap share. This removes all of the share's snapshots from client view.

After you designate the share as a replica-snap share, you create a replica-snap-rule to create managed snapshots on it. You can use the gbl-ns-vol snapshot replica-snap-rule command to create this rule. As mentioned above, clients see the snapshots from this share instead of the share's files and directories, and access those snapshots through their usual method (such as the "Previous Versions" tab). To present any snapshots created previously, you can use the snapshot manage command.

Important

The managed volume cannot write to most replica-snap shares, so it cannot check a replica-snap share to see if another ARX owns it (see the documentation for enable (gbl-ns-vol-shr) ... take-ownership*). Do not use the same replica-snap share behind two different ARXes, unless they are a redundant pair.*

Guidelines (Cont.)

If your managed volume uses tiered storage where recently-changed files reside on tier-1 shares and unchanged files reside on tier 2, we recommend one or more replica-snap shares per tier-1 share. The tier-2 shares have files that are unchanged, so they do not require snapshots. (You can use the place-rule command to configure tiered storage.)

Guidelines: Some Filer-Replication Applications Overwrite ARX Snapshots

Some filer-replication applications are volume-level, and copy the entire contents of the back-end share to the replica-snap share, including the filer snapshots. This would overwrite any snapshots that the ARX takes at the replica-snap share, and defeat the snapshot replica-snap-rule. Do not use volume-level replication to copy the source share's contents to the replica-snap share.

Guidelines: Filer Latency

Slow responses from a replica-snap share may cause slow responses for client access to files, too. The filer host for your replica-snap shares should not have an excessively long file-access latency.

Samples

```
bstnA(gbl-ns-vol-shr[medarcv~/lab_equipment~equipSnap])# replica-snap
```

designates the "equipSnap" share as a replica-snap share. The managed volume only presents snapshots from this share to its clients.

```
bstnA(gbl-ns-vol-shr[ns1~/vol~shr6])# no replica-snap
```

resets the 'shr6' share to be a standard file-storage share.

Related Commands

namespace -> volume -> share
enable (gbl-ns-vol-shr)
snapshot replica-snap-rule
snapshot manage
nsck ... destage

reserve files

Purpose

A managed volume requires one *file credit* for each of its files and directories. By default, the volume automatically increases its file credits as its number of files increases, but there are rare circumstances where you must manually set a volume's file credits. After you turn off auto reserve files for such a volume, you can use the reserve files command to manually set its file credits.

You can also use this command with auto-reserve enabled, though the auto-reserve feature may later change the number of reserved files.

Use the no form of the command to revert to a static default for the managed volume.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
reserve files files
no reserve files
```

files (4,096 to 128,000,000 (ARX-500) or to 256M (all other platforms)) is the new number of files to reserve for the volume. The maximums on your chassis may be lower due to your software license; use show active-license to see the limits imposed by your software license.

Default(s)

4,000,000 (4M) files if auto reserve files is set

64M files if no auto reserve files is set

Guidelines

Each file and directory in a managed volume (as opposed to a direct volume) uses one *file credit*. By default, the managed volume automatically takes more file credits as it grows. This automatic growth is the auto reserve files feature. The reserve files command sets the *current* number of reserved files if the auto-reserve feature is enabled; the auto reserve feature may then increase the number of file credits from there.

This command is more effective in a volume where you must turn off the auto-reserve feature; it sets a permanent file-credit reservation in this case. You must disable the feature (with no auto reserve files) in a volume used as a "filer" by an NFS-supporting direct volume. The documentation for auto reserve files discusses this further.

Each volume group supports a maximum of 384 Million (M) file credits. The ARX-500 can support 2 volume groups, and the ARX-2000 supports 12. The ARX-4000 supports 16 volume groups, but each volume group on the ARX-4000 supports a lower maximum (256M). This is the ceiling for the sum of all reserved files for all volumes in the volume group. You can use the show volume-group command to list the current number of reserved files for each volume; see the File credits section of the output.

Sample

```
bstnA(gbl-ns-vol[archives~/multimedia])# reserve files 5000000
```

reserves 5000000 files for the '/multimedia' volume.

Related Commands

namespace -> volume
auto reserve files
show volume-group

restart namespace ... volume

Purpose Use the restart namespace ... volume command to restart processing in a specific volume. This also restarts all other volumes that run in the same volume group.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`restart namespace`** ***`name`*** **`volume`** ***`volume`*** **`[dump-core]`**

name (1-30 characters) identifies the namespace.

volume (1-1024 characters) is the volume to restart. All volumes in the same volume group will also restart. Use to show volume-group see the assignments of volumes to volume groups.

dump-core (optional) causes the namespace process to write out its memory contents in a core-memory file. F5 personnel can examine this core file to diagnose problems with the namespace. This slows the restart; use this option only on the advice of F5 personnel.

Default(s) None.

Guidelines The CLI prompts for confirmation before restarting the namespace volume(s). If any front-end services export any of the affected volumes, the prompt lists the services. Enter **yes** to proceed.

This restarts all of the volumes in the chosen volume's group. A *volume group* can fail over from one chassis to another in a redundant configuration. It is also a failure domain for a group of volumes in the same namespace; if one of the volumes restarts, they all restart together. Use the show volume-group command to show all volume groups and their volumes. Before a volume is enabled, you can use volume-group to assign it to a group.

Sample

```
bstnA# restart namespace wwmed volume /acct

The Volume Group hosting namespace 'wwmed', volume '/acct'
is in use by NFS global service 'ac1.medarch.org'.

These services will experience a brief outage as the Volume Group
hosting the namespace 'wwmed', volume '/acct' is restarted.

Do you wish to proceed? [yes/no] yes
```

restarts the "wwmed~/acct" volume, along with any other volumes in its volume group.

Related Commands show volume-group
volume-group

share

Purpose Use the no form of the command to delete a managed-volume share.

Mode gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax `share` ***name***

name (1-64 characters) is the name you choose for the share.

`no share` ***name***

`no share` ***name*** `relocate-dirs` ***target-share*** `[remove-file-entries [verbose] [offline]]`

name (1-64 characters) identifies the share to remove.

relocate-dirs is required if you are removing an already-imported share from a managed volume. This is never required in a direct volume or a replica-snap share.

target-share (1-64 characters) is another share in the same managed volume. The volume migrates all master directories in the current share to this target.

You can enter the optional flags (**remove-file-entries**, **offline**, and **verbose**) in any order. As above, these options only apply to managed volumes and never apply to replica-snap shares:

remove-file-entries removes all files from volume metadata that still reside on this back-end share.

verbose (optional, if you choose **remove-file-entries**) causes the operation to list all removed files in its "removeShare" report.

offline (optional, if you choose **remove-file-entries**) is only for back-end shares that are offline or otherwise unreachable. This forces the disconnect without scanning the back-end for its directory attributes. Relocated master directories therefore have all of their file attributes set to 0 (zero).

Default(s) None

Guidelines A *share* maps to a CIFS share or NFS export on back-end storage. Use the share command to add a share in a managed volume or a direct volume (see direct).

The CLI prompts for confirmation before creating a new share; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new objects.) This places you in gbl-ns-vol-shr mode, where you must use the filer command to identify a filer and export/share, and then use the enable (gbl-ns-vol-shr) command to import the export/share. In a direct volume, you also use the attach command to map one or more attach-point directories in the share to physical directories on the back-end filer.

Guidelines: no share

You can remove the share with the simple no share command before it is first imported. There is no import for a share in a direct volume or for a replica-snap share, so you can use this simple form any time in those cases.

After a managed-volume share has been imported, you can use a place-rule to migrate all of its files to other shares in the volume. If the volume supports snapshots (with a snapshot rule, notification rule, or similar rule), we recommend waiting until all retained snapshots have "aged out" before removing the volume. After the last of the current snapshots is gone, none of the snapshots on this share have any record of any files in the volume. At that point, you can safely remove the share.

After the files (and, possibly, snapshots) are cleared, you use the relocate-dirs argument in no share to relocate the share's master directories, too. (A volume typically has multiple copies of its directories in each share, so that it can migrate files between them; a *master directory* is the first-imported (or only) instance of a directory in a volume.) The volume scans the back-end share for the file attributes of these directories, to be duplicated at the *target-share*. If this is impossible because the back-end share is offline, use the offline flag (along with the remove-file-entries flag) to create new instances of the directories with zeroed-out file attributes.

The CLI creates a "removeShare" report to show the progress of the no share command. Each report is named "removeShare.*share-name*.rpt," where *share-name* is the ARX-share name (not the name of the share at the filer).

The no share command fails if any of the client-visible files in the volume are on the share; the remove-file-entries flag removes the files from the volume metadata and allows the no share operation to succeed. The number of files removed appears in the removeShare report. The additional verbose flag adds the names of each removed file into the report. The additional offline flag, described above, applies to unreachable filers.

As an alternative to no share in a managed volume, you can use remove-share migrate to remove an imported managed-volume share, remove-share nomigrate to remove a managed-volume share that failed to import, or remove-share offline to remove a share that is unreachable.

Samples

```
bstnA(gbl-ns-vol[archives~/multimedia])# share lun77
This will create a new share.

Create share 'lun77'? [yes/no] yes
bstnA(gbl-ns-vol-shr[archives~/multimedia~lun77])#
```

creates a share named "lun77" in the current managed volume.

```
bstnA(gbl-ns-vol[archives~/multimedia])# no share nas15
```

deletes the "nas15" share from the volume.

```
bstnA(gbl-ns-vol[ns1~/])# no share test relocate-dirs share4
```

removes the "test" share, migrating all of its master directories to "share4" in the "ns1~/" volume.

Related Commands

namespace `->` volume
show namespace
show global-config **`namespace`**
show namespace status
filer
enable (gbl-ns-vol-shr)
place-rule
remove-share migrate
remove-share nomigrate

show share status

Purpose Use the show share status command to view the current status of all namespace shares on the ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show share status`

Guidelines The output contains a small table for every volume on the ARX. This information is an summary of the data from show namespace.

Each volume starts with a single-line summary. The summary contains the following fields:

NS identifies the namespace.

Vol is the path name of the volume.

MD is a high-level status of the volume's metadata share. This appears in *used*/*free* format, where *used* is the amount of space used for the volume's metadata, and *free* is the space remaining on the metadata share. You can ignore these numbers for a direct volume, which does not use a metadata share.

I/VG is the volume's *Instance* ID (used internally) and volume-group number, separated by a slash (/).

Files shows the number of files used and free, along with the maximum file allowable in volume. The maximum is labeled "(automatic)" for a volume with auto reserve files enabled.

Below each summary line is a table of the volume's shares, one row for each share. Each row contains two fields:

Share Name is the name established by the share command.

Status is the same share status documented for the show namespace command (see Guidelines: Share-Import Status, on page 21-36, Guidelines: Disable/Removal Status, on page 21-37, and Table 21.1 on page 21-49). You can use this to watch the progress of a share import.

Sample `bstnA# show share status`

shows the share status for all configured volumes. For sample output, see Figure 22.4 on page 22-84.

Related Commands show namespace status
show namespace

Figure 22.4 *Sample Output: show share status*

```
bstnA# show share status

NS:"medco", Vol:"/vol", MD:28 kB/-, I/VG:1/1, Files:1 used (1 dir), 31 M free

    Share Name         Status
    ------------------ ------------------------------
    corporate          Online
```

```
    generic            Online
    sales              Online

NS:"wwmed", Vol:"/acct", MD:2.0 MB/62 GB, I/VG:2/1, Files:4,433 used (439 dirs), 3.9 M free, 252 M
max (automatic)

    Share Name         Status
    ------------------ ------------------------------
    bills              Online
    bills2             Online
    budget             Online
    it5                Online

NS:"medarcv", Vol:"/lab_equipment", MD:92 kB/62 GB, I/VG:3/2, Files:66 used (7 dirs), 3.9 M free,
248 M max (automatic)

    Share Name         Status
    ------------------ ------------------------------
    backlots           Online
    equip              Online
    leased             Online
    scanners           Online

NS:"medarcv", Vol:"/rcrds", MD:136 kB/62 GB, I/VG:3/2, Files:188 used (25 dirs), 3.9 M free, 248 M
max (automatic)

    Share Name         Status
    ------------------ ------------------------------
    bulk               Online
    charts             Online
    rx                 Online

NS:"medarcv", Vol:"/test_results", MD:28 kB/62 GB, I/VG:3/2, Files:1 used (1 dir), 31 M free

    Share Name         Status
    ------------------ ------------------------------
    2005_charts        Online
    chemistry          Online
    hematology         Online

NS:"insur", Vol:"/claims", MD:156 kB/62 GB, I/VG:4/2, Files:172 used (21 dirs), 63 M free

    Share Name         Status
    ------------------ ------------------------------
    shr1-next          Online
    shr1-old           Online
```

show sid-translation

Purpose Use the show sid-translation command to view the translation of a user or group name to a numeric security ID (SID), or from SID to name. This shows the translation for every share in the current volume.

Mode gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax **`show sid-translation {`*`user`*` | `*`group`*` | `*`sid`*`}`**

user **|** ***group*** **|** ***sid*** chooses the entity to translate:

user (1-256 characters) is a Windows user name.

group (1-256 characters) is a Windows group name.

sid (1-256 characters) is a numeric SID; this does a reverse translation to the corresponding user or group name.

Guidelines This command does not apply to an NFS-only namespace.

The output shows the SID-to-name translation from each share in the current volume:

Share is the name of the namespace share.

SID is the globally-unique ID to identify the principal (user or group) below.

Name is the name of the requested user, group, or SID. In parenthesis is the type of this principal: “user,” “group,” “domain,” “alias” (a locally-defined principal), “well-known group” (such as “everyone”), “deleted account,” “invalid,” or “unknown.”

If any of the shares have unique translations, some or all of the shares are configured for Local Groups. If the volume migrates a file to or from a share with Local Groups, the volume must translate the SID to a name on the source share and then translate it back on the destination share.

For any share whose filer uses Local Groups, use the gbl-ns-vol-shr sid-translation command to tell the volume to translate all of its SIDs.

Sample

```
bstnA(gbl-ns-vol[medarcv~/rcrds])# show sid-translation jqpublic
SID Translations:

  Share rx
    SID:  S-1-5-21-448539723-1844237615-725345543-1217
    Name: MEDARCH\jqpublic (user)

  Share charts
    SID:  S-1-5-21-448539723-1844237615-725345543-1217
    Name: MEDARCH\jqpublic (user)

  Share bulk
    SID:  S-1-5-21-448539723-1844237615-725345543-1217
    Name: MEDARCH\jqpublic (user)
```

Related Commands namespace -> volume -> share
sid-translation

show volume-group

Purpose A volume is permanently assigned to a *volume group* when it is enabled. The volume group shares a memory pool as well as CPU cycles and other resources. Use the show volume-group command to view the current volume-group assignments.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show volume-group [id] [detailed] [legacy]`

id (optional, 1-255) identifies a particular volume group. If you omit the number, this command displays all volume groups.

detailed (optional) adds details to the output: the current CPU and memory usage for each processor behind the volume group.

legacy (optional) presents the same output using terminology from former ARX releases, where volume groups were called "VPU domains," and there were two or more of them in each "VPU."

Guidelines This applies to the current switch only, not both switches in a redundant pair. To see the volume groups in the redundant peer, log into that peer and run the command from its CLI.

System Credits is a table of system-wide limits for the ARX:

Share credits lists the number of shares configured in the system, share credits remaining, and the total number of shares allowed. These are managed-volume shares, not direct-volume shares.

Direct Share credits lists the number of direct shares in the system, direct-share credits remaining in this ARX, and the total number of direct shares allowed in the ARX. Direct shares are the shares in a direct volume; a direct volume is declared with the direct command.

Volume credits lists the number of volumes currently in the system, volume credits remaining, and the total number of volumes allowed.

File credits lists the number of files currently in the system, file credits remaining, and the total number of files allowed. These credits only apply to a to the managed volumes on the system. A volume automatically sets its file-credit reservation if it has auto reserve files enabled. You can manually set the number of file credits with the reserve files command.

The "*n* credits remain" is not a guarantee for any given volume in the system. Other volumes in other volume groups draw from the same pool of file credits. If the system's file credits are sufficiently low, all volumes in the system share the same credits; if volume A uses 100 file credits, volumes B and C each lose 100 remaining credits, too.

Guidelines (Cont.) Volume Group *n* is a table to describe one of the volume groups on this system. One of these tables appears per volume group:

Physical Processor shows the actual CPU where the volume group is running. If you used the detailed keyword, this also shows current CPU and memory usage.

State is one of the following:

- off means that there is no physical CPU to back this volume group.
- empty means that there are no volumes assigned to the group.
- Normal indicates that there is room for more namespaces or volumes on this volume group. If credits have run out for one or more of these categories, one or more qualifiers appears here:
  - partially used indicates there is room for at least one more volume.
  - maximum instances means that you can add no more volume groups to the namespace. Both volume groups have a namespace assigned to them.
  - maximum volume capacity means that you cannot add any more volumes to the group.
- Degraded means that too many volume groups are sharing the same memory and processor resources. This only occurs after a failover between redundant peers.
- Unrunnable is very rare. This appears only if multiple hardware failures reduce the hardware resources to the level where there is not enough memory and/or CPU cycles to continue running the current volume group.
- Error is an unexpected error; please contact Customer Support if you see this.

Share credits lists the number of managed-volume shares configured in the volume group, share credits remaining, and the total number of shares allowed.

Direct Share credits lists the number of direct shares in the volume group, direct-share credits remaining, and the total number of direct shares allowed in the volume group. Direct shares are the shares in a direct volume; a direct volume is declared with the direct command.

Volume credits lists the number of volumes currently in the volume group, volume credits remaining, and the total number of volumes allowed.

File credits lists the number of files currently in the volume group, file credits remaining, and the total number of files allowed. These credits only apply to the managed volumes on the system. A volume automatically sets its file-credit reservation if it has auto reserve files enabled. You can manually set the number of file credits with the reserve files command.

Then there is a table of namespaces and volumes assigned to the volume group, showing the state of each volume. If you select the legacy option, an additional column shows the VPU and VPU-domain number of the volume group; these are offered for users of former ARX releases that used this terminology.

Use the volume-group command to assign a volume to a group. Use the nsck ... migrate-volume command to re-assign a volume from one group to another.

Samples

`bstnA# show volume-group`

shows all volume groups on the system. See Figure 22.5 on page 22-89.

`bstnA# show volume-group 1`

focuses on volume group 1. See Figure 22.6 on page 22-90.

`stkbrgA# show volume-group 1 detailed`

shows volume group 1 on another chassis, with CPU and memory usage. See Figure 22.7 on page 22-91.

Related Commands

volume-group
nsck ... migrate-volume
auto reserve files
reserve files

Figure 22.5 *Sample Output: show volume-group*

```
bstnA# show volume-group

Switch: bstnA
----------------------------------------------------------------------
System Credits
--------------
Share credits: 15 shares used (1009 credits remain of total 1024)
Direct share credits: 9 direct shares used (16375 credits remain of total 16384)
Volume credits: 8 volumes used (248 credits remain of total 256)
File credits: 76 M files reserved (1.9 G credits remain of total 2.0 G)

Volume Group 1
------------
Physical Processor: 1.1
State: Normal; maximum instances
Share credits: 4 shares used (124 credits remain of total 128)
Direct share credits: 0 direct shares used (2048 credits remain of total 2048)
Volume credits: 1 volumes used (31 credits remain of total 32)
File credits: 4.0 M files reserved (252 M credits remain of total 256 M)

Namespace         Volume Group  Volume              State
---------         ------------  ------              -----
wwmed              1            /acct               Enabled

1 Namespace                   1 Volume

Volume Group 2
------------
Physical Processor: 1.1
State: Normal; maximum instances
Share credits: 9 shares used (119 credits remain of total 128)
Direct share credits: 5 direct shares used (2043 credits remain of total 2048)
Volume credits: 4 volumes used (28 credits remain of total 32)
File credits: 8.0 M files reserved (248 M credits remain of total 256 M)

Namespace         Volume Group  Volume              State
---------         ------------  ------              -----
medarcv            2            /rcrds              Enabled
medarcv            2            /lab_equipment      Enabled
medarcv            2            /test_results       Enabled
```

```
1 Namespace                   3 Volumes

Volume Group 9
------------
Physical Processor: 1.1
State: Normal; maximum instances
Share credits: 0 shares used (128 credits remain of total 128)
Direct share credits: 3 direct shares used (2045 credits remain of total 2048)
Volume credits: 1 volumes used (31 credits remain of total 32)
File credits: 0  files reserved (256 M credits remain of total 256 M)

Namespace         Volume Group  Volume               State
---------         ------------  ------               -----
medco              9            /vol                 Enabled

1 Namespace                   1 Volume

Volume Group 10
------------
Physical Processor: 1.1
State: Normal; maximum instances
Share credits: 2 shares used (126 credits remain of total 128)
Direct share credits: 1 direct shares used (2047 credits remain of total 2048)
Volume credits: 2 volumes used (30 credits remain of total 32)
File credits: 64 M files reserved (192 M credits remain of total 256 M)

Namespace         Volume Group  Volume               State
---------         ------------  ------               -----
insur              10           /claims              Enabled

1 Namespace                   1 Volume
```

Figure 22.6 *Sample Output: show volume-group 1*

```
bstnA# show volume-group 1

Switch: bstnA
-----------------------------------------------------------------------
System Credits
--------------
Share credits: 15 shares used (1009 credits remain of total 1024)
Direct share credits: 9 direct shares used (16375 credits remain of total 16384)
Volume credits: 8 volumes used (248 credits remain of total 256)
File credits: 76 M files reserved (1.9 G credits remain of total 2.0 G)

Volume Group 1
------------
Physical Processor: 1.1
State: Normal; maximum instances
Share credits: 4 shares used (124 credits remain of total 128)
Direct share credits: 0 direct shares used (2048 credits remain of total 2048)
Volume credits: 1 volumes used (31 credits remain of total 32)
File credits: 4.0 M files reserved (252 M credits remain of total 256 M)

Namespace         Volume Group  Volume               State
---------         ------------  ------               -----
wwmed              1            /acct                Enabled

1 Namespace                   1 Volume
```

Figure 22.7 *Sample Output: show volume-group 1 detailed*

```
stkbrgA# show volume-group 1 detailed

Switch: stkbrgA
----------------------------------------------------------------------
System Credits
--------------
Share credits: 4 shares used (380 credits remain of total 384)
Direct share credits: 1 direct shares used (6143 credits remain of total 6144)
Volume credits: 2 volumes used (94 credits remain of total 96)
File credits: 4.0 M files reserved (764 M credits remain of total 768 M)

Volume Group 1
------------
Physical Processor: 1.1 (17% CPU, 39% MEM)
State: Normal; partially used
Share credits: 4 shares used (188 credits remain of total 192)
Direct share credits: 1 direct shares used (3071 credits remain of total 3072)
Volume credits: 2 volumes used (46 credits remain of total 48)
File credits: 4.0 M files reserved (380 M credits remain of total 384 M)

Namespace         Volume Group  Volume              State
---------         ------------  ------              -----
bgh                1            /naumkeag_wing      Enabled

1 Namespace                  1 Volume
```

sid-translation

Purpose

One or more of your back-end CIFS filers may be configured for *Local Groups* and users. These filers use local Security IDs (SIDs) for their group/user names instead of those issued by the Domain Controller (DC). To migrate a file to or from a share with Local Groups, the volume must translate the group SIDs in the file's Access Control List (ACL). Use the sid-translation command to enable SID translation for the current share.

To disable SID translation for this share, use no sid-translation.

Mode

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
sid-translation
no sid-translation
```

Default(s)

```
no sid-translation
```

Guidelines

This command does not apply to an NFS-only namespace.

Whenever the volume migrates a file to or from a share tagged for SID translation, it translates each SID to a group name (at the source share), then translates the group name back to a SID (at the destination share). This ensures that the final SID is one that is properly configured at the destination filer.

All filers behind the volume should support all of each other's local-group names. If a file tagged with local-group name, "staff", migrates to another filer without any definition for a "staff" group, the volume copies the group SID without translating it. The SID is invalid on the filer, though it will be valid if it migrates back later. While the file remains on that filer, all Access Control Entries (ACEs) using that SID are ignored. To avoid these translation failures, duplicate all local group names and user names on all of the filers behind the volume.

Some filers return SID errors for unknown SIDs. This can occur if SID translation fails, due to one filer lacking a local-group or local-user name that is present on the other filers. By default, the policy engine cancels the migration if it receives a SID-translation error. Some file servers (for example, EMC Celerra servers) return errors for unknown SIDs but accept the file or directory anyway. You can use the ignore-sid-errors command to ignore the SID errors from these file servers.

Use the gbl-ns-vol show sid-translation command to translate SIDs to names or names to SIDs for every share in the volume.

Samples

```
bstnA(gbl-ns-vol-shr[medarcv~/rcrds~bulk])# sid-translation
```

causes the volume to translate SIDs for all files migrated to the "bulk" share.

```
bstnA(gbl-ns-vol-shr[insur~/claims~shr1-next])# sid-translation
```

causes the volume to translate SIDs for all files migrated to the "shr1-next" share.

Related Commands

namespace -> volume -> share
ignore-sid-errors
show sid-translation

sparse-files

Purpose

Some applications create "holes" in files with no data (that is, all zeros); a volume that supports *sparse files* like these does not use any disk space for those holes. If any back-end-CIFS filer does not support sparse files, you must disable the feature for its namespace volume. This applies to volumes in namespaces that support CIFS, not in NFS-only namespaces. Use the no sparse-files command to stop the volume from using sparse files.

Use the affirmative form, sparse-files, to reinstate sparse files.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no sparse-files
sparse-files
```

Default(s)

The sparse-files setting at the first-enabled share.

Guidelines

If the volume does not support sparse files, files with large blocks of zero's (or holes) use up disk space as though the holes are filled with substantive data.

You cannot enable a share if it supports sparse files and its back-end-CIFS filer does not. The enable operation fails with an error that lists all CIFS features that must be disabled, possibly including this one. Use the enable (gbl-ns-vol-shr) command to enable a share.

If you remove the share(s) or upgrade the back-end filer(s), you can reinstate this feature for the volume. See the ARX CLI Maintenance Guide for instructions on removing a share from a namespace.

You can use the show exports command to see all CIFS options for the share.

Samples

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# no sparse-files
```

shuts off sparse files for the "/lab_equipment" volume in the "medarcv" namespace.

```
bstnA(gbl-ns-vol[ns1~/])# sparse-files
```

reinstates sparse files for the "/" volume in the "ns1" namespace.

Related Commands

namespace -> volume
enable (gbl-ns-vol-shr)
show exports

unicode-on-disk

Purpose

A volume that supports *unicode* on disk can support file names with any of the multi-byte characters (such as Korean or Japanese characters) supported by the Unicode character set. If any back-end-CIFS filer does not support Unicode file names on disk, you must disable the feature for its namespace volume. This applies to volumes in namespaces that support CIFS, not in NFS-only namespaces. Use the no unicode-on-disk command to stop the volume from using Unicode file names on disk.

Use the affirmative form, unicode-on-disk, to reinstate Unicode file names on disk.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no unicode-on-disk
unicode-on-disk
```

Default(s)

The volume probes the filer behind its first-enabled share and uses the unicode-on-disk setting from there.

Guidelines

A Windows client cannot use multi-byte Unicode characters in any file or directory names unless the volume supports unicode on disk.

You cannot enable a share if it supports Unicode file names on disk and its back-end-CIFS filer does not. The enable operation fails with an error that lists all CIFS features that must be disabled, possibly including this one. Use the enable (gbl-ns-vol-shr) command to enable a share.

If you remove the share(s) or upgrade the back-end filer(s), you can reinstate this feature for the volume. See the ARX CLI Maintenance Guide for instructions on removing a share from a namespace.

You can use the show exports command to see all CIFS options for the share.

Guidelines: Advertised File-System Type

This CIFS option, together with three others, determines the File-System type (NTFS or a FAT file system) that the managed volume advertises to its CIFS clients. Client applications may query the file-system type and use it to customize the way it presents files; for example, if the managed volume advertises that it is a FAT file system, a client application may show old-style file-access times for the volume's files.

The volume advertises one of three file-system types, depending on the CIFS options supported by its back-end shares:

- NTFS - if the volume supports named-streams, persistent-acls, unicode-on-disk, and all of the volume's back-end shares preserve case in their file/directory names ("MyFile.txt" is *not* converted to "MYFILE.TXT").
- FAT32 - if the volume supports unicode-on-disk and case preservation, but not the others.
- FAT - if the volume supports none of the above options.

On the advice of F5 Support, you can use the cifs file-system-name command to manually set this file-system name.

Samples

```
bstnA(gbl-ns-vol[medarcv~/usr])# no unicode-on-disk
```

shuts off Unicode file names on disk for the "/usr" volume in the "medarcv" namespace.

```
bstnA(gbl-ns-vol[ns1~/])# unicode-on-disk
```

reinstates Unicode file names on disk for the "/" volume in the "ns1" namespace.

Related Commands

namespace -> volume
enable (gbl-ns-vol-shr)
show exports

volume

Purpose

Use the volume command to create a new volume or edit an existing one. A volume appears to clients as a discrete file system in the namespace. If it is a managed volume, it contains imported shares from various back-end filers. A direct volume contains attach points to the filers behind it.

Use the no form of the command to delete a volume.

Mode

gbl-ns

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
volume path
no volume path
```

path (1-256 characters) is a directory path you choose for the managed volume (for example, "/multimedia"). This is the root directory for the managed volume.

Default(s)

None

Guidelines

The CLI prompts for confirmation before creating a new volume; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new objects.)

This places you in gbl-ns-vol mode, from which you include one or more exports or shares from back-end filers. From gbl-ns-vol mode, you must use the share command to create at least one share in the current volume; each share connects to one export/share from a back-end filer.

The remaining configuration options are dependant on the type of volume you are configuring. A *managed volume* imports the files and directories from its shares, manages metadata for all of them, and supports policies that migrate the files. A *direct volume* has virtual subpaths that attach to directories on the filer; this does not manage metadata or support any policy.

Guidelines: Managed Volumes

In a managed volume, the files are imported from each filer export/share to the root ***path*** of the managed volume. A file in the root of the imported export/share also appears in the root of the managed volume. This introduces the possibility of naming collisions. If your first filer share has a file named "myFile.txt" in its root, you cannot import another filer share to the same managed volume if it also has a "myFile.txt" file in its root. Optionally, you can allow the volume to rename the second file: use the modify command to do this.

A managed volume offers the option to group the shares into share farms; you can then apply capacity-balancing policies to each share farm. Use the share (gbl-ns-vol-sfarm) command to create a share farm in the current managed volume.

Guidelines: Direct Volumes

In a direct volume, the ***path*** is a read-only directory that contains the volume's *attach points*. Each attach point is a directory that directly connects a client to a back-end share. Use the direct command to make a volume into a direct volume. After you create a share with the share command, use attach to attach a directory from the volume to a back-end filer.

Guidelines: no volume

To prepare for removing a volume, all of the volume's contents must be removed first. That is, it must be deactivated (nsck ... destage), it can contain no rules (various commands), share farms (no share-farm), filesets (various commands), or shares (use remove-share migrate or no share). Additionally, no front-end service can be exporting the volume (no cifs or no nfs), and no CIFS service for this namespace can have browsing enabled. (You can work around this problem by disabling the volume with no enable (gbl-ns, gbl-ns-vol)).

The remove namespace ... volume command performs all of the above steps for you. Best practices dictate that you use that command instead.

Samples

```
bstnA(gbl-ns[archives])# volume /multimedia
This will create a new volume.

Create volume '/multimedia'? [yes/no] yes
bstnA(gbl-ns-vol[archives~/multimedia])#
```

creates a managed volume at /multimedia.

```
bstnA(gbl-ns[archives])# no volume /radio
```

deletes the "/radio" volume.

Related Commands

namespace
share
share (gbl-ns-vol-sfarm)
modify
direct
attach
remove namespace
show global-config **namespace**
show namespace
show namespace mapping
show namespace status

volume-group

Purpose

A *volume group* shares a single memory pool along with CPU-processing time and other resources. If certain catastrophic failures cause a single volume to fail, they can stop processing for all other volumes in the same group. You should therefore group your volumes carefully, to insulate volumes from one another as needed. Use the volume-group command to explicitly choose the volume's group rather than allowing a default assignment.

Use no volume-group to allow default-group assignment.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
volume-group id
no volume-group
```

id is the group ID you assign to the current volume. All the volumes in a given group must belong to the same namespace. On an installation with many volumes in a small number of namespaces, you have the option to divide each namespace's volumes among multiple volume groups.

On the ARX-500 you are limited to half of your potential volume groups. This ensures that the volume software has enough memory to function in most cases, but it reduces the potential number of failure domains that you can use. On those platforms, you can increase the maximum number of volume groups with max-volume-groups.

Default(s)

Internal software chooses a default volume-group assignment for each volume. We recommend choosing volume groups explicitly with this command or nsck ... migrate-volume instead of using the default.

Guidelines

On every platform except the ARX-2000 or ARX-4000, each volume group can hold up to 64 volumes. An ARX-2000 volume group can hold up to 48 volumes, and an ARX-4000 volume group can hold up to 32 volumes.

The namespace software assigns a volume to its group when the volume is enabled (see enable (gbl-ns, gbl-ns-vol)). If necessary, you can use the nsck ... migrate-volume command to can change the volume-group assignment after the volume is enabled. Use the show volume-group command to view the current volume-group assignments on the current switch.

It is possible for default volume-group assignments to artificially limit the maximum number of namespaces on your system. Use this command to make explicit volume-group assignments before the volumes are enabled.

Guidelines: Splitting a Namespace

If a volume has a catastrophic metadata failure (such as losing contact with the back-end filer), all other volumes in the same volume group become unusable. To mitigate this potential problem, you can divide a namespace's volumes into multiple independent groups. If one volume has a catastrophic failure in one group, volumes in the other group(s) are unaffected.

On the ARX-500 you are limited to half of your potential volume groups. This ensures that the volume software has enough memory to function in most cases, but it reduces the potential number of failure domains that you can use. On those platforms, you can increase the maximum number of volume groups with max-volume-groups. The number of volume groups affects the memory usage for volume processing (see metadata cache-size), so you should only change this on the advice of F5 personnel.

Samples

```
bstnA(gbl-ns[archives])# volume /multimedia
bstnA(gbl-ns-vol[archives~/multimedia])# volume-group 1
bstnA(gbl-ns-vol[archives~/multimedia])# exit
bstnA(gbl-ns[archives])# volume /radio
bstnA(gbl-ns-vol[archives~/radio])# volume-group 1
bstnA(gbl-ns-vol[archives~/radio])# ...
```

assigns two volumes to volume group 1. That volume group is now dedicated to volumes from the "archives" namespace.

```
bstnA(gbl-ns[ns3])# volume /vol1
bstnA(gbl-ns-vol[ns3~/vol1])# volume-group 2
bstnA(gbl-ns-vol[ns3~/vol1])# exit
bstnA(gbl-ns[ns3])# volume /vol2
bstnA(gbl-ns-vol[ns3~/vol1])# volume-group 3
bstnA(gbl-ns-vol[ns3~/vol1])# ...
```

assigns two volumes to different volume groups. If /vol1 has a catastrophic metadata failure, /vol2 is unaffected. This also means that both volume groups are dedicated to the "ns3" namespace.

```
bstnA(gbl-ns[archives])# volume /multimedia
bstnA(gbl-ns-vol[archives~/multimedia])# no volume-group
bstnA(gbl-ns-vol[archives~/multimedia])# ...
```

commits the "/multimedia" volume to default-group assignment.

Related Commands

namespace -> volume
show volume-group
enable (gbl-ns, gbl-ns-vol)
nsck ... migrate-volume
share (gbl-ns-vol-sfarm)
max-volume-groups

wait-for shares-online

Purpose Use the wait-for shares-online command to wait until all shares in a volume are enabled, and have started their imports.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`wait-for shares-online`** ***`namespace volume`*** **`[timeout`** ***`timeout`*****`]`**

namespace (1-30 characters) identifies the namespace.

volume (1-1024 characters) is the name of the volume.

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - 0 (zero, meaning that the command should wait indefinitely)

Guidelines This command is useful in CLI scripts.

When enabling a volume (enable (gbl-ns, gbl-ns-vol)) or all of its shares, you can use the wait-for shares-online command to wait for the volume's shares to be enabled. This command waits for all shares in the given volume.

Once this command returns, clients can access the shares and use them. The volume continues to import the shares as the clients access them.

If you set a ***timeout*** and it expires before all of the shares are online, the command exits with a warning.

Sample

```
bstnA(gbl-ns-vol[ns77~/vol31])# wait-for shares-online ns77 /vol31
timeout 60
```

waits for all shares from the 'ns77~/vol31' volume to come online. If 60 seconds expire before this happens, the command exits with a warning.

Related Commands enable (gbl-ns-vol-shr)

wait-for volume-disable

Purpose After disabling a volume or namespace, you can use the wait-for volume disable command to wait for the volume(s) to go offline.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`wait-for volume-disable`** ***`namespace`*** **`[`**`volume`**`] [timeout`** ***`timeout`*** **`]`**

namespace (1-30 characters) identifies the namespace.

volume (optional, 1-1024 characters) is a particular volume.

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - 0 (zero, meaning that the command should wait indefinitely)

Guidelines This command is useful in CLI scripts.

When disabling a volume or a namespace (no enable (gbl-ns, gbl-ns-vol)), you can use the wait-for volume-disable command to wait for one or all volumes to go offline.

If you set a ***timeout*** and it expires before all the chosen volumes are enabled, the command exits with a warning.

Sample `bstnA(gbl-ns[medarcv])# wait-for volume-disable medarcv`

waits for all volumes in the 'medarcv' namespace to be disabled. There is no timeout specified, so the CLI blocks until all volumes are offline.

Related Commands enable (gbl-ns, gbl-ns-vol)
wait-for volume-enable

wait-for volume-enable

Purpose Use the wait-for volume enable command to wait for one or more volumes to come online.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`wait-for volume-enable`** ***`namespace`*** **`[`**`volume`**`] [timeout`** ***`timeout`*** **`]`**

namespace (1-30 characters) identifies the namespace.

volume (optional, 1-1024 characters) is a particular volume.

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - 0 (zero, meaning that the command should wait indefinitely)

Guidelines This command is useful in CLI scripts.

When enabling a volume or a namespace (enable (gbl-ns, gbl-ns-vol)), you can use the wait-for volume-enable command to wait for one or all volumes to come online. To wait for all of the shares in a single volume, you can use wait-for shares-online instead.

If you set a ***timeout*** and it expires before all the chosen volumes are enabled, the command exits with a warning.

Sample `bstnA>` **`wait-for volume-enable wwmed timeout 120`**

waits for all volumes in the 'wwmed' namespace to be enabled. If 120 seconds pass before the last volume comes online, the command exits with a warning.

Related Commands enable (gbl-ns, gbl-ns-vol)
wait-for shares-online
wait-for volume-disable

23

Global Server

active-directory alias

Purpose

A Windows client typically provides a *service-principal name* (SPN, such as "fs1" or "fs1.medarch.org") when connecting to a virtual server. An Active-Directory (AD) SPN, as opposed to an IP address, is required for Kerberos authentication. Use the active-directory alias command to set an official SPN (or multiple SPNs) for this server.

Use the no form of the command to remove one of this server's SPNs from the AD database.

Mode

gbl-gs-vs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
active-directory alias name
no active-directory [alias]
```

name (1-256 characters; optional in the no form) is a SPN for the current global server. If you omit this from the no form of the command, it removes all SPNs for the virtual server. The CLI prompts for confirmation before removing all of the server's SPNs.

Default(s)

None

Guidelines

This command is only relevant to a virtual server whose global server offers cifs service, and whose clients use Kerberos authentication.

The global server requires strong credentials for setting SPNs at a Domain Controller (DC). Before you run this command, go to gbl-gs mode and use the active-directory proxy-user command to identify a proxy-user with permissions to do this.

This command causes the ARX to set the given SPN at the local DC. To identify the DC where this is set, use show global server to find the Windows domain, then use show active-directory domain *domain-name* to map the domain to its DC(s).

The ARX software sets both a "HOST" SPN and a "CIFS" SPN, mapping them back to the virtual server's VIP. This occurs if (and only if) both the virtual server and the global server are enabled (enable (gbl-gs, gbl-gs-vs)). You can run this command to store the SPN in the ARX database, and then enable the global and virtual server later to set the SPN at the DC. While both servers are running, you can re-set the SPN at the DC by removing it (with no active-directory alias) and adding it back.

Guidelines: Setting the Same Aliases in DNS and/or WINS

This command makes Kerberos authentication possible for clients that use the given SPN, but it has no effect on routing. That is, the SPN must also be set as a DNS alias and/or a WINS alias in your network. You can set up the global server so that it automatically registers its aliases at your local DNS or WINS server:

- from gbl-cifs mode, use dynamic-dns to register aliases with a DNS server;
- from gbl-gs mode, use wins-name or wins-alias to register aliases with a WINS server.

Conversely, if a DNS-based name exists for this global server that is different from the global server's FQDN, clients can connect to the ARX service but Kerberos authentication always fails. In this case, use the active-directory alias command to create the SPNs for the server's DNS-based FQDN. For example, if clients can use the DNS-based "ac1.wwmed.com" to reach the "ac1.medarch.org" global server, use active-directory alias ac1.wwmed.com to create the SPNs for that DNS-based name.

Sample

```
bstnA(gbl-gs-vs[ac1.medarch.org~192.168.25.15])# active-directory alias
fs2
```

sets a new SPN, "fs2," for the virtual server at "192.168.25.15."

Related Commands

active-directory proxy-user
enable (gbl-gs, gbl-gs-vs)
dynamic-dns
wins-name
wins-alias

active-directory proxy-user

Purpose

A Windows client typically provides a *service-principal name* (SPN, such as "insur.medarch.org") when connecting to a global server. Kerberos authentication fails if the SPN is unregistered in the Active-Directory (AD). You use the active-directory alias command to register a SPN (or multiple SPNs) for this server. Use this command to assign Windows credentials (username, password, and Windows domain) for setting these SPNs.

Use the no form of the command to remove the proxy-user from this global server. This makes it impossible to set any more SPNs with the active-directory alias command.

Mode

gbl-gs

Security Role(s)

storage-engineer or crypto-officer

Syntax

`active-directory proxy-user` ***`name`***
`no active-directory proxy-user`

name (1-32 characters) is the Windows proxy user to associate with the current global server.

Default(s)

None

Guidelines

This command is only relevant to a global server that offers cifs service with Kerberos authentication.

From gbl mode, use the proxy-user command to add a proxy-user configuration to the ARX. Use the show proxy-user command to view all configured proxy-users and their associated usernames and Windows domains.

Choose a proxy user with privileges for setting SPNs at the local Domain Controller (DC). These privileges are stronger than those needed for a namespace. To find the local DC where the SPNs will be set, use show global server to find the Windows domain for this global server, then use show active-directory domain *domain-name* to map the domain to one or more DCs.

Sample

```
bstnA(gbl-gs[ac1.medarch.org])# active-directory proxy-user acoProxy2
```

associates proxy user "acoProxy2" with the current global server.

Related Commands

active-directory alias
proxy-user
show global server
show active-directory

description (gbl-gs)

Purpose Use the optional description command to set a descriptive string for the current global server. This appears in the show command.

Use the no form of the command to delete the description.

Mode gbl-gs

Security Role(s) storage-engineer or crypto-officer

Syntax

```
description text
no description
```

text (1-255 characters) is your description. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s) `no description`

Guidelines The description appears in the output for show global server.

Sample

```
bstnA(gbl-gs[www.wwmed.com])# description “global NFS server for
network hospitals”
```

specifies a description for the current global server.

Related Commands global server
show global server

enable (gbl-gs, gbl-gs-vs)

Purpose

Use the enable command to activate the current global or virtual server.

Use no enable to shut down the current server.

Modes

gbl-gs

gbl-gs-vs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
enable
no enable
```

Default(s)

```
no enable
```

Guidelines

The gbl-gs no enable command shuts down the global server. This stops all front-end services (nfs or cifs) for the global server. This is a graceful shutdown: the services block any new connections and wait for current clients to complete their current transactions.

The gbl-gs-vs no enable command shuts down the virtual server. This stops all services, too.

Each virtual server listens for clients at a Virtual-IP (*VIP*) address. After enabling a virtual server and its global server, you can use the wait-for vip-enable command to wait for a particular VIP to come online. You can use the wait-for vip-disable command to wait for a particular VIP to shut down.

Samples

```
bstnA(gbl-gs[www.wwmed.com])# enable
```

enables the current global server.

```
bstnA(gbl-gs-vs[www.wwmed.com~192.168.25.10])# no enable
```

disables the current virtual server.

Related Commands

global server
virtual server
wait-for vip-enable
wait-for vip-disable

global server

Purpose

A *global server* links a front-end NFS and/or CIFS service to a virtual-IP (VIP) address. Clients can access the front-end service(s) through the VIP. Use the global server command to create a global server.

Use the no form of the command to delete a global server.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
global server fqdn
no global server fqdn
```

fqdn (1-128 characters) is the fully-qualified domain name (for example, "eng-nfs.company.com") that you choose for the global server. By convention, this is the name that clients use to access the global server.

Default(s)

None

Guidelines

When you use this command to either create or delete a global server, the CLI prompts for confirmation. Enter **yes** to proceed.

This command places you in gbl-gs mode. From there, you must use the virtual server command to bind the global server to the ARX. The virtual server runs the global server's front-end services on its virtual-IP (VIP) address. Once you have a virtual server, you must configure and enable at least one front-end service to run on the global server; for example, you can use the nfs command to instantiate NFS service or the cifs command for CIFS service. You then use the enable (gbl-gs, gbl-gs-vs) command in gbl-gs mode to start the global server.

Global servers that offer CIFS as a front-end service require a Windows domain; configure the Windows domain with the windows-domain (gbl-gs) command. To support Kerberos authentication for multiple CIFS-service aliases, you will need to set each alias as a *service-principal name* (SPN); use active-directory proxy-user to choose a proxy user with permission to set these SPNs. (You can later use active-directory alias, in gbl-gs-vs mode, to set each SPN.)

You can configure an indeterminate number of global servers, but the number of *useful* global servers is limited by the maximum number of front-end services. The ARX-500 supports a maximum of 16 front-end services, and the remaining platforms support up to 64. This limit applies to the sum of all cifs services and nfs services.

To remove a global server, use the no enable command from gbl-gs mode, then exit to gbl mode and use the no global server command.

Use show global server to view the current configuration of the global server. To map the global server to the physical servers (filers) behind it, use show server-mapping. The show statistics global server command shows the volume of client traffic to the global server.

Samples

```
bstnA(gbl)# global server www.wwmed.com
This will create a new global server.

Create global server 'www.wwmed.com'? [yes/no] yes
bstnA(gbl-gs[www.wwmed.com])#
```

instantiates a global server at www.wwmed.com.

```
bstnA(gbl)# global server www.defunctcompany.com
bstnA(gbl-gs[www.defunctcompany.com])# no enable
bstnA(gbl-gs[www.defunctcompany.com])# exit
bstnA(gbl)# no global server www.defunctcompany.com
Delete global server 'www.defunctcompany.com'? [yes/no] yes
bstnA(gbl)# ...
```

removes the global server that was running at www.defunctcompany.com.

Related Commands

enable (gbl-gs, gbl-gs-vs)
virtual server
windows-domain (gbl-gs)
active-directory proxy-user
show global server
show server-mapping
show statistics global server

show global server

Purpose Use the show global server command to display configuration information about global servers.

Mode (any)

Security Role(s) operator

Syntax `show global server [fqdn]`

fqdn (1-128 characters) is the fully-qualified domain name (for example, "www.organization.org") for one global server.

If you do not include the global server FQDN, the command displays all global server configurations.

Guidelines The output of this command is a large table with sub-tables. The table's fields and sub-tables are described below.

Domain Name is the FQDN name for the global server.

State is the global server's administrative state (enabled or disabled). Use enable (gbl-gs, gbl-gs-vs) to enable the global server.

Windows Domain is the Windows Domain through which cifs service (if any) is provided. By default, the "short-domain name" for the service is the first component of this domain (for example, "MYCO" is the short name for "MYCO.COM"). If something other than the default is used for this, it appears in parentheses after the name. Use the windows-domain (gbl-gs) command to change the domain and/or the short-domain name.

The global-server table also contains a sub-table describing all of its virtual servers:

Switch is hostname for the virtual server's host switch. The *virtual server* runs the front-end services (NFS and/or CIFS) on behalf of the global server. The virtual server has its own administrative state, VIP, and VMAC.

State is enabled or disabled. Use enable (gbl-gs, gbl-gs-vs) to enable the virtual server.

VIP is the virtual IP (VIP) address, where clients connect to the virtual server. Use virtual server to set this.

VLAN is the VIP's VLAN number, also set with virtual server.

VMAC is the virtual MAC (VMAC) address for the VIP.

Guidelines (Cont.) Two additional sub-tables, labeled Aliases, show the various CIFS aliases for the service.

The SPNs table shows the server's Active-Directory aliases. These are also known as Service-Principal Names (SPNs). If any are defined, the table contains two columns:

Name is an alias for this global server. If a client uses this name to connect to the global server, the client can authenticate with Kerberos. Each alias (or SPN) has two instances: one "CIFS" instance and one "HOST" instance. You can use the active-directory alias command to add an alias for the global server.

Active Directory Status indicates the status of the SPN in both the AD database and the ARX database. "OK" indicates that the alias is in both databases, "Not Found" indicates that the AD database is missing the SPN, "Unmanaged" means that the AD database has the SPN but it was not configured on the ARX, and "Unknown" appears when the ARX failed to contact the DC or the DC returned an empty list of SPNs.

The WINS table shows the WINS alias for the global server.

WINS Name is the NETBIOS name of the CIFS service, set by the wins-name command.

WINS Server is the address of the local WINS server, set by the wins command.

Description is a string to describe the global server. You can set this with the description (gbl-gs) command.

Samples

```
bstnA> show global server
```

shows all the global servers on the ARX. See Figure 23.1, below.

```
bstnA> show global server www.wwmed.com
```

shows one global server, "www.wwmed.com."

Related Commands global server
virtual server
description (gbl-gs)

Figure 23.1 *Sample Output: show global server*

```
bstnA# show global server

Domain Name             State                Windows Domain
-------------------------------------------------------------------------
ac1.MEDARCH.ORG            Enabled              MEDARCH.ORG

  Switch                State    VIP                VLAN  VMAC
  ------------------------------------------------------------------------
  bstnA                 Enabled  192.168.25.15/24    25   00:0a:49:17:85:c0

  Aliases:

    SPNs
      Name                                              AD Status
      ------------------------------------------------  ----------------
      CIFS/insur                                              OK
      HOST/insur                                              OK
      CIFS/fs1                                                OK
      HOST/fs1                                                OK
      CIFS/fs1.MEDARCH.ORG                                    OK
```

```
      HOST/fs1.MEDARCH.ORG                                          OK
      CIFS/fs2                                                      OK
      HOST/fs2                                                      OK
      CIFS/fs2.MEDARCH.ORG                                          OK
      HOST/fs2.MEDARCH.ORG                                          OK
      CIFS/fs5                                                      OK
      HOST/fs5                                                      OK
      CIFS/fs5.MEDARCH.ORG                                          OK
      HOST/fs5.MEDARCH.ORG                                          OK
      CIFS/pubs_test_insur                                          OK
      HOST/pubs_test_insur                                          OK
      CIFS/pubs_test_insur.MEDARCH.ORG                              OK
      HOST/pubs_test_insur.MEDARCH.ORG                              OK

    WINS
      Name                                                  WINS Server IP
      ------------------------------------------------------------------------
      INSURANCE                                             192.168.25.102


  Description
  ------------------------------------------------------------------------
  CIFS and NFS server for hospital insurance claims

Domain Name             State                Windows Domain
------------------------------------------------------------------------
acopiaFiler             Enabled

  Switch                State    VIP                VLAN  VMAC
  ------------------------------------------------------------------------
  bstnA                 Enabled  192.168.25.12/24    25   00:0a:49:17:85:c0

  Aliases:

    SPNs      No SPNs are defined.
    WINS
      Name                                              WINS Server IP
      ------------------------------------------------------------------------
      ACOPIAFILER                                       (none)


  Description
  ------------------------------------------------------------------------
```

show server-mapping

Purpose Use the show server-mapping command to map shares on the client side of the switch to the physical servers behind it.

Mode (any)

Security Role(s) operator

Syntax

```
show server-mapping [namespace name | virtual-ip vip]
[ip-addresses]
show server-mapping status [ip-addresses]
```

name (optional, 1-30 characters) identifies one namespace on which to focus.

vip (optional) specifies that you want to focus on the chosen Virtual IP (VIP) address.

ip-addresses (optional) converts all external-filer names into IP addresses in the output.

status (optional) shows a one-word status for each back-end server.

If you do not include any options, the command shows all VIPs and namespaces.

Guidelines Virtual Server is the VIP address and name for one share on one virtual server. Use virtual server to create a virtual server. Then use cifs and export (gbl-cifs) to create a CIFS share, and/or use nfs and export (gbl-nfs) to create an NFS export.

Namespace/Volume is served by the front-end share. This is a virtual storage pool behind the share which connects to the Physical Server(s) below.

Virtual Path appears only for a direct volume. This is the *attach point* from the volume to a mount point or CIFS share on one of the filers. Clients see each of these paths as a sub-tree under the root of the front-end export/share. You can use the attach command to create an attach point to a filer share.

Physical Server is one of the server shares behind a namespace. If this has an asterisk (*) at the end, it is used as a metadata-only share (created with metadata share). A double-asterisk (**) indicates that the server is a managed volume that a direct volume is using as though it were a back-end filer; you can use the managed-volume command to use a managed volume this way.

For the **status** command, the output shows the status of each Virtual Server and Physical Server in the above output. A new field, Status, has a different interpretation for each server type:

> Status for the Virtual Server is an abbreviation for the show virtual service status. A status of "Ready" indicates that the service is functional. Refer to that command's documentation for details.
>
> Status for the Physical Server is an abbreviation for the share's import status, as seen in show namespace. See Table 21.1 on page 21-49 for details. A status of "Online" indicates that the back-end share is functional.

Samples

`bstnA> **show server-mapping**`

shows the physical servers behind all the global servers on the ARX. See Figure 23.2, below.

`bstnA> **show server-mapping namespace medarcv ip-addresses**`

shows the virtual server(s) in front of the medarcv namespace, in addition to the physical server(s) behind it. The display shows the IP addresses for the physical servers instead of their external-filer names. See Figure 23.3 on page 23-17.

`bstnA> **show server-mapping status**`

shows the status of all the virtual and physical servers. See Figure 23.4 on page 23-18.

Related Commands

show virtual service
show namespace

Figure 23.2 *Sample Output: show server-mapping*

```
bstnA> show server-mapping

Virtual Server               Namespace/Volume
   Virtual Path                 Physical Server
-------------------------------------------------------------------------
192.168.25.12:/vol           medco:/vol

   vol1/corp                     nas1:/vol/vol2/shr
   vol1/notes                    nas1:/vol/vol2/notes
   vol2                          nas3:/exports/data
   vol3/mtgMinutes               nas2:/vol/datavol1/direct/mtgs
   vol3/sales                    nas2:/vol/datavol1/direct/export


192.168.25.15:/acct          wwmed:/acct

                                 das1:/exports/budget
                                 das3:/exports/acct2
                                 das7:/lhome/it5
                                 das8:/work1/accting
                                 nas3:/exports/meta7*


192.168.25.15:/acct/wksheets wwmed:/acct

                                 das1:/exports/budget
                                 das3:/exports/acct2
                                 das7:/lhome/it5
                                 das8:/work1/accting
                                 nas3:/exports/meta7*


192.168.25.15:/claims        insur:/claims

                                 nas1:/vol/vol2/insurance
                                 nas1:/vol/vol2/meta2*
                                 nasE1:/patient_records
```

```
\\192.168.25.15\ARCHIVES     medarcv:/rcrds

                                \\fs1\histories
                                \\fs2\bulkstorage
                                \\fs4\prescriptions
                                nas1:/vol/vol2/meta3*


\\192.168.25.15\bulkstorage  medarcv:/rcrds

                                \\fs1\histories
                                \\fs2\bulkstorage
                                \\fs4\prescriptions
                                nas1:/vol/vol2/meta3*


\\192.168.25.15\CELEBS       medarcv:/rcrds

                                \\fs1\histories
                                \\fs2\bulkstorage
                                \\fs4\prescriptions
                                nas1:/vol/vol2/meta3*


\\192.168.25.15\chem_results medarcv:/test_results

   2005charts                   ** medarcv:/rcrds/2005
   chemLab                      \\fs1\chem_results/.
   hematologyLab                \\fs3\hematology_results/.


\\192.168.25.15\CLAIMS       insur:/claims

                                nas1:/vol/vol2/meta2*
                                \\nas1\insurance
                                \\nasE1\patient_records


\\192.168.25.15\labs         medarcv:/lab_equipment

                                \\fs2\backlot_records
                                \\fs5\xraysScanners
                                nas1:/vol/vol2/meta6*
                                \\nas10\equipment
                                \\nas10\for_lease
                                \\nas11\equipBkup
                                \\nas11\leasedBkup


\\192.168.25.15\MP3S         medarcv:/rcrds

                                \\fs1\histories
                                \\fs2\bulkstorage
                                \\fs4\prescriptions
                                nas1:/vol/vol2/meta3*


\\192.168.25.15\SPECS        insur:/claims

                                nas1:/vol/vol2/meta2*
                                \\nas1\insurance
                                \\nasE1\patient_records
```

```
\\192.168.25.15\STATS        insur:/claims

                                 nas1:/vol/vol2/meta2*
                                 \\nas1\insurance
                                 \\nasE1\patient_records


\\192.168.25.15\xraysScannersmedarcv:/lab_equipment

                                 \\fs2\backlot_records
                                 \\fs5\xraysScanners
                                 nas1:/vol/vol2/meta6*
                                 \\nas10\equipment
                                 \\nas10\for_lease
                                 \\nas11\equipBkup
                                 \\nas11\leasedBkup


\\192.168.25.15\Y2004        medarcv:/rcrds

                                 \\fs1\histories
                                 \\fs2\bulkstorage
                                 \\fs4\prescriptions
                                 nas1:/vol/vol2/meta3*


\\192.168.25.15\Y2005        medarcv:/rcrds

                                 \\fs1\histories
                                 \\fs2\bulkstorage
                                 \\fs4\prescriptions
                                 nas1:/vol/vol2/meta3*


\\192.168.25.15\Y2008        medarcv:/rcrds

                                 \\fs1\histories
                                 \\fs2\bulkstorage
                                 \\fs4\prescriptions
                                 nas1:/vol/vol2/meta3*


\\192.168.25.15\Y2010        medarcv:/rcrds

                                 \\fs1\histories
                                 \\fs2\bulkstorage
                                 \\fs4\prescriptions
                                 nas1:/vol/vol2/meta3*


\\192.168.25.16\MICE         labResearch:/labMice

                                 \\fs3\mice_lab_results
                                 \\fs4\white_mice
                                 nas1:/vol/vol2/meta8*


\\192.168.25.16\RATS         labResearch:/labRats

                                 \\fs6\labRats
```

```
                                \\fs7\lrTests
                                nas1:/vol/vol2/meta9*


Where * denotes metadata only physical server.
Where ** denotes a direct volume mapped to a namespace.
```

Figure 23.3 *Sample Output: show server-mapping namespace medarcv ip-addresses*

```
bstnA> show server-mapping namespace medarcv ip-addresses

Virtual Server               Namespace/Volume
   Virtual Path                 Physical Server
-------------------------------------------------------------------------
\\192.168.25.15\ARCHIVES     medarcv:/rcrds

                                \\192.168.25.20\histories
                                \\192.168.25.27\bulkstorage
                                \\192.168.25.29\prescriptions
                                192.168.25.21:/vol/vol2/meta3*


\\192.168.25.15\bulkstorage  medarcv:/rcrds

                                \\192.168.25.20\histories
                                \\192.168.25.27\bulkstorage
                                \\192.168.25.29\prescriptions
                                192.168.25.21:/vol/vol2/meta3*


\\192.168.25.15\CELEBS       medarcv:/rcrds

                                \\192.168.25.20\histories
                                \\192.168.25.27\bulkstorage
                                \\192.168.25.29\prescriptions
                                192.168.25.21:/vol/vol2/meta3*


\\192.168.25.15\chem_results medarcv:/test_results

   2005charts                   ** medarcv:/rcrds/2005
   chemLab                      \\192.168.25.20\chem_results/.
   hematologyLab                \\192.168.25.28\hematology_results/.


\\192.168.25.15\labs         medarcv:/lab_equipment

                                \\192.168.25.27\backlot_records
                                \\192.168.25.71\xraysScanners
                                192.168.25.21:/vol/vol2/meta6*
                                \\192.168.25.49\equipment
                                \\192.168.25.49\for_lease
                                \\192.168.25.48\equipBkup
                                \\192.168.25.48\leasedBkup


\\192.168.25.15\MP3S         medarcv:/rcrds

                                \\192.168.25.20\histories
                                \\192.168.25.27\bulkstorage
```

```
                                \\192.168.25.29\prescriptions
                                192.168.25.21:/vol/vol2/meta3*


\\192.168.25.15\xraysScannersmedarcv:/lab_equipment

                                \\192.168.25.27\backlot_records
                                \\192.168.25.71\xraysScanners
                                192.168.25.21:/vol/vol2/meta6*
                                \\192.168.25.49\equipment
                                \\192.168.25.49\for_lease
                                \\192.168.25.48\equipBkup
                                \\192.168.25.48\leasedBkup


\\192.168.25.15\Y2004        medarcv:/rcrds

                                \\192.168.25.20\histories
                                \\192.168.25.27\bulkstorage
                                \\192.168.25.29\prescriptions
                                192.168.25.21:/vol/vol2/meta3*


\\192.168.25.15\Y2005        medarcv:/rcrds

                                \\192.168.25.20\histories
                                \\192.168.25.27\bulkstorage
                                \\192.168.25.29\prescriptions
                                192.168.25.21:/vol/vol2/meta3*


\\192.168.25.15\Y2008        medarcv:/rcrds

                                \\192.168.25.20\histories
                                \\192.168.25.27\bulkstorage
                                \\192.168.25.29\prescriptions
                                192.168.25.21:/vol/vol2/meta3*


\\192.168.25.15\Y2010        medarcv:/rcrds

                                \\192.168.25.20\histories
                                \\192.168.25.27\bulkstorage
                                \\192.168.25.29\prescriptions
                                192.168.25.21:/vol/vol2/meta3*



Where * denotes metadata only physical server.
Where ** denotes a direct volume mapped to a namespace.
```

Figure 23.4 *Sample Output: show server-mapping status*

```
bstnA> show server-mapping status

Virtual Server
        Physical Server                                                  Status
-----------------------------------------------------------------    --------
192.168.25.12:/vol                                                    Ready

        nas1:/vol/vol2/shr                                            Online
        nas2:/vol/datavol1/direct/mtgs                                Online
```

```
        nas1:/vol/vol2/notes                                    Online
        nas2:/vol/datavol1/direct/export                        Online
        nas3:/exports/data                                      Online

192.168.25.15:/acct                                             Ready

        das1:/exports/budget                                    Online
        das8:/work1/accting                                     Online
        das3:/exports/acct2                                     Online
        das7:/lhome/it5                                         Online

192.168.25.15:/acct/wksheets                                    Ready

        das1:/exports/budget                                    Online
        das8:/work1/accting                                     Online
        das3:/exports/acct2                                     Online
        das7:/lhome/it5                                         Online

192.168.25.15:/claims                                           Ready

        nas1:/vol/vol2/insurance                                Online
        nasE1:/patient_records                                  Online

\\192.168.25.15\CLAIMS                                          Ready

        \\nas1\insurance                                        Online
        \\nasE1\patient_records                                 Online

\\192.168.25.15\SPECS                                           Ready

        \\nas1\insurance                                        Online
        \\nasE1\patient_records                                 Online

\\192.168.25.15\STATS                                           Ready

        \\nas1\insurance                                        Online
        \\nasE1\patient_records                                 Online

\\192.168.25.15\labs                                            Ready

        \\nas10\equipment                                       Online
        \\nas10\for_lease                                       Online
        \\fs2\backlot_records                                   Online
        \\fs5\xraysScanners                                     Online
        \\nas11\equipBkup                                       Unknown
        \\nas11\leasedBkup                                      Unknown

\\192.168.25.15\xraysScanners                                   Ready

        \\nas10\equipment                                       Online
        \\nas10\for_lease                                       Online
        \\fs2\backlot_records                                   Online
        \\fs5\xraysScanners                                     Online
        \\nas11\equipBkup                                       Unknown
        \\nas11\leasedBkup                                      Unknown

\\192.168.25.15\ARCHIVES                                        Ready

        \\fs4\prescriptions                                     Online
        \\fs1\histories                                         Online
        \\fs2\bulkstorage                                       Online
```

```
\\192.168.25.15\bulkstorage                                    Ready

        \\fs4\prescriptions                                    Online
        \\fs1\histories                                        Online
        \\fs2\bulkstorage                                      Online

\\192.168.25.15\Y2004                                          Ready

        \\fs4\prescriptions                                    Online
        \\fs1\histories                                        Online
        \\fs2\bulkstorage                                      Online

\\192.168.25.15\Y2005                                          Ready

        \\fs4\prescriptions                                    Online
        \\fs1\histories                                        Online
        \\fs2\bulkstorage                                      Online

\\192.168.25.15\Y2008                                          Ready

        \\fs4\prescriptions                                    Online
        \\fs1\histories                                        Online
        \\fs2\bulkstorage                                      Online

\\192.168.25.15\Y2010                                          Ready

        \\fs4\prescriptions                                    Online
        \\fs1\histories                                        Online
        \\fs2\bulkstorage                                      Online

\\192.168.25.15\MP3S                                           Ready

        \\fs4\prescriptions                                    Online
        \\fs1\histories                                        Online
        \\fs2\bulkstorage                                      Online

\\192.168.25.15\CELEBS                                         Ready

        \\fs4\prescriptions                                    Online
        \\fs1\histories                                        Online
        \\fs2\bulkstorage                                      Online

\\192.168.25.15\chem_results                                   Ready

        ** medarcv:/rcrds/2005                                 Online
        \\fs1\chem_results/.                                   Online
        \\fs3\hematology_results/.                             Online

\\192.168.25.16\MICE                                           Ready

        \\fs3\mice_lab_results                                 Online
        \\fs4\white_mice                                       Online

\\192.168.25.16\RATS                                           Ready

        \\fs6\labRats                                          Online
        \\fs7\lrTests                                          Online
```

virtual server

Purpose

A *virtual server* is the ARX that hosts a global server. Each virtual server has a virtual IP (*VIP*) address where the global server's front-end services listen for clients. Use the virtual server command to create a virtual server for the current global server.

Use the no form of the command to delete a virtual server from the current global server.

Mode

gbl-gs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
virtual server switch-name vip mask [vlan vlan-id]
  [cluster cluster-name]

no virtual server switch-name vip [cluster cluster-name]
```

switch-name (1-128 characters) is host name for the ARX.

vip is the IP address for the ARX to use when it runs front-end services for the current global server.

mask reveals the subnet-part of the VIP. All clients for this VIP should reside on this subnet or be reachable through this subnet via a static route (ip route).

vlan *vlan-id* (optional; 1-65535) identifies the VLAN that carries the above subnet.

cluster-name (optional, 1-64 characters) is only relevant if the ARX is part of a disaster-recovery (DR) configuration. In a DR configuration, there is an active ARX cluster with one set of filers and a backup cluster with a mirrored set of filers. This determines which cluster uses the *vip* and *mask*. Run the virtual server command twice per global server if you use DR: once to designate the *vip* and *mask* at the active cluster, and again to determine the *vip* and *mask* at the backup cluster. Use show cluster for a list of configured clusters. If you omit this, the CLI applies the change to the local cluster.

Default(s)

vlan-id - 1

Guidelines

This command places you in gbl-gs-vs mode. From there, you can establish the server's presence in your Windows network (if necessary), and start the virtual server. If there is a local WINS server, you can use the wins command to identify it. The default NetBIOS name for the virtual server is derived from the global server's FQDN (for example, "\\MYSERVER"); you can use wins-name to change the NetBIOS name. The wins-alias command registers one or more additional names for the server. To start the server, use the enable (gbl-gs, gbl-gs-vs) command.

Clients can access the global server's front-end service(s) through the VIP configured in this command. They must be on the same subnet established by the *vip* and *mask*, or they must be reachable through a static route (see ip route) where the next-hop gateway is in that subnet.

The CLI prompts you for confirmation if you use no virtual server to delete a virtual server; enter **yes** to proceed.

Guidelines: Clusters

The CLI keeps the Windows names and aliases consistent for both clusters. Whether you enter gbl-gs-vs mode for cluster A or cluster B, there is only one set of wins, wins-name, or wins-alias parameters that are shared with both clusters. Only the *vip* and *mask* are unique to each cluster.

Guidelines: External DNS Config

Clients can access the same front-end services through the global server's FQDN (for example, www.mybusiness.com) if you use a naming service (such as DNS) to map the global server's FQDN to the VIP address. Set up an alias on your external name server.

Guidelines: DNS Config Required for Kerberos

For virtual servers that offer CIFS service and use Kerberos authentication, the above naming-service configuration is required. Kerberos does not allow the use of IP addresses for connecting with the virtual IP. To simplify maintenance, you can configure dynamic DNS as part of the CIFS service: use name-server to identify the DNS server for each domain in the Active-Directory (AD) forest, then use dynamic-dns to map a host name to the VIP. Finally, use active-directory alias to register the same name as a valid SPN for this server. With this configuration, the CIFS service automatically updates the local DNS servers with any configuration changes.

Samples

```
bstnA(gbl-gs[www.wwmed.com])# virtual server bstnA 192.168.25.10
255.255.255.0 vlan 25
bstnA(gbl-gs-vs[www.wwmed.com~192.168.25.10])#
```

adds the current ARX, "bstnA," as a virtual server for the current global server. The switch listens at VIP 192.168.25.10; it offers all of the global server's front-end services at that IP address. Each front-end service listens, by default, at its well-known TCP and/or UDP port on the VIP.

```
provA(gbl-gs[provmed.medarch.org])# virtual server newptA 192.168.8.145
255.255.255.0 vlan 80 cluster newport

% INFO: Cluster 'newport' is not the local cluster. This configuration
option will not take effect on this switch.

provA(gbl-gs-vs[provmed.medarch.org])#
```

adds a backup ARX, "newptA," as a virtual server for the "newport" cluster. The switch listens at VIP 192.168.8.145; it offers all of the global server's front-end services at that IP address.

```
bstnA(gbl-gs[www.wwmed.com])# no virtual server bstnA 192.168.25.14
Delete virtual server ''192.168.25.14'' on switch ''bstnA''? [yes/no]
yes
```

removes the current switch from the current global server.

Related Commands

ip route
wins
wins-name
wins-alias
active-directory alias
enable (gbl-gs, gbl-gs-vs)
windows-domain (gbl-gs)
show global server

wait-for vip-disable

Purpose When you use no enable to disable a virtual server, its Virtual IP (VIP) is disabled in the background while you continue to use the CLI. Use the wait-for vip-disable command to wait until the VIP is disabled.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`wait-for vip-disable`** ***`vip`*** **`[timeout`** ***`timeout`*** **`]`**

vip (ipV4: 0.0.0.0. to 255.255.255.255) is the VIP.

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - 0 (zero, meaning that the command should wait indefinitely)

Guidelines This command is useful in CLI scripts.

You can use this command after you use no enable (gbl-gs, gbl-gs-vs) to stop a virtual server. The CLI blocks until the virtual server is offline and unavailable to clients.

If you set a ***timeout*** and it expires before the VIP is offline, the command exits with a warning. To interrupt the wait-for vip-disable command, press **<Ctrl-C>**.

Sample

```
bstnA(gbl-gs-vs[www.wwmed.com~192.168.25.10])# wait-for vip-disable
192.168.25.10
```

waits until this VIP is offline.

Related Commands enable (gbl-gs, gbl-gs-vs)
wait-for vip-enable

wait-for vip-enable

Purpose Use the wait-for vip-enable command to wait until the Virtual IP (VIP) is ready.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`wait-for vip-enable`** ***`vip`*** **`[timeout`** ***`timeout`*****`]`**

vip (0.0.0.0. to 255.255.255.255) is the VIP.

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - 0 (zero, meaning that the command should wait indefinitely)

Guidelines This command is useful in CLI scripts.

You can use this command after you use enable (gbl-gs, gbl-gs-vs) to start a virtual server. The CLI blocks until the VIP is available on the network.

This does not indicate that the virtual server behind the VIP is fully online. enabling all of the software in the virtual service and its backing namespace(s) can take several additional seconds.

If you set a ***timeout*** and it expires before the VIP is available, the command exits with a warning. To interrupt the wait-for vip-enable command, press **<Ctrl-C>**.

Sample

```
bstnA(gbl-gs-vs[www.wwmed.com~192.168.25.10])# wait-for vip-enable
192.168.25.10
```

waits until this VIP is available.

Related Commands enable (gbl-gs, gbl-gs-vs)
wait-for vip-disable

windows-domain (gbl-gs)

Purpose

For global servers running CIFS as a front-end service (see cifs), use the windows-domain command to set the global server's Windows domain.

Use no windows-domain to remove any Windows domain from the global server.

Mode

gbl-gs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
windows-domain domain [pre-win2k-name short-name]
no windows-domain
```

domain (1-64 characters) is the domain name (for example, MYCOMPANY.COM).

short-name (optional, 1-15 characters) is the short, pre-Windows 2000 name for the domain. This is the domain's NetBIOS name. The default (below) should suffice for most situations, and this option should be unnecessary.

Default(s)

short-name - The old-style name discovered with active-directory update seed-domain. If the old-style name was never discovered for this domain, the ARX uses the first part (before the first "."), up to 15 characters) of the FQDN in the ***domain-name***.

Guidelines

You must set the Windows domain for global servers that offer CIFS as a front-end service.

The default for the short-domain name (sometimes called an "NT domain name") is sufficient for most installations. This name is automatically discovered with the active-directory update seed-domain operation. If that operation did not run or discovered no short name for this domain, this uses the first component in the long domain name, converted to uppercase: for example, if the domain is "group.bigco.com" the short domain name is "GROUP." You can use the pre-win2k-name option for rare cases where the short name is *not* one of these options (for example, the short name is "TEAM" instead of "GROUP", and the "TEAM" name is not set in the AD).

If ***short-name*** is completely different from the ***domain*** name, and the cifs service does not use constrained delegation (see domain-join), each backing namespace requires an additional ntlm-auth-server. The second ntlm-auth-server object is a copy of the one for the long-domain name ("group.bigco.com"), where the windows-domain (gbl-ntlm-auth-srv) is the short-domain name ("TEAM") and all other parameters are the same. Use the ntlm-auth-server (gbl-ns) command to associate the additional ntlm-auth-server object to the backing namespace.

Samples

```
bstnA(gbl-gs[insur.medarch.org])# windows-domain MEDARCH.ORG
```

sets the Windows domain to MEDARCH.ORG for the current global server.

```
bstnA(gbl-gs[www.test.com])# no windows-domain
```

removes the Windows domain from the global server at www.test.com.

Related Commands

global server
export (gbl-cifs)

wins

Purpose A WINS server resolves domain-based names (such as www.mycompany.com) with IP addresses (such as 192.168.95.91). This is analogous to Domain Name Service (DNS), and is sometimes used instead of DNS in Windows networks. The virtual server registers its NetBIOS names with the WINS server when the service is enabled.

Use the wins command to identify the local WINS server.

Mode gbl-gs-vs

Security Role(s) storage-engineer or crypto-officer

Syntax **`wins`** ***`ip-address`***

ip-address is the IP address of the WINS server (for example, 192.168.70.65). This address must be on the proxy-IP subnet (established with the ip proxy-address command) or reachable through a gateway on that subnet (via static route: see ip route to create a static route).

Default(s) None

Guidelines The virtual server registers its NetBIOS names with the WINS server when you use enable (gbl-gs, gbl-gs-vs) to enable the virtual server. The registered NetBIOS names are the *computer* name and the *group* name. The default computer name is the first component of the global server's FQDN (for example, "\\FS1" is the default for a global server at "fs1.mycompany.com"). You can use the wins-name command to reset this computer name and wins-alias to add additional NetBIOS computer names. The NetBIOS group name is the pre-win2k-name from the windows-domain (gbl-gs) command (for example, "MYCOMPANY").

To set the character encoding for these name registrations, you can use cifs *fqdn* to reach gbl-cifs mode, then use wins-name-encoding.

Sample

```
bstnA(gbl-gs-vs[ac1.medarch.org~192.168.25.15])# wins 192.168.25.20
```

sets the WINS server to 192.168.25.20 for the current virtual server.

Related Commands

global server **->** virtual server
wins-name
wins-alias
cifs
wins-name-encoding

wins-alias

Purpose

You can use the wins-name command to set the NetBIOS computer name for the virtual server, and you can use the wins-alias command to add an additional NetBIOS alias.

Use no wins-alias to remove one or all NetBIOS aliases.

Mode

gbl-gs-vs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
wins-alias alias
no wins-alias [alias]
```

alias (1-15 bytes; see below) is a NetBIOS alias you choose for the current virtual server. The alias must start with a letter, but can contain numbers and underscores (_) after the first character.

If you omit the ***alias*** from the no form of the command, the CLI deletes all aliases.

Default(s)

None

Guidelines

This only applies to virtual servers that support CIFS front-end services.

CIFS clients can connect to the virtual server using either the wins-name or any aliases you set with this command. You can issue the command multiple times to create multiple NetBIOS aliases.

When you run enable (gbl-gs, gbl-gs-vs) for the virtual server, the virtual server registers its computer name and all aliases with the WINS server.

The WINS server may use multi-byte character encoding to allow registration of NetBIOS names with Japanese, Korean, or other non-Latin characters. At the cifs service for this global server, use the wins-name-encoding command to set the character encoding for this registration.

For Latin-1 characters, a NetBIOS alias can be 15 characters (one character per byte); for non-latin, multi-byte characters, you are limited to a smaller number of characters.

Guidelines: Kerberos Support

If any of the CIFS front-end services support Kerberos authentication, additional configuration is required in the Active Directory for these WINS aliases. The domain-join command registers its CIFS-service FQDN at the Active-Directory (AD) domain; if a client connects to the service using a WINS alias, the local Domain Controller (DC) must be able to translate the WINS alias into the registered FQDN. Otherwise, Kerberos authentication fails.

You can use the active-directory alias command to set a SPN for the CIFS service's global server.

Samples

```
bstnA(gbl-gs-vs[fs.nt.org~192.168.25.12])# wins-alias HR_SERVER
```

sets a NetBIOS computer name of “HR_SERVER” for the current virtual server. A Windows “Network Neighborhood” will display “\\HR_SERVER” for the current virtual server.

```
bstnA(gbl-gs-vs[fs.nt.org~192.168.25.12])# no wins-alias HUMAN_RES
```

removes one NetBIOS alias.

Related Commands

global server `->` virtual server
wins-name
cifs
wins-name-encoding

wins-name

Purpose

The default NetBIOS computer name for a virtual server is the first component of the global server's FQDN (for example, "\\FS1" is the default for a global server at "fs1.mycompany.com"). Use the wins-name command to set another NetBIOS computer name.

Use no wins-name to revert to the default NetBIOS name.

Mode

gbl-gs-vs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
wins-name name
no wins-name
```

name (1-15 bytes; see below) is the NetBIOS name you choose for the current virtual server. The name must start with a letter, but can contain numbers and underscores (_) after the first character.

Default(s)

The default NetBIOS name is the first component of the global server's FQDN (for example, "\\MYSERVER" is the default for a global server at "myserver.myorg.org"). If this is longer than 15 bytes, the switch only uses the first 15 bytes.

Guidelines

This only applies to virtual servers that support CIFS front-end services.

The virtual server registers its NetBIOS names with the WINS server when you enable (gbl-gs, gbl-gs-vs) the virtual server. The registered NetBIOS names include both the computer name and the domain name.

The WINS server may use multi-byte character encoding to allow registration of NetBIOS names with Japanese, Korean, or other non-Latin characters. At the cifs service for this global server, use the wins-name-encoding command to set the character encoding for this registration.

For Latin characters, a NetBIOS name can be 15 characters (one character per byte); for non-latin, multi-byte characters, you are limited to a smaller number of characters. The limit is 15 bytes, not simply 15 characters.

To add additional NetBIOS names to this virtual server, use the wins-alias command.

Guidelines: Kerberos Support

If any of the CIFS front-end services support Kerberos, additional configuration is required in the Active Directory for this WINS name. The domain-join command registers its CIFS-service FQDN at the Active-Directory (AD) domain; if a client connects to the service using a different WINS name, the local Domain Controller (DC) must be able to translate the WINS name into the registered FQDN. Otherwise, Kerberos authentication fails.

You can use the active-directory alias command to set a SPN for the CIFS service's global server.

Sample

```
bstnA(gbl-gs-vs[fs.nt.org~192.168.25.14])# wins-name PERSONNEL
```

sets the NetBIOS computer name to "PERSONNEL" for the current virtual server. A Windows "Network Neighborhood" displays "\\PERSONNEL" for the current virtual server.

Related Commands global server -> virtual server
cifs
wins-name-encoding

24

Front-End Services

browsing

Purpose

An MMC user can share out storage from a remote CIFS filer, and MMC offers an opportunity to browse the filer's volumes so that the user can choose one of them for sharing. The CIFS service does not support this browsing feature by default. Use the browsing command to enable this type of MMC browsing for one namespace behind the current CIFS service.

Use no browsing to disallow this MMC browsing for one namespace (or all namespaces) behind the CIFS service.

Mode

gbl-cifs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
browsing namespace
no browsing [namespace]
```

namespace (1-30 characters) chooses a namespace behind this CIFS service. This makes the namespace's volumes visible to MMC clients who "browse" the namespace's volumes before choosing one for sharing. (In a multi-protocol namespace, the backing namespace might be established by an NFS export for the same global server.)

Default(s)

`no browsing`

Guidelines

After you enable browsing, authorized Windows clients can use MMC browsing to share any volume in the chosen namespace. If you want MMC clients to use this feature for multiple namespaces behind this CIFS service, repeat this command for each desired namespace.

You authorize a client by adding it to a Windows-management-authorization (*WMA*) group; use the windows-mgmt-auth command to create one such group. The show windows-mgmt-auth command shows all members and permissions for a group. Use the windows-mgmt-auth (gbl-ns) command to assign a WMA group to the backing namespace(s); you can use the command more than once to associate multiple groups with the namespace. All namespaces behind this CIFS service must have a matching set of WMA groups.

Authorized WMA users can share the service's backing volumes with MMC instead of the export (gbl-cifs) command. With browsing enabled, MMC clients have an added convenience: they can browse all of the namespace's volumes and share the one(s) they want.

Guidelines: CIFS-Share Names for Backing Namespaces

From the Windows-management interface, management clients see all of a namespace's managed volumes as directories under a single CIFS share. The share names are "C$," "D$," and so on. There is one such share for each namespace. If the namespace is exported by multiple CIFS services, it uses that same share name for all of them.

Each direct volume appears as a separate root, "Z$," "Y$," and so forth. These share names are also the same for every CIFS service that exports the namespace.

The ARX software assigns these share names to a namespace and its direct volume(s) when you issue this command. The share-name assignments persist as long as browsing is enabled for the namespace from any CIFS service.

There is a maximum of 26 such share names. Each namespace with browsing enabled uses one share name for all of its managed volumes plus one for each direct volume. You cannot enable browsing for a namespace if it contains enough volumes to exceed 26 share names. If you disable browsing for a given namespace in all CIFS services, its share names are available for use by other namespaces.

Samples

```
bstnA(gbl-cifs[ac1.medarch.org])# browsing medarcv
```

enables MMC browsing for the "medarcv" namespace behind the "ac1.medarch.org" service. This is a convenience for using MMC to share out medarcv's volumes.

```
bstnA(gbl-cifs[testSrvr])# no browsing demo2
```

disables MMC browsing for a namespace behind the "testSrvr" service.

Related Commands

export (gbl-cifs)
windows-mgmt-auth
windows-mgmt-auth (gbl-ns)
show windows-mgmt-auth

cifs

Purpose

Use the cifs command to instantiate a CIFS service for a global server.

Use the no form of the command to delete a CIFS-service instance.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
cifs fqdn
no cifs fqdn
```

fqdn (1-128 characters) is the fully-qualified domain name (for example, "myserver.organization.org") for an existing global server. Use show global server to see a list of global servers.

Default(s)

None

Guidelines

This command prompts for confirmation before creating the CIFS service; enter **yes** to continue. (To prevent confirmations for creating these sorts of configuration objects, use the terminal expert command.)

After you confirm, the CLI places you in gbl-cifs mode. From gbl-cifs mode, you must use the export (gbl-cifs) command to share at least one managed volume through CIFS. Use domain-join to register as a valid CIFS service with a Windows Domain Controller (DC). This creates a machine account in the Active Directory for the *fqdn* service, or uses an existing machine account if one already exists for *fqdn*. At a DC, an authorized administrator must raise the "Trusted for Delegation" flag for this service, and should also list the filers to which it can delegate; from priv-exec mode, you can use probe delegate-to to find all filers behind the CIFS service.

You then use the enable (gbl-cifs, gbl-nfs) command to start the CIFS service. Finally, map the service's *fqdn* to its VIP with the dynamic-dns command. Windows clients can then map their network drives to the volumes exported by the CIFS service. Each exported volume appears on the Windows network as one share on the CIFS service's *fqdn*. To add additional service-principal names for this CIFS service, to which clients could also connect, use the active-directory alias command from gbl-gs-vs mode.

Dynamic DNS names for a CIFS service are registered at external DCs, and all must be removed (with no dynamic-dns) before you can remove the CIFS service (with no cifs).

Guidelines: SMB2

If all the filers behind a CIFS service support SMB2, and if SMB2 is enabled on the ARX, the service can offer SMB2 to its clients.

Use the smb2-allowed command to enable SMB2 for the entire system. To disable SMB2 on a particular CIFS service, you can use the no smb2-allowed (gbl-cifs) command.

The show cifs-service *fqdn* command shows if SMB2 is possible for the *fqdn* service. If it is not possible, you can use the show cifs-service filer-capabilities *fqdn* command to find the specific problem (for example, it might be that one of the filers behind the CIFS service is not capable of supporting SMB2).

Guidelines: NetBIOS Name(s)

The NetBIOS name for the CIFS service is tied to the virtual server. Each virtual server registers a NetBIOS name with its local WINS server. From gbl-gs-vs mode, use the wins command to identify the WINS server for a virtual server. From the same mode, you can optionally use wins-name to set the NetBIOS name for the virtual server and wins-alias to set additional aliases; by default, the NetBIOS name is the first part of the global server's FQDN (for example, "\\ATHOS" for the global server at "athos.myco.com").

Guidelines: Maximum Front-End Services

The ARX-500 supports a maximum of 16 front-end services, and the other platforms support up to 64. This limit applies to the sum of all CIFS services (created by this command) and NFS services (created with the nfs command).

Guidelines: Show Commands

There are several CLI commands for monitoring and managing client access to CIFS shares. The show cifs-service exports command shows the connection statistics for each CIFS share/export. Use show cifs-service user-sessions to find clients with open CIFS sessions. To disconnect a client session, use drop cifs-service user-session. The show cifs-service open-files command shows all files that are currently open through CIFS (and locked against reads). The policy engine cannot migrate or copy a file in this state; use close cifs file to close one from the CLI.

For CIFS-authentication statistics, including numbers for client authentications with the CIFS service as well as ARX authentications at back-end filers, use the show statistics cifs authentication command. For NTLM-authentication statistics, use show statistics domain-controller.

Samples

```
bstnA(gbl)# cifs www.mycompany.com
This will create a new CIFS service.

Create CIFS service 'www.mycompany.com'? [yes/no] yes
bstnA(gbl-cifs[www.mycompany.com])#
```

instantiates CIFS service at www.mycompany.com.

```
bstnA(gbl)# no cifs www.defunctcompany.com
Delete CIFS service on 'www.defunctcompany.com'? [yes/no] yes
bstnA(gbl)#
```

removes the CIFS-service instance that was running at www.defunctcompany.com.

Related Commands

export (gbl-cifs)
domain-join
probe delegate-to
enable (gbl-cifs, gbl-nfs)
dynamic-dns
smb2-allowed
smb2-allowed (gbl-cifs)
global server `->` virtual server `->` active-directory alias
global server `->` virtual server `->` wins

Optional, descriptive commands:

global server `->` virtual server `->` wins-name
description (gbl-cifs)
show cifs-service exports
show cifs-service user-sessions
drop cifs-service user-session
show cifs-service open-files
close cifs file
show statistics cifs authentication
show statistics domain-controller
show cifs-service filer-capabilities

cifs rekey

Purpose Some Kerberos DCs have an option to assign a "maximum age" for their machine-account passwords. Each machine uses its secret password to authenticate with the AD domain. (An ARX-CIFS service is treated as a machine in this context.) If the secret password expires, the CIFS service cannot use Kerberos for client authentication. You can use the cifs rekey all command to refresh the secret password (or key) for every CIFS service on the ARX.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `cifs rekey all`

all is required. If you rekey any CIFS service on the ARX, you rekey all of them at once.

Default(s) None

Guidelines If the ARX has a redundant peer, you can only perform this command on the active peer.

Most sites do not set expiration times for these machine-account keys. This command is only required if the local AD policy sets a maximum age for these keys. This only affects CIFS services that use Kerberos and have joined a domain (with the domain-join command).

Note

This causes a brief outage for every Kerberos-supporting CIFS service as it resets its secret key and restarts. The outage is brief enough that it should not be perceptible to any ARX clients.

A CLI prompt warns you of the service outage; enter **yes** to continue with the rekey operation. If you proceed, the CLI lists all of the affected CIFS services.

Run this command before any CIFS service's password expires. If a CIFS service's password has already expired, use domain-join to rejoin the AD domain.

Sample

```
bstnA# cifs rekey all

WARNING: Confirming this operation will cause the ARX to reset the
machine account password for all Kerberos-enabled CIFS services on the
switch. This will cause an outage of 30-60 seconds for all CIFS/NFS
services on the switch.

Proceed [yes/no] yes
  The machine account password for CIFS service 'ac1.medarch.org'
successfully changed.
bstnA# ...
```

resets the machine-account keys for all CIFS services on "bstnA."

Related Commands at
domain-join

clear dynamic-dns

Purpose A CIFS service can use dynamic DNS to register its host name and/or aliases in DNS. When you remove a host name or change the VIP, the CIFS service attempts to remove the DNS entry. If the remove fails after several retries, the show dynamic-dns output keeps a record of the failed removal indefinitely, to remind you to remove the "A" record from the DNS server itself. Use the clear dynamic-dns command to clear all failed DNS removals from the system.

Mode priv-exec

Security Role(s) network-engineer, storage-engineer, or crypto-officer

Syntax `clear dynamic-dns`

Default(s) None

Guidelines The CLI prompts for confirmation before clearing all failed DNS-remove records from the ARX. Enter **yes** to proceed.

The dynamic-dns command causes a CIFS service to add or remove DNS-host names from a local DNS server. The CIFS service retries every minute on failure; it retries add operations indefinitely, but eventually stops retrying for remove operations. These remove operations show up with "Failed" status in the show dynamic-dns output. "Failed" removals persist in the show output until someone clears them. Use the clear dynamic-dns command to clear all failed DNS removals from the system.

Sample

```
bstnA# clear dynamic-dns
Clear failed dynamic DNS entries? [yes/no] yes
bstnA#
```

clears all failed DNS-removal entries from the system.

Related Commands dynamic-dns
show dynamic-dns

clear nlm locks

Purpose The NFS Lock Manager (NLM) manages file locks for NFS clients. Use the clear nlm locks command to clear all locks for all clients, or for a given client machine.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `clear nlm locks {all | host-name}`

all | *hostname* is a required choice:

all clears all locks for all NFS clients, system-wide.

hostname (1-192 characters) clears all locks held by one client's host machine (for example, "myhost").

Default(s) None

Guidelines An NFS client can lock a file or a region within a file, so that other NLM-compliant users cannot access the region until the lock is released. The ARX's NLM server manages all the locks for all NFS front-end services and volumes.

Important

Never clear the locks for a running client. If the client still has locks, it has no way of knowing that those locks have been invalidated: the NLM protocol does not support a "clear all locks" message to the client. This could result in two different clients having an exclusive lock on the same file region, which typically corrupts the file.

To see all the client machines that hold locks for a given file, use the show nlm file command. For a list locks held by a given client machine (or a list of all the client machines that have locks), use the show nlm client command. Use show nlm statistics to see statistics for all NLM activity.

To disable NLM for a given NFS service, disable NFS first (with enable (gbl-cifs, gbl-nfs)) and then use no nlm enable.

Sample

```
bstnA# clear nlm locks client18
```

clears all NFS locks held by the client machine, "client18."

Related Commands

show nlm file
show nlm client
show nlm statistics
nlm enable

description (gbl-cifs)

Purpose Use the optional description command to set a descriptive string that appears as a comment in the Windows "Network Neighborhood."

Use the no form of the command to delete the description.

Mode gbl-cifs

Security Role(s) storage-engineer or crypto-officer

Syntax

```
description text
no description
```

text (1-48 characters) is your description. Surround the text with quotation marks ("") if it contains any spaces.

Default(s) no description

Guidelines The server description appears as a comment for all browser views of the Windows network. It describes this CIFS service. This description also appears in the output for show cifs-service.

Sample

```
bstnA(gbl-cifs[www.medarcv.com])# description "archives for Longwood Ave"
```

specifies a description for the current CIFS service.

Related Commands cifs
show cifs-service

description (gbl-nfs)

Purpose

Use the optional description command to set a descriptive string for the NFS service. This appears in the show command.

Use the no form of the command to delete the description.

Mode

gbl-nfs

Security Role(s)

storage-engineer or crypto-officer

Syntax

`description` ***`text`***
`no description`

text (1-255 characters) is your description. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s)

no description

Guidelines

The description appears in the output for show nfs-service.

Samples

```
bstnA(gbl-nfs[www.wwmed.com])# description “mount point for wwmed storage”
```

specifies a description for the current NFS service.

Related Commands

nfs
show nfs-service

domain-join

Purpose Use the domain join command to register a CIFS front-end service with its Active-Directory (AD) domain. This is similar to adding client computers to the Active Directory domain, and it requires a sufficiently-privileged username and password. An administrator with stronger credentials can pre-create an account at the DC before you run this command, or the command can automatically enable delegation privileges at the DC.

Mode gbl-cifs

Security Role(s) storage-engineer or crypto-officer

Syntax

```
domain-join domain-name [ou "organizational-unit"]
  [delegation {constrained | unconstrained | none}]
  [timeout seconds]
```

domain-name (1-256 characters) is the Windows domain to join. This domain must be the global server's domain, and it must be part of the Active-Directory forest configuration (see active-directory update seed-domain or active-directory-forest). The CIFS service registers with one of the DCs that manage this domain.

organizational-unit (optional, 1-512 characters) is the *Organizational Unit* (OU) to which this service should be added. A Windows domain is divided into OUs for administrative convenience: each OU is a group of machines or accounts that are managed by a single administrator. Surround the OU name with quotation marks (""); OU names often contain spaces. Use a slash (/) to separate each layer in a nested OU: for example, ou "Virtual Services/pharmaceuticals" is a valid nested OU. This is the reverse of the LDAP path shown on the Windows server, which would be ".../cn=ac1,**ou=pharmaceuticals,ou=Virtual Services**,dc=medarch,dc=org."

delegation {constrained | unconstrained | none} (optional) is unnecessary if the CIFS-service account already exists in the local AD. Omit this option if an AD administrator pre-created the machine account for this service. Otherwise, the ARX automatically generates the account on a DC that manages this domain. This option determines which "Trust for delegation" flag (if any) to be used for the CIFS service. The command enables this delegation at the DC if (and only if) you have sufficient privileges to do so:

> **constrained** is recommended, though it is only possible if the AD forest has a "Domain Functional Level" of Windows 2003 or later. This sets the "Trust ... for delegation for specified services" flag for the CIFS service. You (or an authorized administrator) can then access the DC and specify all of the back-end filers to whom this service can delegate. Once all the service's filers are on its "delegate-to" list at the DC, clients can authenticate once and then access CIFS storage on all those filers. (You can use the probe delegate-to command to confirm that all of the filers behind the CIFS service are properly configured at the DC.) The DCs allow clients to use NTLM, NTLMv2, or Kerberos to authenticate; at the namespace(s) behind the CIFS service, you can use cifs authentication to support any or all of these authentication types.

Syntax (Cont.)

unconstrained only works for Kerberos authentication (by default). This sets the "Trust ... for delegation to any service" flag for this CIFS service. Clients cannot connect to any back-end storage with NTLM or NTLMv2 if this flag is raised at the DCs. To support NTLM authentication together with this option, you must install a Secure Agent at one of your DCs; see ntlm-auth-server.

none does not trust this CIFS service for any delegation. This makes it impossible for Kerberos clients to authenticate to this service. (NTLM/NTLMv2 clients cannot authenticate either, unless you install a Secure Agent on your DCs and use the ntlm-auth-server command.) Use this when you do not have sufficiently strong Windows credentials to raise any "Trust for delegation" flag. This presumes that an authorized administrator will raise the appropriate "Trust for delegation" flag at a later time, thereby opening access to CIFS clients.

seconds (optional, 1-300) sets a time limit for the domain-join operation. Use this in case the DC is too slow to respond before the command times out.

Default(s)

organizational-unit, if omitted, defaults to "Computers" at the domain's root.

delegation, if omitted, is **unconstrained**.

timeout *seconds* is 90 seconds

Guidelines: Before You Join

To support CIFS-client authentication, the namespace(s) behind this CIFS service must have the desired authentication types enabled. Use the gbl-ns cifs authentication command to enable Kerberos, NTLM, and/or NTLMv2 for a namespace.

Additionally, the Active-Directory forest must be accurately reflected in the switch configuration (see active-directory update seed-domain or active-directory-forest), and the global server's FQDN must be in one of the forest's domains (the forest-root, a tree-domain, or a child-domain). You can use show active-directory to see all of the domains discovered in the AD forest.

Guidelines: Minimal Windows Privileges Needed for a Domain-Join

After you issue this command, the CLI prompts for a username and password. These are the credentials to be presented to the DC, called the *domain-joiner* credentials. In some sites, the domain-joiner credentials are sufficiently privileged to create a machine account on the DC and raise all of the proper flags. In other sites, a more-privileged user creates the machine account before the domain-joiner runs this command. These next subsections describe the minimal access privileges required to join the domain with either a new machine account or a pre-created account.

New Machine Account

If the *domain-joiner* has the following privileges, the domain-join operation can create a new machine account for the CIFS service. A pre-created machine account is unnecessary.

- In the "Computers" OU for the domain (or the OU chosen by this command), the user must have permission to "Create Computer Objects."

  On a Windows Server 2003 DC, you can set this in the "Active Directory Users and Computers" interface.

- The user must be able to "Enable computer and user accounts to be trusted for delegation."

  On a Windows Server 2003 DC, you can find this under "Domain Controller Security Policy."

  This is only required if you use some form of delegation, either constrained or unconstrained. It is unnecessary if you specify **delegation none**.

Guidelines: Minimal Windows Privileges Needed for a Domain-Join (Cont.)

Pre-Created Machine Account

If someone in another group pre-creates the machine account at the DC, that person must have the minimal privileges above. On a Windows Server 2003 DC, that person creates a new machine account in the "Active Directory Users and Computers" interface. The machine account requires the following minimal settings:

- Set the machine account's "Delegation" tab as desired.

  We recommend "Trust this computer for delegation to specified services only" together with "Use any authentication protocol." This is what the domain-join operation sets when it creates the account with the **delegation constrained** options.
- Add the domain-joiner's user account into the machine account's "Security" tab. Give the following permissions to that user:
  - Read
  - Write
  - Reset Password

Do not use any **delegation** options with a pre-created machine account. The delegation is already set at the DC, as described above.

Guidelines: Switching a Joined Service to Constrained Delegation

If the CIFS service is already joined to the domain and you want to upgrade to constrained delegation, you can upgrade at the DC. From the DC, use the "Active Directory Users and Computers" interface to access the machine account for *fqdn*. (The machine account name is the first part of the *fqdn*: for example, the machine account name for "ac1.medarch.org" is "ac1".) In the machine account's properties, access the machine account's "Delegation" tab and enter

- "Trust this computer for delegation to specified services only,"
- "Use any authentication protocol," and
- the "delegate to" list at the bottom, which requires all the filers behind this CIFS service. All of the filers must be joined to the same Windows Domain as this CIFS service. See below if you are not aware of all the filers behind this service.

The CIFS service probes the DCs every few minutes to find these settings, so it may not be up-to-date right after you change them. You can use the sync cifs delegation command to update the CIFS service immediately. To verify success, wait for the DCs to synchronize with one another and then use the show cifs-service *fqdn* command to confirm the change. You can also use the probe delegate-to *fqdn* command to probe the DC and confirm that all the necessary filers are on the "delegate to" list. If any of these filers show up as "failed" on that list, enter those filers at the DC.

Note

When you enter your filers at a Windows 2003 or 2008 DC, the following error may appear in a pop-up:

```
The following Active Directory error occurred: The
specified directory service attribute or value
already exists.
```

This indicates a software issue at the DC. Click OK *to return to the Delegation tab, then check the* Expanded *box at the bottom. Look for duplicate SPNs in this expanded view; each filer should have only 2 SPNs, one with its FQDN and another with only its hostname. Remove all duplicate SPNs and then click* Apply.

Guidelines: Switching a Joined Service to Constrained Delegation (Cont.)

Once this is done for all of your cifs services, you can remove all ARX Secure Agent applications from your DCs. When constrained delegation is active in all CIFS services, the Secure Agent applications are ignored. For more information about Secure Agent, see the ARX® Secure Agent Installation Guide.

Guidelines: Verifying the Domain Join

You can use the show cifs-service *fqdn* command to verify the success of the domain-join operation. This also shows the delegation type (constrained, unconstrained, or none) configured at the DC. If the service uses constrained delegation, this also lists the filers to which the service can delegate. The CIFS service probes the DCs every few minutes to update this information, so it may not be up-to-date right after you run the domain-join command.

The probe delegate-to *fqdn* command verifies that all of the filers behind the *fqdn* service are properly listed at the DC. This verifies that the "delegate to" list is correct and complete on the DC. If not, copy all failed filers from the ARX's probe delegate-to output to the DC's "delegate to" list.

◆ Note

As described in the note above, an error may appear in a pop-up when you enter your filers at a Windows 2003 or 2008 DC. If the error appears, remove any duplicate filer SPNs (visible in the Expanded *view) as described in the above note.*

Guidelines: Name Length

The CIFS service's machine-account name is its FQDN by default, unless the hostname in the FQDN ("ac1" in "ac1.medarch.com") is too long. If the hostname exceeds 15 bytes and the CIFS service uses the recommended constrained-delegation feature, DCs will reject the name for NTLM or NTLMv2 authentications. Therefore, the CIFS service uses only the first 15 bytes for the machine-account name. If the name is too long, a prompt informs you of the truncated account name that was registered at the DCs.

The shorter account name appears in the output of show cifs-service.

Guideline: Coexistence with NTLM

If you use the delegate constrained option and properly configured the CIFS service at the DC, the service can support NTLM and Kerberos simultaneously. Each client negotiates the authentication protocol at the start of the CIFS session. The CIFS service uses Kerberos to authenticate with its back-end filers.

If you use the delegate unconstrained option, NTLM and NTLMv2 authentication is not supported by default. In this case, you can install a Secure Agent on some of your DCs to allow NTLM and/or NTLMv2 authentication. Then use ntlm-auth-server to identify each of these DCs, and ntlm-auth-server (gbl-ns) to use them in the namespace(s) behind this CIFS service. (The show cifs-service *fqdn* command shows the namespace(s) behind the *fqdn* service.) In this case, if a client connects with NTLM or NTLMv2, the CIFS service uses one of the NTLM protocols (preferably NTLMv2) to connect to the back-end filers.

All policy-initiated transactions with CIFS filers (such as file migrations) attempt to use Kerberos. The backing namespace negotiates the authentication method (Kerberos, NTLMv2, or NTLM) with each back-end filer. Kerberos requires that the backing namespace uses a proxy user configured with a full FQDN for its windows-domain (gbl-proxy-user). Kerberos also requires that it can discover the service-principal name (SPN) of each filer; you can use the show exports command to find if the filer allows for SPN discovery, or use spn to set the SPN manually.

Guidelines: Resetting Machine-Account Passwords

The domain-join operation involves the exchange of a shared password (or *key*). It is possible to set a "maximum age" for this type of machine-account key at the DC. If a maximum age is set at your site, you can later use the cifs rekey command to reset the CIFS service's key before it expires. This makes the CIFS service automatically restart, so it causes a brief service outage.

You can also use the at command to reset the machine-account key on a regular schedule.

Samples

```
bstnA(gbl-cifs[ac1.medarch.org])# enable
bstnA(gbl-cifs[ac1.medarch.org])# domain-join MEDARCH.ORG
Username: jsmith
Password: jspasswd

% INFO: Service 'ac1' successfully joined the domain using a
pre-created computer account.


% INFO: Service 'ac1' joined the domain with delegation type set to
Constrained. Only certain selected services are allowed to be delegated
on 'ac1'. The Active Directory administrator must configure the
services 'ac1' may delegate to, which are the filer(s) this CIFS
service virtualizes.

bstnA(gbl-cifs[ac1.medarch.org])#
```

joins the current CIFS service to its Windows domain using the credentials for jsmith. This uses a pre-created machine account; before issuing this CLI command, someone created a machine account for the CIFS service on a DC.

```
stkbrgA(gbl-cifs[bgh.medarch.org])# domain-join MEDARCH.ORG
Username: Administrator
Password: adminpasswd

% INFO: Service 'bgh' successfully joined the domain. The new computer
account was created.


% INFO: Service 'bgh' joined the domain with delegation type set to
Unconstrained. The delegation of services on 'bgh' is allowed only with
the Kerberos protocol.
```

joins another CIFS service to its Windows domain using the credentials for Administrator. This creates a new machine account with unconstrained delegation. Unconstrained delegation does not support NTLM or NTLMv2 clients unless you also install a Secure Agent (see ntlm-auth-server) for the domain.

Related Commands

cifs
active-directory-forest
show cifs-service
sync cifs delegation
probe delegate-to
show global server
cifs authentication **kerberos**
cifs rekey

dynamic-dns

Purpose

Kerberos authentication requires up-to-date DNS records for every front-end CIFS service. Use the dynamic-dns command to register a host name or DNS alias for the current CIFS service.

Use the no form of the command to withdraw an alias from the local DNS database.

Mode

gbl-cifs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
dynamic-dns alias
no dynamic-dns alias
```

alias (1-255 characters) is an alias for this CIFS service. This can be a simple host name (for example, "myserver") or an FQDN (for example, "myserver.mycompany.com"). If it is a simple host name, the CIFS service appends the service's Windows domain (set with windows-domain (gbl-gs)) before registering the name. If it is an FQDN, it must match that Windows domain or the CLI rejects it. Use show global server to see the Windows domain for the CIFS service (the Windows domain is set at the global server for the service).

Default(s)

None

Guidelines

Before you use this command, you must configure at least one DNS server for this service's Windows domain. You can do this as part of setting up the Active-Directory (AD) forest; use the gbl-forest name-server command to identify the name server for the service's domain. The Windows domain for this CIFS service is set in the service's global server. Use show global server to see the Windows domain. You must also join an Active-Directory domain with domain-join. Finally, the CIFS service must be enabled (with enable (gbl-cifs, gbl-nfs)) before you issue this command.

This command immediately registers an address (A) record with one of the configured name servers. The "A" record maps the service's virtual-IP (VIP) address to the alias you set in this command. (Use the virtual server command to set the VIP.) You can repeat this command to map multiple aliases to the VIP, so that clients can connect through multiple service names. All clients that connect through these names can use Kerberos authentication.

The CIFS service renews the "A" record registration at 1:00 AM (local time) each morning. It adds and or withdraws "A" records whenever the global server's VIP changes (with [no] virtual server) or the CIFS service is removed (with no cifs). Use the priv-exec dynamic-dns update command to update all "A" records immediately.

The no dynamic-dns command causes the CIFS service to remove the A record from DNS. As with the affirmative syntax, you must issue the command while the CIFS service is enabled.

The show dynamic-dns command shows the dynamic-DNS configuration as well as the current status of all host-name registrations.

Basic DNS is defined in RFCs 1034 and 1035. RFC 3645 defines the Microsoft-specific authentication extensions for dynamic DNS. The ARX implementation adheres to all three RFCs.

The CIFS services only register "A" records, not "PTR" records for reverse lookups.

Guidelines: Alias Support

If the CIFS service has more valid hostnames than its FQDN, each additional hostname is a *DNS alias*. Each DNS alias must be established externally, in the Active Directory. The domain-join command registers its CIFS-service FQDN at the Active-Directory (AD) domain; if a client connects to the service using a DNS alias, the local Domain Controller (DC) must be able to translate the alias into the registered FQDN. Otherwise, Kerberos authentication fails.

From gbl-gs-vs (virtual server) mode, you can use the active-directory alias command to set a DNS alias for the CIFS service's global server. This registers the alias in the external Active Directory. Follow these instructions to reach gbl-gs-vs mode for the correct virtual server:

- Go back to gbl mode.
- Use the global server *fqdn*, command, where *fqdn* is the one that identifies the CIFS service, too.
- Use the virtual server command to access the single virtual server under the global server.

Samples

```
bstnA(gbl-cifs[sales.myco.com])# dynamic-dns sales
bstnA(gbl-cifs[sales.myco.com])# dynamic-dns fs5.myco.com
```

registers two names for the current CIFS service.

```
bstnA(gbl-cifs[fs.myorg.org])# no dynamic-dns betafs
```

removes an alias, "betafs," from the current CIFS service.

Related Commands

cifs
name-server
domain-join
active-directory-forest
active-directory alias
dynamic-dns update
show dynamic-dns

dynamic-dns update

Purpose Kerberos authentication requires up-to-date DNS records for every front-end CIFS service. If you use dynamic-dns for the current CIFS service, the CIFS service automatically updates its DNS server daily, or whenever a configuration change necessitates a change in DNS configuration. If you require an immediate update to your DNS server(s), use the dynamic-dns update command.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer or crypto-officer

Syntax `dynamic-dns update [`***`fqdn`***`]`

fqdn (optional, 1-128 characters) identifies a single CIFS service for the update. If you omit this, all CIFS services send DNS updates to their local name servers.

Default(s) Perform the update for all CIFS services

Guidelines If you send DNS updates from all CIFS services, the CLI prompts for confirmation. If you have a large number of CIFS services, each with multiple host names to register, this can be a network-intensive operation. Enter **yes** to proceed.

You can register a new host name or DNS alias with the dynamic-dns command. This command sends an immediate update for all such host names. Under most circumstances, the update is redundant; the CIFS service sends updates daily, and sends changes whenever a VIP changes, a host name is removed, or some other IP/name change occurs. You can use this command to force an update if, for example, your external DNS servers have returned from an extended outage.

The show dynamic-dns command shows the dynamic-DNS configuration as well as the current status of all host-name registrations.

This command only affects CIFS services that are enabled (enable (gbl-cifs, gbl-nfs)).

Samples

```
bstnA# dynamic-dns update as1.medarch.org
```

sends DNS updates for the “as1.medarch.org”service. One “A” record is registered for each alias.

```
bstnA# dynamic-dns update
Update DNS host names for ALL CIFS services? [yes/no] yes
bstnA# ...
```

sends DNS updates for all CIFS services on the ARX.

Related Commands

dynamic-dns
show dynamic-dns

enable (gbl-cifs, gbl-nfs)

Purpose

Use the enable command to activate the current CIFS or NFS service.

Use the no form of the command to stop the current CIFS or NFS service.

Modes

gbl-cifs

gbl-nfs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
enable
no enable
```

Default(s)

Disabled

Guidelines

The enable command makes all declared CIFS shares or NFS exports available to clients. Clients can access the share through the virtual-IP (VIP) address of the global server's virtual server.

The CIFS service does not send any dynamic-dns updates unless it is enabled.

Samples

```
bstnA(gbl-cifs[ac1.medarch.org])# enable
```

enables the current CIFS service.

```
bstnA(gbl-nfs[www.myorg.org])# enable
```

enables the current NFS service.

```
bstnA(gbl-nfs[www.wwmed.com])# no enable
```

disables the current NFS service.

Related Commands

cifs
nfs

export (gbl-cifs)

Purpose

Use the export command to share a namespace's managed volume through CIFS.

Use the no form of the command to stop sharing a managed volume.

Mode

gbl-cifs

Security Role(s)

storage-engineer or crypto-officer

Syntax: Standard CIFS Share

```
export namespace vol-path [as share-name] [description description]
no export namespace vol-path
no export as share-name
```

namespace (1-30 characters) can be any namespace that supports CIFS. If the same CIFS service exports more than one namespace, all of the namespaces behind the service must have matching CIFS settings; see *Guidelines: Exporting from Multiple Namespaces*, below.

vol-path (1-1024 characters) is the path to one of the namespace's managed volumes (for example, "/multimedia") or inside the volume ("/multimedia/data/apps").

as *share-name* (optional; 1-1024 characters) sets an advertised path to the volume. If entered, CIFS clients see this path as an available export instead of the ***volume*** path. You can re-enter the command to share the volume under more than one share name.

description *description* (optional; 1-64 characters) sets a description for this share. This appears as a comment in Windows network browsers such as "Network Neighborhood" and the **net view** command. Surround the text with quotation marks ("") if it contains any spaces.

Syntax: Filer Subshare

```
export namespace vol-path as subshare-name filer-subshare
[hidden] [description description]
no export namespace vol-path
no export as subshare-name
```

is the syntax for exporting a filer subshare. A filer *subshare* is a CIFS share inside the volume's directory tree, below the root. At the back-end filer(s), a subshare often has a different share-level ACL than the root share.

namespace and ***vol-path*** are described above. The ***vol-path*** must be a valid path in a managed volume with filer-subshares enabled.

as *subshare-name* (optional; 1-1024 characters) is required for filer subshares. This is the name of the subshare for ***vol-path*** (for example, "MYSHARE"). If the CIFS service is enabled, this subshare is replicated on all back-end shares behind ***vol-path***. This replicates the subshare's directory path, ACL, and attributes from the master directory to all other back-end shares. You can use the find command to find a directory's master directory. This replication operation is equivalent to running sync subshares from-service.

filer-subshare is also required for filer subshares.

Syntax: Filer Subshare (Cont.)

hidden (optional) tells the CIFS service to find a hidden version of the ***subshare-name*** behind the volume, and expose it to front-end clients. The hidden subshare has a dollar-sign at the end of its name (such as "hidden_share$"). With this option, the CIFS service exposes the subshare by exporting it without the "$" suffix.

description *description* (optional; 1-64 characters) sets a description for this subshare, as explained above.

Default(s)

share-name defaults to the volume path without the leading slash (/), all in uppercase. For example, the "/var" volume would default to the share name, "VAR."

subshare-name has no default.

description defaults to an empty string.

Guidelines

This shares a managed volume or one of its subshares from any CIFS namespace. Clients can access this share through CIFS by accessing the share-name from their Windows Network Neighborhood.

You can enter this command multiple times to share multiple managed volumes or subshares. You can also share one volume under multiple share names; for subshares, however, the share name must match the one used at the back-end filers. The ARX supports a maximum of 16000 CIFS exports.

You must enable the CIFS service for any shares to be visible to clients. Use the enable (gbl-cifs, gbl-nfs) command in gbl-cifs mode to enable CIFS service.

Guidelines: Configuration at a Local DC

A new CIFS front-end share often requires some configuration at a local Domain Controller (DC). Best practices dictate that a CIFS service is trusted to delegate CIFS connections to all of the filers behind it, and it is not trusted to delegate CIFS to any other filer. This is called *constrained delegation*. The DCs manage the CIFS service's ability to delegate, as well as the back-end filers to whom it can delegate. An authorized Windows administrator must therefore go to a DC and add this new volume's filers to its CIFS service's "delegate to" list.

After you export the volume, you can use probe delegate-to *fqdn* to check the DC(s) and determine whether or not delegation is properly configured for the new filers. If not, a properly authorized administrator must access the DC and add the new filer(s) to the CIFS service's "delegate to" list.

Every new filer must be joined to the same Windows Domain as the CIFS service, and it must support Kerberos authentication. The show cifs-service *fqdn* command shows the domain to which the *fqdn* service is joined.

Guidelines: Subshares

Filer subshares are defined at back-end filers, along with all of their share-level ACLs. If you support subshares in the CIFS service, a client that connects to the subshare on the front end passes through to the corresponding subshare at the back-end; the share-level ACL is enforced there. Before using the filer-subshares syntax, you must prepare the volume with the filer-subshares command. Without the subshare preparation at both the CIFS service and the volume, a client always connects to the root of the back-end share, even after connecting to a lower-level share at the front-end. You can use the show cifs-service subshares command to view all filer subshares corresponding to each front-end subshare.

As an alternative to export ... filer-subshares, you can use sync subshares from-namespace to export all back-end subshares at once.

Guidelines: Exporting from Multiple Namespaces

If you share out volumes from more than one namespace, the following configuration options must match in all namespaces:

- same set of cifs authentication protocols,
- same setting for cifs anonymous-access ,
- same ntlm-auth-server (gbl-ns) or servers,
- same ntlm-auth-db (gbl-ns),
- same set of windows-mgmt-auth (gbl-ns) (WMA) groups, and
- same sam-reference filer for all the namespaces behind the service. This SAM-reference filer must define all local groups for all filers behind all namespaces. (The CIFS service does not strictly enforce this restriction.)

When you use the export command to share a volume from a new namespace, the ARX software confirms that the new namespace is consistent with all currently-exported namespaces. If any of the above options do not match, the CLI offers the opportunity to synchronize the namespace settings. Enter **yes** to proceed. This provides all backing namespaces with the same superset of options; for example, if one namespace supported Kerberos and the other supported Kerberos and NTLM, the synchronization process would configure all namespaces to have both Kerberos and NTLM.

Guidelines: Unsupported Windows File-System Features

Microsoft Windows supports several file-system features that are not supported by an ARX CIFS export. Included among these unsupported features are:

- Encryption
- Disk Quotas
- Mount Points
- Hard Links
- Distributed Link Tracking
- Change Journals
- System ACLs (SACLs), used for auditing and access alarms - an ARX CIFS service does not copy a SACL when migrating a file or striping a directory.
- OS/2 Extended Attributes (EAs)
- Object ids
- Reparse points

Samples

```
bstnA(gbl-cifs[www.medarcv.com])# export medarcv /rcrds description
“clerical records”
```

shares the “/rcrds” volume of the medarcv namespace.

```
bstnA(gbl-cifs[www.myco.com])# export homes /lhome as HOMEDIRS
bstnA(gbl-cifs[www.myco.com])# export homes /lhome as ALLSHARES
```

shares the “/lhome” volume under two different share names.

```
bstnA(gbl-cifs[www.myco.com])# no export as ALLSHARES
```

removes one of the shares for the “/lhome” volume. The other share name (HOMEDIRS, above) is still accessible to clients.

```
bstnA(gbl-cifs[www.medarcv.com])# no export medarcv /schedules
```

stops sharing the “/schedules” volume of the medarcv namespace.

```
bstnA(gbl-cifs[www.medarcv.com])# export medarcv /rcrds/2004 as Y2004
filer-subshare
```

exports a subshare behind the “/rcrds” volume, “Y2004.”

```
bstnA(gbl-cifs[www.medarcv.com])# export medarcv /rcrds/VIP_wing as
CELEBS filer-subshare hidden
```

exports a hidden subshare behind the “/rcrds” volume, “CELEBS$,” but exports it as “CELEBS” so that front-end clients can see it.

Related Commands

cifs
enable (gbl-cifs, gbl-nfs)
filer-subshares
sync subshares from-namespace
sync subshares from-service
show cifs-service subshares

export (gbl-nfs)

Purpose

Use the export command to export a namespace volume through NFS.

Use the no form of the command to stop exporting a volume.

Mode

gbl-nfs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
export namespace vol-path [as export-path] [access-list name]
no export namespace vol-path
no export as export-path
```

namespace (1-30 characters) can be any namespace that supports NFS. If the same NFS service exports more than one namespace, all of the namespaces behind the service must have the same character-encoding nfs settings.

vol-path (1-1024 characters) is the path to one of the namespace's managed volumes (for example, "/var") or inside the volume ("/var/log/archive").

as *export-path* (optional, 1-1024 characters) sets an advertised path to the NFS-share mount point. If entered, NFS clients see this path as an available export or share instead of the ***volume*** path. You cannot export the same *vol-path* under more than one name; you can only use this option to change the export path that is visible to clients.

access-list *name* (optional, 1-64 characters) is the access list to associate with this share.Use the show nfs-access-list command to view the configured access lists.

Default(s)

as *export-path*, if omitted, defaults to the name of the ***volume***.

access-list *name*, if omitted, defaults to allowing all clients access to the volume/share with read/write permission, root-squash enabled, and root squashed to UID 65534 and GID 65534.

Guidelines

This command exports a volume from any namespace. Clients can mount this volume through NFS by accessing the virtual-IP (VIP) address of the global server's virtual server.

You can enter this command multiple times to export multiple volumes.

You must enable the NFS service for any shares to be visible to clients. Use the enable (gbl-cifs, gbl-nfs) command in gbl-nfs mode to enable NFS service.

If you export volumes from more than one namespace, the character-encoding nfs settings must match in all namespaces:

Before you remove a volume with no export, you must first disable it using the no enable command in gbl-nfs mode (see enable (gbl-cifs, gbl-nfs)).

Guidelines: Behavior When a Back-End Filer Goes Offline

If any filer goes offline behind this export, a client cannot access its files or directories. The NFS service could return an access error to the client (typically preferred by human users), or the service could let the client retry until the filer answers (often preferred by scripted clients). You can use the offline-behavior command to choose the behavior of this export when a back-end filer goes offline. You can define a default behavior for the entire NFS service, and you can redefine the behavior for each export.

Sample

```
bstnA(gbl-nfs[ac1.medarch.org])# export wwmed /acct
```

exports the "/acct" volume of the wwmed namespace.

Related Commands

nfs
enable (gbl-cifs, gbl-nfs)
offline-behavior
show nfs-access-list

export offline-access

Purpose

When a Windows client uses *offline access* for a file on a remote CIFS share, Windows creates a local copy of the file for the client. The client can use that local copy whenever the client machine is disconnected from the CIFS share, and can later *sync* the local copy with the original whenever the CIFS-share connection is up.

By default, clients can manually select directories and files for offline access. You can use the export offline-access command to also automatically enable offline access for any file the client opens (with or without network optimization), or to disable all offline access.

Use the no form of the command to revert to the default setting.

Mode

gbl-cifs

Security Role(s)

storage-engineer or crypto-officer

Syntax

`export offline-access` ***`share-name`*** **`{manual|auto|auto-local|none}`**
`no export offline-access` ***`share-name`***

share-name (1-1024 characters) is the name of a share or subshare from this CIFS service, established with the export (gbl-cifs) command. You can use show cifs-service *fqdn* for a list of all shares and subshares in the *fqdn* service.

manual | auto | auto-local | none is a required choice. This maps to the offline settings in the MMC interface for managing CIFS shares. Use this command in place of those options:

manual means that clients must set the option for offline files and directories themselves. This maps to the MMC option (in Windows 7), "Only the files and programs that users specify are offline."

auto indicates that client software is allowed to automatically enable offline access for every file the client opens. Client software is also allowed to automatically sync its local files with the files on the CIFS share whenever the connection is up. This can produce a great deal of network traffic between clients, the ARX, and the back-end filers behind this CIFS share. It maps to the Windows 7 MMC option, "All files and programs that users open from the shared folder are automatically available offline."

auto-local extends the **auto** feature. This maps to the same MMC option above together with an additional "Optimized for performance" option. This permits the client software to reduce its network usage by downloading all executable files immediately after connecting. For newer Windows clients, this is equivalent to the **auto** feature.

none means that the client software is not permitted to offer any offline-access options to the client. This maps to the MMC option, "No files or programs from the shared folder are available offline."

no export offline-access ***share-name*** is equivalent to **export offline-access** ***share-name*** **manual**.

Default(s)

`manual`

Guidelines

The default is sufficient for most sites. Windows servers use the same default behavior for offline access.

Both of the **auto** options have the potential for creating high network traffic, since they ensure that every opened file is locally cached, and they sync local files whenever possible. The CLI prints an informational message about this if you choose either of those options. Use these options only on the advice of F5 Support.

Samples

```
stkbrgA(gbl-cifs[bgh.medarch.org])# export offline-access naumkeag
auto-local
% INFO: Setting offline-access to automatic may result in increased
traffic loads through the ARX, if users choose to cache a very large
number of files.
```

makes the "naumkeag" share provide automatic offline access in every directory, but also makes the client software optimize its network usage.

```
bstnA(gbl-cifs[ac1.medarcv.com])# export offline-access ARCHIVES none
```

disables offline access for the "ARCHIVES" share.

Related Commands

cifs
export (gbl-cifs)
show cifs-service

nfs

Purpose

Use the nfs command to instantiate an NFS service for a global server.

Use the no form of the command to delete an instance of NFS service.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

nfs *fqdn*
no nfs *fqdn*

fqdn (1-128 characters) is the fully-qualified domain name (for example, "www.organization.org") for an existing global server. Use show global server to see a list of global servers.

Default(s)

None

Guidelines

This command prompts for confirmation before creating the NFS service; enter **yes** to continue. (To prevent confirmations for creating these sorts of configuration objects, use the terminal expert command.)

This places you in gbl-nfs mode. You must use the export (gbl-nfs) command to export at least one ARX volume through NFS. Optionally (while the service is disabled), you can use no nlm enable to disable the NFS Lock Manager (NLM) service. Once you have at least one export, you must use the enable (gbl-cifs, gbl-nfs) command to start the NFS service. Clients can then mount the export(s) through the global server's virtual-IP (VIP) address (created with the virtual server command).

If any filer goes offline behind this service, a client cannot access its files or directories. The service could return an access error to the client (typically preferred by human users), or the service could let the client retry until the filer answers (often preferred by scripted clients). You can use the offline-behavior command to choose the behavior of this service when a back-end filer goes offline. You can also use the command to customize the behavior for various NFS exports.

The ARX-500 supports a maximum of 16 front-end services, and the other platforms support up to 64. This limit applies to the sum of all NFS services (created by this command) and CIFS services (created with the cifs command).

You can use show nfs-service to show the configuration of an NFS service. The show nfs-service mounts command shows all clients who are mounted to an NFS service.

Samples

```
bstnA(gbl)# nfs www.mycompany.com
This will create a new NFS service.

Create NFS service 'www.mycompany.com'? [yes/no] yes
bstnA(gbl-nfs[www.mycompany.com])#
```

instantiates NFS service at www.mycompany.com.

```
bstnA(gbl)# no nfs www.defunctcompany.com
Delete NFS service on 'www.defunctcompany.com'? [yes/no] yes
```

removes the NFS-service instance that was running at www.defunctcompany.com.

Related Commands

export (gbl-nfs)
no nlm enable
enable (gbl-cifs, gbl-nfs)
show global server
global server **->** virtual server
nfs-access-list
show nfs-service
show nfs-service mounts

nfs tcp timeout

Purpose

If an NFS/TCP connection to a back-end filer times out, the NFS service disconnects from the client. Use this command to reset the timeout and change this client-side behavior; return an NFS I/O error instead of disconnecting.

Use the no form of this command to go back to the default timeout and behavior.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

`nfs tcp timeout` ***`seconds`***
`no nfs tcp timeout`

seconds (5-104) is the new timeout period.

Default(s)

105-second timeout period

disconnect the client on a back-end-filer timeout

Guidelines

If a non-default timeout is set, the NFS service returns an NFSERR_IO (NFSv2) or NFS3ERR_IO (NFSv3) to its client. If the default timeout is sent, NFS service disconnects from clients, with the expectation that they will retry.

This command sets the timeout and behavior for *all* NFS services and filers.

Use the show nfs tcp command to show the current timeout period and behavior.

Samples

```
bstnA(gbl)# nfs tcp timeout 90
```

sets the timeout period to 90 seconds on the current switch. If an NFS/TCP connection times out on the back end, the NFS service returns an NFS I/O error to the client instead of disconnecting.

```
bstnA(gbl)# no nfs tcp timeout
```

reverts to the default timeout and behavior.

Related Commands

show nfs tcp

nlm enable

Purpose

The NFS Lock Manager (NLM) manages file locks for NFS clients. While the NFS service is disabled, you can use the no nlm enable command to disable NLM.

Use the affirmative command, nlm enable, to enable NLM for this NFS service.

Mode

gbl-nfs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no nlm enable
nlm enable
```

Default(s)

```
nlm enable
```

Guidelines

The current NFS service must be disabled (through no enable (gbl-cifs, gbl-nfs)) for this command to function.

The CLI prompts for confirmation before stopping NLM; enter **yes** to proceed.

An NFS-client application can lock a file or a region within a file, so that other NLM-compliant applications cannot access the region until the lock is released. You can disable NLM on a service-by-service basis.

It is often appropriate to disable NLM for a direct volume, where clients are allowed to access back-end files without going through the ARX. Each volume manages its own set of NLM locks, unknown to any of the filers behind the volume. If any client accesses a file by going directly to the filer, the filer is unaware of any NLM locks that the volume may have granted. This is also true if a client accesses the same file through another direct volume. In either case, the volume behind this NFS service is incapable of enforcing its NLM locks.

To support NLM locking for direct volumes, you must be sure that clients exclusively use the volume to access all of the direct shares behind it.

If multiple NFS services export the same volume, consistently enable or disable NLM for all of them. This applies to managed volumes as well as direct volumes.

To see all the client machines that hold locks for a given file, use the show nlm file command. For a list of locks held by a given client machine (or a list of all the client machines that have locks), use the show nlm client command. If a client machine is down, you can clear all of its locks with clear nlm locks. Use show nlm statistics to see statistics for all NLM activity.

Samples

```
bstnA(gbl-nfs[acopiaFiler])# no nlm enable
Disable Network Lock Manager on acopiaFiler? [yes/no] yes
```

disables NLM for the NFS service, “acopiaFiler.”

```
bstnA(gbl-nfs[myserver.com])# nlm enable
```

enables NLM for the NFS service, “myserver.com.”

Related Commands

nfs
enable (gbl-cifs, gbl-nfs)
show nlm file
show nlm client
clear nlm locks
show nlm statistics

offline-behavior

Purpose When an NFS client attempts to access a file on an offline filer, the ARX-NFS service ignores the request and lets the client retry; this can block the NFS client for a long time. Use the offline-behavior ... deny-access command to immediately return an access error (typically NFSERR_ACCES or NFS3ERR_ACCES) instead. This provides a definitive answer to NFS-client software, so that it can return an error for an offline file and immediately move on to other file requests.

To silently wait for a client retry instead of returning an access error, use the offline-behavior ... retry command. This setting prevents sparse returns (a mixture of available files and access errors) for multi-file requests.

Use the no offline-behavior command to revert to the inherited offline behavior for the entire NFS service, one namespace behind the service, one volume behind it, or one particular export from one of its volumes.

Mode gbl-nfs

Security Role(s) storage-engineer or crypto-officer

Syntax

```
offline-behavior [namespace ns [volume /vol [path expt-path]]]
  { retry | deny-access }
no offline-behavior [namespace ns [volume /vol
  [path expt-path]]]
```

ns (optional, 1-30 characters) narrows the scope of this command to a particular namespace behind this NFS service.

/vol (optional if you choose namespace, 1-1024 characters) narrows the scope further, to a particular volume in the above namespace.

expt-path (optional, if you choose a volume, 1-64 characters) narrows the scope to one explicitly-exported path in the volume. You must choose an export path that is explicitly exported with the export (gbl-nfs) command. The command is accepted, but the NFS service does not honor the offline behavior until someone exports this exact path.

retry | deny-access is a required choice:

> **retry** causes the service to accept requests to offline filers without sending an error back to the NFS client. The client must therefore retry until the filer comes back online, or until the client software times out. This causes NFS requests to hang if they involve a mix of online and offline filers.
>
> **deny-access** causes the service to return an error (typically NFSERR_ACCES or NFS3ERR_ACCES) for any client request that requires an offline filer. An NFS request succeeds (with some access errors) if it accesses both online and offline filers.

no offline-behavior reverts to the behavior inherited from the parent path, volume, namespace, or NFS service. Inheritance is explained below.

Default(s) `no offline-behavior`

At the NFS-service level, this is equivalent to **offline-behavior retry**.

Guidelines

If one filer fails behind an NFS service on the ARX, only a subset of the service's files are unavailable. Clients may prefer to work with the subset of files instead of waiting until all files are available again. By default, the ARX service drops all requests to offline filers: soft-mounted clients retransmit the request until they time out, and hard-mounted clients retransmit indefinitely. You can use this command to make the ARX service return errors to its clients instead of allowing them to retry.

For multi-file requests, this behavior can result in sparse returns (a mixture of access errors and available files).

For example, suppose an NFS client mounts an ARX export with one offline filer and runs the Unix **ls -l** command. If the export is configured to deny-access for offline files, the user would see access errors in place of the offline files (two of them, in this case), along with full information for all of the accessible files:

```
juser@client2:/mnt/arx$ ls -l
ls: process.html: Permission denied
ls: time.html: Permission denied
total 160
drwxrwxrwx  11 juser engineering  4096 2009-08-31 10:52 acctProcs
-rwxrwxrwx   1 muser medteam      9927 2006-04-11 22:19 ctrlProc.html
drwxrwxrwx   2 muser medteam      4096 2006-03-22 10:07 images
-rwxrwxrwx   1 xuser srnurses     2058 2006-03-26 21:07 index.html
-rwxrwxrwx   1 juser engineering  1743 2006-03-22 10:17 intro.html
drwxrwxrwx   3 muser medteam      4096 2009-08-24 16:01 payable
-rwxrwxrwx   1 muser medteam     20326 2005-05-06 09:31
timeSheetForm.xsl
juser@client2:/mnt/arx$ ...
```

For an ARX export that waits for the client to retry, the above command would hang until the back-end filer came back online or the **ls** command timed out.

Guidelines: Retry vs. Deny-Access Behavior

For volumes where an NFS client desires access to some filers whether or not other filers are available, we recommend the deny-access behavior. This behavior always allows NFS access to filers that are online, even for directories that are partially hosted by an offline filer. This can be useful in tiered volumes where files on the lower-tier filers are rarely accessed.

Guidelines: Offline-Behavior Inheritance

The best practice is to choose a general rule for the NFS service (either clients want access to a subset of files when a filer is offline, or they would prefer no access if any filer is offline), and apply the opposite rule to specific NFS exports where you want the opposite behavior. An NFS export inherits its offline behavior as follows:

- NFS-service
- namespace
- volume
- export path
- export sub path

For example, suppose a service hosts a software repository, and is mostly used by scripts that build the software into executable files. If any files are missing from the build, the build would fail. The default for this service should therefore be

```
offline-behavior retry
```

to support the build scripts. If any filer is offline, the build should wait for it to come back online before it continues.

Suppose that one of its exports was used by engineers for collecting the executable output files and copying them to test environments; that particular export could override the service's offline behavior. The engineers would likely prefer to see all the available files even when some are unavailable. If the export was "/usr/local/bin" in the "eng~/builds" volume, you could use

```
offline-behavior namespace eng volume /builds path /usr/local/bin
deny-access
```

in the same NFS service. The second command changes the offline behavior for the /usr/local/bin export only.

For a sub-export that is used by other scripts, you could set the offline behavior at that export level. For example,

```
offline-behavior namespace eng volume /builds path
/usr/local/bin/auto-drop retry
```

sets the retry behavior for an export under /usr/local/bin.

Guidelines: Offline Behavior Applies to Exact Export Paths

The offline behavior setting is only useful if there is an export (gbl-nfs) for the exact path named in the offline-behavior command. The above examples would be ineffective until someone also issued both of the following export commands:

```
export eng /builds/usr/local/bin
export eng /builds/usr/local/bin/auto-drop
```

If /usr/local/bin was exported but not /usr/local/bin/auto-drop/, an NFS client that mounts /usr/local/bin/auto-drop would get the offline behavior for /usr/local/bin, "deny access." You can use the show nfs-service ... detailed command to confirm that all of your offline behavior settings are assigned to at least one NFS export.

Samples

```
bstnA(gbl-nfs[ac1.medarch.org])# offline-behavior deny-access
```

causes the “ac1.medarch.org” service to return an access error to NFS clients who access an offline filer. Presumably, the service has predominantly human clients who do not require access to all files in the volume all the time.

```
bstnA(gbl-nfs[ac1.medarch.org])# offline-behavior namespace wwmed
volume /acct path /wksheets retry

% WARNING: This command did not match any existing export. It will have
no effect until an export matching '/acct/wksheets' is entered in this
NFS service.
```

overrides this default behavior for clients that connect to the /acct/wksheets export. The warning indicates that the export does not yet exist, so NFS clients that mount /acct/wksheets still get the access errors for offline filers. You can use the export (gbl-nfs) command to create the /acct/wksheets export later, and that export would then get the retry behavior.

```
bstnA(gbl-nfs[ac1.medarch.org])# offline-behavior namespace insur retry
```

sets the retry behavior for exports from the ’insur’ namespace. As above, this overrides the default set for the NFS service.

Related Commands

nfs
export (gbl-nfs)
show nfs-service

probe delegate-to

Purpose A front-end CIFS service can delegate CIFS to the filers behind it. A CIFS client can then authenticate once and access any of the service's back-end filers without re-authenticating. This single-authentication mechanism, facilitated by *constrained delegation*, requires some configuration at the local DC. After the CIFS service joins its domain using constrained delegation (with domain-join), you must go to the DC and inform it of the filers to which the CIFS service can delegate. This command lists all of the filers behind a given CIFS service, along with their current delegation status. You can use this list to confirm that all filers are correctly set at the DC.

Mode priv-exec

Security Role(s) crypto-officer or storage-engineer

Syntax **`probe delegate-to`** ***`fqdn`*** **`[user`** ***`user-name`*** **`domain`** ***`domain-name`*****`]`**

fqdn (1-128 characters) is the fully-qualified domain name for a CIFS service (for example, "www.organization.org").

user ***user-name*** **domain** ***domain-name*** (optional) is a Windows-user identity for the probe operation. The CLI uses this Windows identity to test whether or not the CIFS service can delegate to each back-end filer. If you omit this, the CLI uses the proxy-user (gbl-ns) from one of the namespaces behind the CIFS service.

user-name and

domain-name each take1-128 characters.

Guidelines This command is only relevant for a CIFS service with constrained delegation. You set up constrained delegation at a DC or with the domain-join command. A background query runs every 10 minutes to confirm that the service has constrained delegation; this command does not function until that query runs successfully at least once.

You can use this command to find all of the filers behind a given CIFS service. At the DC, you can apply each of these filers in the Delegation list for the CIFS service: select the "Trust ... for delegation for specified services" and the "Use any authentication protocol" flags, then add the filers to the list of services. The CIFS service must be allowed to delegate to all of these filers so that its clients can authenticate to them. The client-authentication mechanism can then be any combination of NTLM, NTLMv2, and Kerberos.

Once the service is properly configured at your DCs, you can use this command to confirm the configuration.

Guidelines (Cont.) The output displays the exact username and domain used for the probe, followed by a table of probe results. The table contains one sub table per filer behind the CIFS service. Each sub table contains the following rows:

External-Filer is the name of the back-end filer. You can use the show external-filer command to list all external filers on the ARX.

Filer SPN is the service-principal name (SPN) of the back-end filer. This is the SPN that the ARX discovered, or the one that was set with the spn command. The next two SPNs are both expected to be on the CIFS service's Delegation list, at the DC.

Probe SPN-1 is another service-principal name (SPN) for the back-end filer. This is the exact text of a SPN that someone must enter at the DC. After the SPN is the status of that SPN on the DC:

- OK means that the DC allows the front-end CIFS service to delegate to this SPN. If Probe SPN-1 and Probe SPN-2 both have this status, the filer is properly entered in the CIFS service's Delegation (or "delegate-to") list.
- Failed (Not in delegate-to list) means that the front-end CIFS service cannot delegate to this filer. Go to one of the DCs for the service's domain and add this filer to the Delegate list for the CIFS service. Use the full filer name from this **Probe SPN** field. Configure the CIFS service to delegate "cifs" to the filer.
- Failed (Filer spn is not configured nor discovered) indicates that the filer's service principal name (SPN) is unknown. You can use show exports to automatically discover a filer's SPN, or spn to manually set it.
- Failed (Filer not in the same domain) means that the back-end filer is not joined to the same domain as the CIFS service. The CIFS service belongs to the Windows domain of its backing namespace(s): use show cifs-service *cifs-service* to find this Windows domain for a given *cifs-service*. The filer must be joined to the same Windows Domain for constrained delegation to function.

Probe SPN-2 is the exact text of a second SPN that should be on the CIFS service's Delegation (or "delegate-to") list. As above, the SPN's status appears after the SPN. If either SPN fails, address the issue as described above.

Samples

```
bstnA# probe delegate-to ac1.medarch.org user juser domain MEDARCH.ORG
```

uses an identity of "juser@MEDARCH.ORG" to test the CIFS delegation for the "insur.medarch.org" service. See Figure 24.1 on page 24-42 for sample output. In this case, neither of the service's back-end filers are on its Delegation list, so the administrator must go to a local DC to add them.

```
stoweA# probe delegate-to patrol.MEDARCH.ORG
```

tests the delegation status for the "patrol.MEDARCH.ORG" service. This uses the credentials defined for the proxy-user in one of the service's backing namespaces; those credentials appear at the top of the output. For sample output, see Figure 24.2 on page 24-42. This service is properly configured at the DCs.

Related Commands

show cifs-service
show exports
show external-filer

Figure 24.1 *Sample Output: probe delegate-to ac1.medarch.org ...*

```
bstnA# probe delegate-to ac1.MEDARCH.ORG user juser domain MEDARCH.ORG

User used for probe:  juser@MEDARCH.ORG

  External Filer : fs1
    Filer SPN    : vm-swp2003s2-04$@MEDARCH.ORG
    Probe SPN-1  : CIFS/vm-swp2003s2-04@MEDARCH.ORG -- OK
    Probe SPN-2  : CIFS/vm-swp2003s2-04.MEDARCH.ORG@MEDARCH.ORG -- OK

  External Filer : fs2
    Filer SPN    : vm-swp2003s1-5$@MEDARCH.ORG
    Probe SPN-1  : CIFS/vm-swp2003s1-5@MEDARCH.ORG -- OK
    Probe SPN-2  : CIFS/vm-swp2003s1-5.MEDARCH.ORG@MEDARCH.ORG -- OK

  External Filer : fs3
    Filer SPN    : VM-SWP2008S-01@MEDARCH.ORG
    Probe SPN-1  : CIFS/VM-SWP2008S-01@MEDARCH.ORG -- OK
    Probe SPN-2  : CIFS/VM-SWP2008S-01.MEDARCH.ORG@MEDARCH.ORG -- OK

  External Filer : fs4
    Filer SPN    : VM-SWP2008S-02@MEDARCH.ORG
    Probe SPN-1  : CIFS/VM-SWP2008S-02@MEDARCH.ORG -- OK
    Probe SPN-2  : CIFS/VM-SWP2008S-02.MEDARCH.ORG@MEDARCH.ORG -- OK

  External Filer : fs5
    Filer SPN    : vm-pv770n-01$@MEDARCH.ORG
    Probe SPN-1  : CIFS/vm-pv770n-01@MEDARCH.ORG -- OK
    Probe SPN-2  : CIFS/vm-pv770n-01.MEDARCH.ORG@MEDARCH.ORG -- OK

  External Filer : nas1
    Filer SPN    : ntap820$@MEDARCH.ORG
    Probe SPN-1  : CIFS/ntap820@MEDARCH.ORG -- OK
    Probe SPN-2  : CIFS/ntap820.MEDARCH.ORG@MEDARCH.ORG -- OK

  External Filer : nas10
    Filer SPN    : enterprise$@MEDARCH.ORG
    Probe SPN-1  : CIFS/enterprise@MEDARCH.ORG -- OK
    Probe SPN-2  : CIFS/enterprise.MEDARCH.ORG@MEDARCH.ORG -- FAILED (Not in delegate-to list)

  External Filer : nas11
    Filer SPN    : ntap-prov$@MEDARCH.ORG
    Probe SPN-1  : CIFS/ntap-prov@MEDARCH.ORG -- OK
    Probe SPN-2  : CIFS/ntap-prov.MEDARCH.ORG@MEDARCH.ORG -- OK

  External Filer : nasE1
    Filer SPN    : engdm$@MEDARCH.ORG
    Probe SPN-1  : CIFS/engdm@MEDARCH.ORG -- OK
    Probe SPN-2  : CIFS/engdm.MEDARCH.ORG@MEDARCH.ORG -- OK
```

Figure 24.2 *Sample Output: probe delegate-to patrol.MEDARCH.ORG*

```
stoweA# probe delegate-to patrol.MEDARCH.ORG

User used for probe:  administrator@MEDARCH.ORG

  External Filer : fsvr_15
    Filer SPN    : cifsslave$@MEDARCH.ORG
    Probe SPN-1  : CIFS/cifsslave@MEDARCH.ORG -- OK
    Probe SPN-2  : CIFS/cifsslave.MEDARCH.ORG@MEDARCH.ORG -- OK

  External Filer : fsvr_16
```

```
   Filer SPN    : sytfl$@MEDARCH.ORG
   Probe SPN-1  : CIFS/sytfl@MEDARCH.ORG -- OK
   Probe SPN-2  : CIFS/sytfl.MEDARCH.ORG@MEDARCH.ORG -- OK

 External Filer : fsvr_17
   Filer SPN    : marley$@MEDARCH.ORG
   Probe SPN-1  : CIFS/marley@MEDARCH.ORG -- OK
   Probe SPN-2  : CIFS/marley.MEDARCH.ORG@MEDARCH.ORG -- OK
```

remove namespace ... volume ... exports-only

Purpose Use the remove namespace ... volume ... exports-only command to remove all of a volume's front-end exports.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`remove namespace`** ***`name`*** **`volume`** ***`volume`*** **`exports-only [timeout`** ***`seconds`*** **`] [sync]`**

name (1-30 characters) identifies the namespace.

volume (optional, 1-1024 characters) focuses on a single volume.

exports-only removes all front-end exports (created with export (gbl-nfs) and/or export (gbl-cifs)) for the specified volume. The volume and namespace configurations remain.

seconds (optional, 300-10,000) sets a time limit for each component operation in the removal process.

sync (optional) shows its progress at the command line. With this option, the CLI prompt does not return until all components have been removed.

Default(s) None.

Guidelines The CLI prompts for confirmation before removing any front-end exports; enter **yes** to proceed.

By default, this command generates a report to show all of the actions it takes to remove the volume(s), in order. The CLI shows the report name after you issue the command, and then returns. You can enter CLI commands as the namespace software removes the exports in the background. Use tail to follow the report as it is written. Use show reports *file-name* to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands. Use the sync option to send the status to the command line instead; the command does not generate a report if you use the sync option.

Use remove namespace to remove an entire namespace or volume, or remove service to remove an entire namespace and *all* configuration objects that are dedicated to the namespace (including the front-end exports). To remove a share from a volume, use remove-share migrate or remove-share nomigrate.

Sample

```
prtlndA# remove namespace insur_bkup volume /insurShdw exports-only

Remove exports for namespace 'insur_bkup', volume '/insurShdw'?
[yes/no] yes
Scheduling report: removeNs_insur_bkup_200605250803.rpt
```

removes all front-end exports for the "insur_bkup~/insurShdw" volume.

Related Commands

remove namespace
remove service
remove-share migrate
remove-share nomigrate

show cifs-service

Purpose Use the show cifs-service command to display configuration information about front-end CIFS service(s).

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show cifs-service [{ fqdn | all } [detailed]]`

fqdn (optional; 1-128 characters) is the fully-qualified domain name for one CIFS service (for example, "www.organization.org").

all (optional) shows summaries of all CIFS-service offerings.

detailed (optional) adds CIFS-share details to the output.

If you enter the command without any of these options, the output is a table of CIFS services.

Guidelines The default output is a table of all CIFS services, one service per row, with the following columns:

Service Name is the FQDN that identifies the CIFS service. The cifs command, which instantiates the CIFS service, establishes this FQDN. It is the same as the FQDN for the global server where this service runs.

Description is set by the description (gbl-cifs) command.

Guidelines: Output for Each Service

If you focus on a single CIFS service, more details appear. Several of these fields are updated by a background query to the DCs. This background query runs every 10 minutes, so you cannot immediately see the results of a domain-join operation:

Service Name is described above.

Domain Join shows whether or not the CIFS service has performed a successful domain-join operation.

- Joined to *domain-name* indicates that the CIFS service has successfully joined the given domain.
- Not Joined indicates that the domain-join operation failed, or that no-one ever ran the command for this service.

This field is updated by the DC query mentioned above. It can take up to 10 minutes after a domain-join for this output to show the results of the operation.

Account Name the name that was used for the machine account at the local DCs. This is typically the hostname part of the CIFS-service FQDN followed by a "$" symbol; for example, "insur$" could be the account name for the "insur.medarch.com" service. If that hostname exceeds 15 bytes, this is the first 15 bytes of the hostname followed by a "$" (for example, the account name for "reallylonghostname.myorg.org" would be "reallylonghostn$"). As above, this field is updated by a DC query that runs every 10 minutes.

Delegation is the machine account's setting for the "Trust for delegation" flag at the DCs. This determines whether or not the DCs trust the CIFS service to delegate CIFS access to its filers.

- Constrained, Any Protocol indicates that this CIFS service is trusted to delegate its CIFS service to a limited number of filers (listed in the next field).
- Unconstrained, Kerberos Only means that this CIFS service is trusted to delegate any CIFS-related service to any filer. Clients cannot authenticate with NTLM or NTLMv2 unless you assign an ntlm-auth-server (gbl-ns) or an ntlm-auth-db (gbl-ns) to the namespace behind the service.
- None indicates that this CIFS service is not trusted to delegate any service to any filer. CIFS clients cannot access any of the storage behind this service. To correct this issue, an administrator with sufficient credentials must access the service's computer account at a local DC and raise its "Trust for delegation" flag. The account name appears in the field above.
- Not Applicable means that the domain-join operation failed, or that no-one ever ran the domain-join command for this service.

As above, this field is updated by a DC query that runs every 10 minutes.

Delegate To shows a list of filers to whom the CIFS service is allowed to delegate. This is "Not Applicable" until someone performs a successful domain-join for the CIFS service. This list is only relevant if the Delegation type (above) is Constrained. Each filer appears on a separate line, in the format "cifs/*filer-name*." The "cifs/" string indicates that the service can delegate "CIFS" to the filer; if another service type appears here, the line is irrelevant. If no filers appear here and the Delegation is Constrained, an administrator with sufficient credentials must access the service's computer account at a local DC and add them. The DCs must trust the CIFS service to delegate CIFS to all the filers behind it. You can use probe delegate-to to get a complete list of the CIFS filers behind the CIFS service. As above, this field is updated by a DC query that runs every 10 minutes.

Guidelines: Output for Each Service (Cont.)

Description, is the same string as explained for the summary output, above.

State is Enabled or Disabled, as set by the enable (gbl-cifs, gbl-nfs) command. If "(Global Server Disabled)" appears, clients cannot access the CIFS service because its global server is disabled. In this case, use enable (gbl-gs, gbl-gs-vs) to enable the global server, too.

SMB2 Protocol indicates whether or not this service supports SMB2. It has one of the following settings:

- Available (*smb-version*) means that SMB2 is available for client use. The *smb-version* is the maximum SMB version supported at all filers behind the service.
- Not Available shows that SMB2 is administratively allowed, but at least one filer behind the CIFS service does not support it. Use the show cifs-service filer-capabilities *fqdn* command to list all filers behind the *fqdn* service and display their SMB2 capabilities, if any.
- Not Allowed (service) means that SMB2 is disabled for this particular CIFS service. Use smb2-allowed (gbl-cifs) to enable SMB2 for a CIFS service.
- Not Allowed (global) indicates that SMB2 is disabled for the entire system. Use smb2-allowed to enable SMB2 for the ARX.

Signatures describes this service's setting for SMB signing, a CIFS security feature. This is Enabled (use SMB signing if and only if the client requires it), Required (only allow connections from clients that agree to use SMB signing), or Not Enabled (only allow connections from clients that do *not* require SMB signing). You can change this with the signatures command.

WINS Name Encoding is the character-encoding method used for the service's WINS-name registration. If the service uses WINS for name resolution, this is the character encoding used to send the CIFS service's name to the WINS server. You can use the wins-name-encoding command to change this.

Then there are sub-tables with the service's CIFS shares. Each sub table shows the shares from one namespace:

Exports for Namespace: *ns-name* identifies one namespace behind this CIFS service. The export (gbl-cifs) command establishes a namespace for the service.

Share Name is the name that is advertised to clients.

S appears for *subshares* of an imported share. To configure a subshare for this service, use the filer-subshare argument in the export (gbl-cifs) command, or use sync subshares from-namespace.

Volume Path is the directory path that is being shared, set by the export (gbl-cifs) command.

Guidelines: Output for Each Service (Cont.)

State describes the state of the volume behind this share:

- Online indicates that the volume is operational.
- Degraded: At least one of the volume's shares is offline, and at least one share is online. CIFS clients can only access the files and directories on the online share(s).
- Offline: All shares in the backing volume are disabled or unreachable. CIFS clients cannot access the front-end share.
- Read-Only: CIFS clients cannot write to this volume because it was enabled with no modify set. If the volume has no import issues, you can go to gbl-ns-vol mode and use modify to allow clients to write.
- Unavail: Some or all backing shares are enabled, but the volume is not starting or stopping. An attempt to map to this share may time out.
- Snapshot: The share is an export of the ~snapshot directory (or its equivalent; you can set the directory name with the snapshot directory cifs-name and/or snapshot directory nfs-name commands). This is a virtual directory created and presented by the volume software.
- Not Found: The Directory named in the export does not exist.

Degraded (at the end of the row) only appears for a CIFS subshare. This indicates that the filer subshares behind this client-visible subshare are not replicated properly. All of the filer subshares must have the same name (or a special name generated by the ARX), the same directory path (relative to the root of each filer share), and the same ACL. If this flag appears for any subshare, the CLI provides a prompt at the bottom suggesting that you use sync subshares from-service ... tentative to get details about the problem.

Guidelines: Detailed Output

Detailed output displays the same configuration and state fields for each CIFS service, followed by a separate table for each of the service's CIFS shares. Each share table has the following fields:

(The name of the share is the header.)

Namespace is the ARX namespace that holds the share's storage.

Volume Path is the ARX volume in the above namespace.

Description is set by an option in the export (gbl-cifs) command.

Export state describes the state of the client-side export. This often depends on the results of the subshare-replication process that occurs during managed-volume import. If the backing volume contains any CIFS subshares, the volume must replicate all subshare definitions on all of its filers:

- Ready: Clients can use the share or subshare. This indicates that the subshare replication process succeeded, or that there are no subshares under this share.
- Idle: The CIFS service is not-yet enabled, so clients cannot connect (or map a drive) to it. You can use enable (gbl-cifs, gbl-nfs) to enable the CIFS service.
- Pending: This only appears for a CIFS subshare. The subshare and its ACL is currently being replicated at all back-end filers behind the volume.
- NSM Wait: Subshare replication has completed successfully for the backing volume, and the service is now programming the volume's shares and subshares on the network (NSM) processors.
- Degraded: This only appears for a CIFS subshare. At least one of the filer subshares for this front-end subshare failed to replicate properly. That is, one of the back-end filers does not have the same subshare name and/or ACL. CIFS clients can only access the files and directories on the replicated subshare(s). As mentioned above, you can use the sync subshares from-service ... tentative command to get details about the problem.
- Unknown: Contact F5 Support if this status appears in the output.

Path State is the same as the State field in the summary output, described above.

Filer-subshare is Yes or No, depending on whether or not the share is implemented as a subshare. To create a subshare, use the filer-subshare argument in the export (gbl-cifs) command, or use sync subshares from-namespace to export all subshares from a particular backing volume.

Offline-access shows the type of offline access allowed to the share's clients. Clients can use this feature to make a local copy of a file on this share, edit the copy while disconnected from the share, then sync the copy with the original on the share once the connection is re-established. This shows the value set by the export offline-access command on the current share.

Samples

`bstnA> show cifs-service`

lists all global servers (by their FQDNs) configured with a CIFS front-end service. See Figure 24.3 on page 24-50 for sample output.

`bstnA> show cifs-service ac1.medarch.org`

shows a summary for the CIFS service running at "ac1.medarch.org." For sample output, see Figure 24.4 on page 24-50.

`bstnA> show cifs-service all`

shows summaries for all the front-end CIFS services on the "bstnA" chassis. See Figure 24.5 on page 24-51 for sample output.

`stoweA> show cifs-service patrol.MEDARCH.ORG detailed`

shows details for the CIFS service at "patrol.MEDARCH.ORG." For sample output, see Figure 24.6 on page 24-53.

Related Commands

cifs
enable (gbl-cifs, gbl-nfs)
export (gbl-cifs)
probe delegate-to
filer-subshares
sync subshares from-namespace
sync subshares from-service
smb2-allowed
smb2-allowed (gbl-cifs)
show cifs-service filer-capabilities
show cifs-service exports
show cifs-service user-sessions
drop cifs-service user-session
show cifs-service open-files
close cifs file
show statistics cifs authentication

Figure 24.3 *Sample Output: show cifs-service*

```
bstnA> show cifs-service

Service Name                Description
--------------------------------------------------
ac1.MEDARCH.ORG             insurance-claim records
labtests.MEDARCH.ORG        tests and records for lab animals
```

Figure 24.4 *Sample Output: show cifs-service ac1.medarch.org*

```
bstnA> show cifs-service ac1.medarch.org

Service Name:          ac1.MEDARCH.ORG
Domain Join:           Joined to MEDARCH.ORG
Account Name:          ac1$
Delegation:            Constrained, Any Protocol
Delegate To:           cifs/PV770N
                cifs/PV770N.MEDARCH.ORG
                cifs/VM-PV770N-01
```

```
                cifs/VM-PV770N-01.MEDARCH.ORG
                cifs/VM-SWP2003S1-5
                cifs/VM-SWP2003S1-5.MEDARCH.ORG
                cifs/VM-SWP2003S2-04
                cifs/VM-SWP2003s2-04.MEDARCH.ORG
                cifs/VM-SWP2008S-01
                cifs/VM-SWP2008S-02
                cifs/VM-SWP2008s-01.MEDARCH.ORG
                cifs/VM-SWP2008s-02.MEDARCH.ORG
                cifs/engdm
                cifs/engdm.MEDARCH.ORG
                cifs/enterprise
                cifs/enterprise.wwmed.com
                cifs/ntap-prov
                cifs/ntap-prov.MEDARCH.ORG
                cifs/ntap820
                cifs/ntap820.MEDARCH.ORG
Description:            insurance-claim records
State:                  Enabled
SMB2 Protocol:          Not Available
Signatures:             Enabled
WINS Name Encoding:     ISO-8859-1

Exports for Namespace: insur

  Share Name                        Volume Path                          State
  --------------------------------- - ------------------------------------ -------
  CLAIMS                            /claims                              Online
  SPECS                             /claims/specs                        Online
  STATS                             /claims/stats                        Online

Exports for Namespace: medarcv

  Share Name                        Volume Path                          State
  --------------------------------- - ------------------------------------ -------
  ARCHIVES                          /rcrds                               Online
  bulkstorage                       /rcrds                               Online
  CELEBS                          S /rcrds/VIP_wing                      Online
  chem_results                      /test_results                        Online
  labs                              /lab_equipment                       Online
  MP3S                            S /rcrds/2011/mp3downloads             Online
  xraysScanners                     /lab_equipment                       Online
  Y2004                           S /rcrds/2004                          Online
  Y2005                           S /rcrds/2005                          Online
  Y2008                           S /rcrds/2008                          Online
  Y2010                           S /rcrds/2010                          Online

  S = filer-subshare export
```

Figure 24.5 *Sample Output: show cifs-service all*

```
bstnA> show cifs-service all

Service Name:           ac1.MEDARCH.ORG
Domain Join:            Joined to MEDARCH.ORG
Account Name:           ac1$
Delegation:             Constrained, Any Protocol
Delegate To:            cifs/PV770N
                cifs/PV770N.MEDARCH.ORG
                cifs/VM-PV770N-01
                cifs/VM-PV770N-01.MEDARCH.ORG
                cifs/VM-SWP2003S1-5
```

```
                cifs/VM-SWP2003S1-5.MEDARCH.ORG
                cifs/VM-SWP2003S2-04
                cifs/VM-SWP2003s2-04.MEDARCH.ORG
                cifs/VM-SWP2008S-01
                cifs/VM-SWP2008S-02
                cifs/VM-SWP2008s-01.MEDARCH.ORG
                cifs/VM-SWP2008s-02.MEDARCH.ORG
                cifs/engdm
                cifs/engdm.MEDARCH.ORG
                cifs/enterprise
                cifs/enterprise.wwmed.com
                cifs/ntap-prov
                cifs/ntap-prov.MEDARCH.ORG
                cifs/ntap820
                cifs/ntap820.MEDARCH.ORG
Description:            insurance-claim records
State:                  Enabled
SMB2 Protocol:          Not Available
Signatures:             Enabled
WINS Name Encoding:     ISO-8859-1

Exports for Namespace: insur

  Share Name                        Volume Path                          State
  --------------------------------- - ------------------------------------ -------
  CLAIMS                            /claims                              Online
  SPECS                             /claims/specs                        Online
  STATS                             /claims/stats                        Online

Exports for Namespace: medarcv

  Share Name                        Volume Path                          State
  --------------------------------- - ------------------------------------ -------
  ARCHIVES                          /rcrds                               Online
  bulkstorage                       /rcrds                               Online
  CELEBS                          S /rcrds/VIP_wing                      Online
  chem_results                      /test_results                        Online
  labs                              /lab_equipment                       Online
  MP3S                            S /rcrds/2011/mp3downloads             Online
  xraysScanners                     /lab_equipment                       Online
  Y2004                           S /rcrds/2004                          Online
  Y2005                           S /rcrds/2005                          Online
  Y2008                           S /rcrds/2008                          Online
  Y2010                           S /rcrds/2010                          Online

  S = filer-subshare export


Service Name:           labtests.MEDARCH.ORG
Domain Join:            Joined to MEDARCH.ORG
Account Name:           labtests$
Delegation:             Constrained, Any Protocol
Delegate To:            cifs/VM-SMB2K8R2-01
                cifs/VM-SMB2K8R2-01.MEDARCH.ORG
                cifs/VM-SMB2K8R2-02
                cifs/VM-SMB2K8R2-02.MEDARCH.ORG
                cifs/VM-SWP2008S-01
                cifs/VM-SWP2008S-02
                cifs/VM-SWP2008s-01.MEDARCH.ORG
                cifs/VM-SWP2008s-02.MEDARCH.ORG
Description:            tests and records for lab animals
State:                  Enabled
```

```
SMB2 Protocol:         Available (2.002)
Signatures:            Not Enabled
WINS Name Encoding:    ISO-8859-1

Exports for Namespace: labResearch

  Share Name                        Volume Path                          State
  --------------------------------- - ------------------------------------ -------
  MICE                              /labMice                             Online
  RATS                              /labRats                             Online
```

Figure 24.6 *Sample Output: show cifs-service ... detailed*

```
stoweA> show cifs-service patrol.MEDARCH.ORG detailed

Service Name:          patrol.MEDARCH.ORG
Domain Join:           Joined to MEDARCH.ORG
Account Name:          patrol$
Delegation:            Constrained, Any Protocol
Delegate To:           cifs/MEDARCH
                cifs/MARLEY
                cifs/MARLEY.MEDARCH.ORG
                cifs/SYTFL
                cifs/cifsslave.MEDARCH.ORG
                cifs/sytfl.MEDARCH.ORG
Description:           medical server for ski patrols
State:                 Enabled
SMB2 Protocol:         Not Allowed (global)
Signatures:            Enabled
WINS Name Encoding:    ISO-8859-1

  Shares
  ------
    MEDIC
      Namespace        lodges
      Volume path      /skiPatrol
      Description      ski patrol medic server
      Export state     Ready
      Path state       Online
      Filer-subshare   No
      Offline-access   Manual

    C$
      Namespace        lodges
      Volume path      /acopia$ns1
      Description      Managed volumes in namespace 'lodges'
      Export state     Ready
      Path state       Online
      Filer-subshare   No
      Offline-access   Manual
```

show cifs-service filer-capabilities

Purpose A front-end CIFS service can support SMB2 if (and only if) all of the filers behind it can also support SMB2. When a CIFS client negotiates for SMB2 service, it may later access storage on any of the CIFS service's back-end filers, so they all need to support SMB2 calls. To confirm that all of the filers behind a given CIFS service support SMB2, or to identify the filers that do not support SMB2, use the show cifs-service filer-capabilities command.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show cifs-service filer-capabilities` ***fqdn***

fqdn (1-128 characters) is the fully-qualified domain name for one CIFS service (for example, "www.mycharity.org").

Guidelines This command displays information collected when each of the shares was enabled (see enable (gbl-ns-vol-shr)). The shares must be enabled in an ARX volume before the information in this output is complete.

Guidelines: Output The output contains one table per filer, followed by a table of Merged capabilities, and a final table of Service capabilities. The Merged capabilities are the capabilities that are supportable by all of the filers, and the Service capabilities are the subset of the Merged capabilities that are also supportable by the ARX service.

Each table contains the following fields to describe the SMB and SMB2 capabilities:

SMB describes the supported capabilities for the original SMB protocol. This shows "not available" if the filer does not support SMB. Otherwise, the next two lines describe the SMB capabilities:

Max Buffer is the negotiated size for buffer memory allocated at each end of the SMB connection. This is the negotiated maximum size of any SMB packet.

The next line is a list of SMB capabilities:

- Large Files,
- Large Read,
- Large Write, and/or
- Info Passthru.

Guidelines: Output (Cont.)

SMB2 describes the supported capabilities for the SMB2 protocol. This shows "not available" if the filer does not support SMB2. Otherwise, the next two lines describe the SMB2 capabilities:

The first line shows three transmission maximums for an SMB2 connection:

- Max Read,
- Max Write, and
- Max Transact.

The next line is a list of SMB2 capabilities. The line only appears if one of these capabilities is found:

- Leasing, and/or
- Large MTU.

These capabilities are accepted from the filer, but not supported by the ARX or advertised to its clients.

Has not been probed by current software appears at the bottom in the following situation: the service existed on an ARX release that did not support SMB2, the ARX was upgraded to a release that supports SMB2, and the ARX software has not yet tested the filers for SMB2 support. The probe should be complete in about 15 minutes.

Guidelines: Consult the Compatibility Matrix

We recommend using only SMB2 filers that have been tested for SMB2 support at F5. The *F5 Data Solutions Compatibility Matrix* (included with this doc set) lists all filers that have been fully qualified for SMB2 support behind the ARX.

Samples

`bstnA> show cifs-service filer-capabilities ac1.medarch.org`

shows the SMB and SMB2 capabilities for all filers behind the "ac1.medarch.org" service. The sample in Figure 24.7 shows that their are several filers behind that service that do not support SMB2.

`bstnA> show cifs-service filer-capabilities labtests.medarch.org`

shows the SMB and SMB2 capabilities for the "ac1.medarch.org" service. Figure 24.8 on page 24-56 shows that this service can fully support SMB2.

Related Commands

cifs

Figure 24.7 *Sample Output: show cifs-service filer-capabilities ac1*

```
bstnA> show cifs-service filer-capabilities ac1.MEDARCH.ORG

Capabilities of CIFS/SMB2 filers used by service ac1.MEDARCH.ORG

Filer fs1 [192.168.25.20]
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2: not available

Filer fs2 [192.168.25.27]
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2: not available

Filer fs3 [192.168.25.28]
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2.002: Max Read 65536, Max Write 65536, Max Transact 65536
```

```
Filer fs4 [192.168.25.29]
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2.002: Max Read 65536, Max Write 65536, Max Transact 65536

Filer fs5 [192.168.25.71]
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2: not available

Filer nas1 [192.168.25.21]
  SMB: Max Buffer 33028
       Large Files, Large Read, Large Write
  SMB2: not available

Filer nas10 [192.168.25.49]
  SMB: Max Buffer 33028
       Large Files, Large Read, Large Write
  SMB2.002: Max Read 65536, Max Write 65536, Max Transact 65536

Filer nasE1 [192.168.25.51]
  SMB: Max Buffer 65535
       Large Files, Large Read, Large Write, Info Passthru
  SMB2.002: Max Read 65536, Max Write 65536, Max Transact 65536

Merged capabilities
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write
  SMB2: not available

Service capabilities
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write
  SMB2: not available
```

Figure 24.8 *Sample Output: show cifs-service filer-capabilities labtests*

```
bstnA> show cifs-service filer-capabilities labtests.MEDARCH.ORG

Capabilities of CIFS/SMB2 filers used by service labtests.MEDARCH.ORG

Filer fs3 [192.168.25.28]
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2.002: Max Read 65536, Max Write 65536, Max Transact 65536

Filer fs4 [192.168.25.29]
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2.002: Max Read 65536, Max Write 65536, Max Transact 65536

Filer fs6 [192.168.25.30]
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2.1: Max Read 1048576, Max Write 1048576, Max Transact 1048576
       Leasing, Large MTU

Filer fs7 [192.168.25.41]
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2.1: Max Read 1048576, Max Write 1048576, Max Transact 1048576
```

```
       Leasing, Large MTU

Merged capabilities
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2.002: Max Read 65536, Max Write 65536, Max Transact 65536

Service capabilities
  SMB: Max Buffer 16644
       Large Files, Large Read, Large Write, Info Passthru
  SMB2.002: Max Read 65536, Max Write 65536, Max Transact 65536
```

show cifs-service subshares

Purpose

A *subshare* is a CIFS share of another CIFS share's subdirectory. You can configure a volume to pass CIFS clients from a front-end subshare directly to the corresponding subshare on the back-end, therefore applying the correct share-level ACL for the client's access. Use the show cifs-service subshare command to show the mapping of front-end subshares to their corresponding back-end subshares.

Mode

(any)

Security Role(s)

crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show cifs-service subshares {fqdn|all} [detailed]
  [state {ready | degraded | idle | pending | nsmwait}]
```

***fqdn*|all** chooses one (or all) CIFS services:

fqdn (1-128 characters) is the fully-qualified domain name for one CIFS service (for example, "www.mycharity.org").

all shows the subshares in all CIFS services.

detailed (optional) expands the output to show details for each subshare.

state {ready | degraded | idle | pending | nsmwait} (optional) filters the output so that it shows only the subshares in the specified state. The *Guidelines* below explain each of these states in detail.

Guidelines: Summary Output

The summary output is one table for each CIFS service. Each table is labeled with the following field:

Service is the FQDN that identifies the CIFS service. The cifs command, which instantiates the CIFS service, establishes this FQDN. It is the same as the FQDN for the global server where this service runs.

The table contains one filer subshare per row. The filer subshares are grouped together with their front-end exports. Each row contains the following fields:

Export is the name of the front-end subshare, seen by clients of the CIFS service. The export (gbl-cifs) ... filer-subshare command establishes this name.

Filer is the name of an external filer that hosts a back-end subshare. This is the external-filer name defined in the ARX configuration; use show external-filer to see a full list of external filers.

Share is the name of the parent share at the back-end filer.

Subshare is the name of the subshare at the back-end filer. Some subshares that are automatically generated by the ARX (with filer-subshares, sync subshares from-namespace, or sync subshares from-service) may have names with the following format: "_acopia_*subshare_id*$," where *subshare* is the name of the original subshare and *id* is a unique integer to identify the replica. If the subshare synchronization is underway and the back-end subshare is temporarily inaccessible to clients, the term [Fenced] appears in place of the subshare name.

Guidelines: Summary Output (Cont.)

State describes the results of the subshare-replication process that occurs during managed-volume import. If the backing volume contains any CIFS subshares, the volume must replicate all subshare definitions on all of its filers:

- Ready: This indicates that the subshare replication process succeeded, or that there are no subshares under this share.
- Idle: The CIFS service is not-yet enabled, so clients cannot connect (or map a drive) to it. You can use enable (gbl-cifs, gbl-nfs) to enable the CIFS service.
- Pending: This only appears for a CIFS subshare. The subshare and its ACL is currently being replicated at all back-end filers behind the volume.
- NSM Wait: Subshare replication has completed successfully for the backing volume, and the service is now programming the volume's shares and subshares on the network (NSM) processors.
- Degraded: This only appears for a CIFS subshare. At least one of the filer subshares for this front-end subshare failed to replicate properly. That is, one of the back-end filers does not have the same subshare name and/or ACL. CIFS clients can only access the files and directories on the replicated subshare(s). As mentioned above, you can use the sync subshares from-service ... tentative command to get details about the problem.
- Unknown: Contact F5 Support if this status appears in the output.

Guidelines: Detailed Output

If you used the detailed flag in any command, the output contains a complete table for each subshare:

Export is the name of the CIFS service's subshare, as seen by CIFS clients.

This table contains one sub table per back-end instance of the subshare. Each sub table contains the following fields:

Filer is the name of the back-end filer, as configured on the ARX. For details on this filer (such as its IP address), use the show external-filer *filer-name* command.

Share is the name of the imported share on the filer. The subshare resides in this share.

Subshare is the name of the subshare on the filer. This is described above.

State is also described above.

Samples

```
bstnA> show cifs-service subshares ac1.medarch.org
```

shows all of the front-end subshares in the "ac1.medarch.org" service, and all of the back-end subshares that correspond with them. The sample in Figure 24.9 shows several front-end subshares, where each is connected to three back-end subshares.

```
bstnA> show cifs-service subshares ac1.medarch.org detailed
```

shows details for the above subshares. Figure 24.9 shows sample output.

```
canbyA> show cifs-service subshares all state degraded

No CIFS subshare mappings.
```

shows all degraded subshares on the "canbyA" chassis. None of the subshares on this system are in the "degraded" state.

Related Commands

cifs
export (gbl-cifs)
sync subshares from-namespace
sync subshares from-service
show cifs-service exports
show cifs-service user-sessions

Figure 24.9 *Sample Output: show cifs-service subshares*

```
bstnA> show cifs-service subshares ac1.MEDARCH.ORG

Service:           ac1.MEDARCH.ORG

Export               Filer                Share                Subshare             State
-------------------- -------------------- -------------------- -------------------- --------
Y2005                 fs4                 prescriptions        Y2005                Ready
                      fs1                 histories            Y2005                Ready
                      fs2                 bulkstorage          Y2005                Ready
CELEBS                fs4                 prescriptions        CELEBS$              Ready
                      fs1                 histories            CELEBS$              Ready
                      fs2                 bulkstorage          CELEBS$              Ready
MP3S                  fs4                 prescriptions        MP3S                 Ready
                      fs1                 histories            MP3S                 Ready
                      fs2                 bulkstorage          MP3S                 Ready
Y2004                 fs4                 prescriptions        Y2004                Ready
                      fs1                 histories            Y2004                Ready
                      fs2                 bulkstorage          Y2004                Ready
Y2010                 fs4                 prescriptions        Y2010                Ready
                      fs1                 histories            Y2010                Ready
                      fs2                 bulkstorage          Y2010                Ready
```

Figure 24.10 *Sample Output: show cifs-service subshares detailed*

```
bstnA> show cifs-service subshares ac1.MEDARCH.ORG detailed

Service: ac1.MEDARCH.ORG

  Export: Y2005
    Filer:    fs4
    Share:    prescriptions
    Subshare: Y2005
    State:    Ready

    Filer:    fs1
    Share:    histories
    Subshare: Y2005
    State:    Ready

    Filer:    fs2
    Share:    bulkstorage
    Subshare: Y2005
    State:    Ready

  Export: CELEBS
    Filer:    fs4
    Share:    prescriptions
    Subshare: CELEBS$
    State:    Ready

    Filer:    fs1
```

```
  Share:    histories
  Subshare: CELEBS$
  State:    Ready

  Filer:    fs2
  Share:    bulkstorage
  Subshare: CELEBS$
  State:    Ready

Export: MP3S
  Filer:    fs4
  Share:    prescriptions
  Subshare: MP3S
  State:    Ready

  Filer:    fs1
  Share:    histories
  Subshare: MP3S
  State:    Ready

  Filer:    fs2
  Share:    bulkstorage
  Subshare: MP3S
  State:    Ready

Export: Y2004
  Filer:    fs4
  Share:    prescriptions
  Subshare: Y2004
  State:    Ready

  Filer:    fs1
  Share:    histories
  Subshare: Y2004
  State:    Ready

  Filer:    fs2
  Share:    bulkstorage
  Subshare: Y2004
  State:    Ready

Export: Y2010
  Filer:    fs4
  Share:    prescriptions
  Subshare: Y2010
  State:    Ready

  Filer:    fs1
  Share:    histories
  Subshare: Y2010
  State:    Ready

  Filer:    fs2
  Share:    bulkstorage
  Subshare: Y2010
  State:    Ready
```

show dynamic-dns

Purpose A CIFS service can use dynamic DNS to register its host name(s) and VIP with the local DNS server. Up-to-date DNS information is required for Kerberos authentication. To display the configuration and current status of all dynamic-DNS connections, use the show dynamic-dns command.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show dynamic-dns [fqdn]`

fqdn (optional; 1-128 characters) is the fully-qualified domain name for one CIFS service (for example, "www.organization.org"). If you omit the FQDN, this command shows DNS summaries for all CIFS services.

Guidelines For every CIFS service that uses dynamic DNS, the output contains one table with a summary for the service and a sub-table with one row per DNS alias. The first table contains three fields:

Svc is the type of front-end service. This is always "CIFS" in the current implementation.

Global Server identifies the global server for the CIFS service. The cifs service uses the same name.

Domain Name is the windows-domain (gbl-gs) for the Global Server. This is also the Windows domain for the CIFS service.

Each CIFS-service table has a sub-table of its host names:

Status is one of the following values:

- "OK" indicates that the CIFS service successfully added this host name.
- "Retry" means none of the DNS name-servers acknowledged the add or remove operation. (The operation type, add or remove, is shown in the Operation field.) The CIFS service retries once per minute.
- "Failed" indicates that a remove operation failed. It retries an add operation forever, but only retries a remove operation 15 times. You must remove the "A" record manually at the DNS server itself. Any Failed removes remain in this output until you clear them all (system-wide) with clear dynamic-dns.
- "Disabled" means the CIFS service is disabled (with no enable (gbl-cifs, gbl-nfs)) and therefore is not adding or removing any records from DNS.
- "Pending" indicates that the operation has started but is incomplete for this host name; another host name is in process. This state should change quickly to one of the other states.

Host Name is a host name that was configured for this CIFS service, set by the dynamic-dns command.

VIP is the virtual-IP address for this CIFS service, set by the gbl-gs-vs virtual server command. The CIFS service registers the Host Name with the VIP in a single "A" record.

Guidelines (Cont.)

Operation is "add" or "remove." The dynamic-dns command triggers an add operation. A remove is triggered by commands that delete a DNS host name (no dynamic-dns) or change the VIP (virtual server). This field is empty if the CIFS service is disabled.

Retries is the number of failed attempts to register the host name. This should be 0 (zero); if it is greater than 0, you should check the DNS server(s) that are configured for the domain. (The name-server command configures one DNS server for an Active-Directory domain.) The syslog catalogs all registration failures; use show logs syslog to view the syslog. As above, this field is empty if the CIFS service is disabled.

Last Update is the time and date of the last attempt to register this host name.

DNS Server is the IP address of the name-server where the Last Update was sent. If the Status is "FAIL," this server may be down. If the Status is "OK" but the Retries are 1 or more, the registration failed at some other DNS server(s) before it succeeded at this one.

Sample

```
bstnA> show dynamic-dns ac1.medarch.org
```

shows the current status of dynamic DNS for one CIFS front-end service. See Figure 24.11 on page 24-63.

Related Commands

dynamic-dns
clear dynamic-dns

Figure 24.11 *Sample Output: show dynamic-dns ac1.medarch.org*

```
bstnA> show dynamic-dns ac1.medarch.org

Svc     Global Server                   Domain Name
-------------------------------------------------------------------------
CIFS    ac1.MEDARCH.ORG                 MEDARCH.ORG

  Status   Host Name                                        VIP
           Operation      Retries   Last Update             DNS Server
  ---------------------------------------------------------------------
  OK       ac1.MEDARCH.ORG                                  192.168.25.15
           Add               0     Sat Jan 15 01:07:36 2011  192.168.25.102

  OK       fs1.MEDARCH.ORG                                  192.168.25.15
           Add               0     Sat Jan 15 01:07:36 2011  192.168.25.102

  OK       fs2.MEDARCH.ORG                                  192.168.25.15
           Add               0     Sat Jan 15 01:07:36 2011  192.168.25.102

  OK       fs5.MEDARCH.ORG                                  192.168.25.15
           Add               0     Sat Jan 15 01:07:37 2011  192.168.25.102

  OK       insur.MEDARCH.ORG                                192.168.25.15
           Add               0     Sat Jan 15 01:07:37 2011  192.168.25.102
```

show global service

Purpose Use the show global service command to cross-reference the front-end services with their global servers.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show global service [`*`fqdn`*`]`**

fqdn (optional; up to 128 characters) is the fully-qualified domain name (for example, "www.organization.org") for one global server.

If no FQDN is specified, this command lists the front-end services for all global servers.

Guidelines This displays a table of front-end services, ordered by global server.

Domain Name is the FQDN that identifies the global server.

Service is the type of front-end service:

- NFS or
- CIFS

State is enabled or disabled. Use enable (gbl-cifs, gbl-nfs) to enable a front-end service.

Samples

```
bstnA> show global service

Domain Name                 Service     State
-------------------------------------------------
acopiaFiler                 NFS         Enabled
ac1.MEDARCH.ORG             NFS         Enabled
ac1.MEDARCH.ORG             CIFS        Enabled
```

shows all front-end services at all global servers.

```
bstnA> show global service ac1.MEDARCH.ORG

Domain Name                 Service     State
-------------------------------------------------
ac1.MEDARCH.ORG             NFS         Enabled
ac1.MEDARCH.ORG             CIFS        Enabled
```

shows the front-end services for one global server, "ac1.MEDARCH.ORG."

Related Commands show global server
show virtual service

show nfs-service

Purpose Use the show nfs-service command to display configuration information about one or more front-end NFS services.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show nfs-service [fqdn | all] [detailed]
show nfs-service fqdn export namespace export-path
```

fqdn (1-128 characters) is the fully-qualified domain name (for example, "www.organization.org") for one NFS service. If no FQDN is specified, this command shows summaries for all NFS-service offerings.

all (optional) shows full configurations for all NFS-service offerings.

detailed (optional) shows full configurations for the selected NFS service(s), along with an additional table to show the behavior of each NFS export when one of its back-end shares is offline. When a client attempts to access a file that resides on an offline share, the NFS export can either send an access error to the client or silently wait for the client to retry. You can configure this behavior with the offline-behavior command.

export *namespace export-path* focuses on a single service export. This shows the full configuration and the offline-behavior details for the selected export.

namespace (1-30 characters) is the namespace that hosts the export, and

export-path (1-1024 characters) is the full path of the NFS export, starting with the volume name.

Guidelines: Summary View The summary output is a table of NFS-service summaries, one row per service. The table has three columns:

Domain Name is the FQDN that identifies the NFS service. The nfs command, which instantiates the NFS service, establishes this FQDN. It is the same as the FQDN for the global server where this service runs.

Description is a string to describe this service. You can change it with the description (gbl-nfs) command.

Guidelines: Full-Configuration View The detailed output shows similar fields, followed by additional configuration details:

Virtual Server is the same as the Domain Name in the summary output; it is the FQDN that identifies the NFS service.

Description also matches the summary output.

NFS State is "Enabled" or "Disabled." This is the administrative state, set with the enable (gbl-cifs, gbl-nfs) command.

NLM State is also "Enabled" or "Disabled." This shows whether or not the NFS Lock Manager (NLM) protocol is running for this NFS service. You can set this with the nlm enable command.

Guidelines: Full-Configuration View (Cont.)

Each NFS-service table has a sub-table for each of its backing namespaces. The sub tables contain all of the NFS exports backed by the given namespace:

Exports for Namespace *name* is the heading for each export table. This is the namespace that holds the back-end storage for the exports below.

Export As is the NFS-export name that is advertised to clients. You can enter this with the export (gbl-nfs) command.

Volume Path is the managed-volume path that is being shared. You also enter this as part of the export (gbl-nfs) command. An asterisk (*) appears after this if the export returns an access error for back-end shares that are offline; you can use the offline-behavior command to change this behavior.

State is the volume's state:

- Online
- Degraded: At least one share is offline, and at least one share is online. NFS clients may be able to access the files and directories on the online share(s).
- Offline: All shares in the volume are disabled or unreachable. If a client attempts to mount this export, the NFS service denies the mount with MNTERR_ACCESS.
- Read-Only: NFS clients cannot write to this volume because it was enabled with no modify set. Go to gbl-ns-vol mode and use modify to allow clients to write.
- Unavailable: Some or all shares are enabled, but the volume is not starting or stopping. An attempt to mount this export may time out.
- Unknown
- Not Found: The Directory named in the export does not exist, An attempt to mount this directory fails with MNTERR_NOENT.

Access shows the permissions that clients have for accessing this service. If an NFS access list is assigned to this service (with the nfs-access-list command), the name of the access list appears here. An "(rw)" appears if no access list is configured; this indicates that all clients have read and write access to the service.

(NIS *ip-address*) appears if the access list refers to at least one NIS netgroup. This is the address of the NIS server used for the most-recent NIS update. The word ERROR appears instead of the IP address if the netgroup is not completely resolved. The error indicates that either a netgroup has caused the access list to grow too large, or one of the hosts in one of the netgroups cannot be resolved through a DNS lookup. Use show nis domain to find more details about the NIS configuration and its current status.

Guidelines: Detailed View

The Offline-behavior table appears if you enter the detailed option in the command. This table shows the export behavior when clients access a back-end filer that is offline: the export either returns an access error to the client or silently allows the client to retry. The top row shows the default offline-behavior for the whole service, and the following rows (if any) each show the offline behavior for an individual namespace, volume, or export in the service. Each row contains the following columns:

Namespace is one of the namespaces behind the current NFS service. This contains "All" in the top row and a specific namespace name in all subsequent rows.

Volume is one volume in the namespace, or "All" if the behavior applies to all of the namespace's volumes.

Path is a volume path that is exported to NFS clients, or "All" to indicate that the behavior applies to all of this volume's exported paths.

Behavior is "deny-access" or "retry." This shows the export behavior when a client accesses a file or directory on a back-end share that is currently offline. Either the client receives an access error or the NFS service silently waits for the client to retry the request.

Count shows the number of NFS exports that use this behavior. If the count is 0 (zero) for a given behavior row, there is no export with the exact Path in this row. You can use the command with the exact path in this row for the offline behavior to take effect.

The NFS client chooses one of these behaviors when mounting the NFS service. If the client's **mount** command exactly matches a Path in any row, the client gets the Behavior in that row. Otherwise, the client gets the behavior in the row that matches the closest, with an ancestor Path, the same Volume, or the same Namespace.

You can use this output to determine the offline behavior that your NFS clients should expect from any given export.

Samples

`bstnA> show nfs-service`

shows all the front-end NFS services on the ARX. Sample output appears in Figure 24.12.

`bstnA> show nfs-service ac1.MEDARCH.ORG`

shows the NFS service running at "ac1.MEDARCH.ORG." See Figure 24.13 for sample output.

`bstnA> show nfs-service ac1.MEDARCH.ORG detailed`

shows the NFS service running at "ac1.MEDARCH.ORG," along with details about the offline-filer behavior. See Figure 24.14 for sample output.

`bstnA> show nfs-service acopiaFiler`

shows the NFS service running at "acopiaFiler," an NFS service where NLM is disabled. See Figure 24.15 for sample output.

Related Commands

nfs
description (gbl-nfs)
nfs-access-list
offline-behavior
show nfs-service mounts

Figure 24.12 *Sample Output: show nfs-service*

```
bstnA> show nfs-service

Domain Name                Description
-----------------------------------------------
acopiaFiler
ac1.MEDARCH.ORG            insurance records for WW Medical
```

Figure 24.13 *Sample Output: show nfs-service ac1.medarch.org*

```
bstnA> show nfs-service ac1.MEDARCH.ORG

Virtual Server: ac1.MEDARCH.ORG
Description:    insurance records for WW Medical

NFS State:      Enabled
NLM State:      Enabled

Exports for Namespace: insur

  Export As             Volume Path                  State       Access
  ------------------------------------------------------------------------------
  /claims               /claims                      Online      eastcoast (NIS 192.168.25.2)

(*) Return permission error for offline back-end shares.

Exports for Namespace: wwmed

  Export As             Volume Path                  State       Access
  ------------------------------------------------------------------------------
  /acct                 /acct                     *  Online      eastcoast (NIS 192.168.25.2)
  /acct/wksheets        /acct/wksheets               Online      eastcoast (NIS 192.168.25.2)

(*) Return permission error for offline back-end shares.
```

Figure 24.14 *Sample Output: show nfs-service ac1.medarch.org detailed*

```
bstnA> show nfs-service ac1.MEDARCH.ORG detailed

Virtual Server: ac1.MEDARCH.ORG
Description:    insurance records for WW Medical

NFS State:      Enabled
NLM State:      Enabled

Exports for Namespace: insur

  Export As             Volume Path                  State       Access
  ------------------------------------------------------------------------------
  /claims               /claims                      Online      eastcoast (NIS 192.168.25.2)

(*) Return permission error for offline back-end shares.

Exports for Namespace: wwmed
```

```
  Export As             Volume Path                   State       Access
  ------------------------------------------------------------------------------
  /acct                 /acct                      *  Online      eastcoast (NIS 192.168.25.2)
  /acct/wksheets        /acct/wksheets                Online      eastcoast (NIS 192.168.25.2)

(*) Return permission error for offline back-end shares.


Offline-behavior:

  Namespace          Volume             Path                          Behavior       Count
  ------------------------------------------------------------------------------------------
  All                All                All                           deny-access        1
  wwmed              /acct              wksheets                      retry              1
  insur              All                All                           retry              1
```

Figure 24.15 *Sample Output: show nfs-service acopiaFiler*

```
bstnA> show nfs-service acopiaFiler

Virtual Server: acopiaFiler
Description:

NFS State:      Enabled
NLM State:      Disabled

Exports for Namespace: medco

  Export As             Volume Path                   State       Access
  ------------------------------------------------------------------------------
  /vol                  /vol                          Online      (rw)

(*) Return permission error for offline back-end shares.
```

show nfs-service mounts

Purpose Use the show nfs-service mounts command to show all NFS mounts to one or more NFS services.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show nfs-service mounts {`*`fqdn`*` | all}`**

fqdn (1-128 characters) specifies a particular NFS service (for example, "www.company.com").

all shows all client connections to all NFS services.

Guidelines The output is a table of NFS mounts from client machines. Two rows of information describe each currently-active client mount:

Global Server identifies the global server that the client mounted. The global server's name appears here as configured; depending on your external DNS configuration, it may not be an FQDN that is visible to clients. Use the show global server command for a full list of global servers, along with configuration details.

Mount Point is where the client connected. This is the volume path established in the export (gbl-nfs) command, possibly with an additional sub path.

Client IP Address identifies the client host.

[VIP] is the Virtual-IP address to which the client connected. This also appears in the output for show global server.

The total number of mount points appears at the bottom of the table.

Note: After a client un mounts, there may be a short time where the clients mount persists in the output.

Use the show nfs-service command to see NFS-service configuration and state.

Samples `bstnA> show nfs-service mounts all`
lists all mounts to all NFS services. For sample output, see Figure 24.16.

`bstnA> show nfs-service mounts acopiaFiler`
lists all mounts to a particular NFS service. See Figure 24.17 for sample output.

Related Commands nfs
show nfs-service

Figure 24.16 *Sample Output: show nfs-service mounts all*

```
bstnA> show nfs-service mounts all

Global Server      Mount Point                    Client IP Address
[VIP]
-------------------------------------------------------------------------------
acopiaFiler                  /vol/vol2                      172.16.100.183
[192.168.25.12]
```

```
ac1.MEDARCH.ORG              /acct                          172.16.100.183
[192.168.25.15]

Total number of mount points displayed is 2.
```

Figure 24.17 *Sample Output: show nfs-service mounts acopiaFiler*

```
bstnA> show nfs-service mounts acopiaFiler

Global Server      Mount Point                    Client IP Address
[VIP]
-------------------------------------------------------------------------------
acopiaFiler                  /vol/vol2                      172.16.100.183
[192.168.25.12]

Total number of mount points displayed is 1.
```

show nfs tcp

Purpose Use the show nfs tcp command to display the current timeout period for NFS/TCP connections to back-end shares, and the current client-side behavior in the event of this server-side timeout (disconnect or send an I/O error).

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show nfs tcp`

Guidelines Consider an NFS client that makes a request, where the ARX forwards the request to a back-end NFS server. If the server times out after 105 seconds, the switch disconnects the client by default. You can use nfs tcp timeout to change this reaction: instead of closing the client connection, NFS services can send an NFSERR_IO (NFSv2) or NFS3ERR_IO (NFSv3) back to the client application.

The output shows the current, system-wide setting for TCP timeouts:

Behavior is either "Close Connection" or "Return I/O Error."

Inactivity is the timeout period.

Sample

```
bstnA> show nfs tcp

Transaction Timeout
  Behavior:    Close Connection
  Inactivity:  105 seconds
```

shows that the timeout is set to its default values.

Related Commands nfs tcp timeout

show nlm client

Purpose The NFS Lock Manager (NLM) manages file locks for NFS clients. Use the show nlm client command to display currently-held locks by one or all NFS-client machines.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show nlm client {all | `*`hostname`*`}`**

all | *hostname* is a required choice:

all shows all NFS clients that hold any locks.

hostname (1-192 characters) identifies one client's host machine (for example, "myhost"). This shows all files locked by the given host.

Guidelines An NFS client can lock a file or a region within a file, so that other NLM-compliant users cannot access the region until the lock is released. The ARX's NLM server manages all the locks for all NFS front-end services and volumes.

The show nlm client all command presents a list of client hosts that currently hold locks, or held them over the last five minutes.

The show nlm client *hostname* command shows the SVID (the process ID of the lock holder), namespace, and file path for all files locked by one client. To see all clients that hold locks for one of these files, you can copy the namespace and file path from this output and paste it into the show nlm file command.

Use show nlm statistics to see statistics for all NLM activity. If a client machine is down, you can clear all of its locks with clear nlm locks.

To disable NLM for a given NFS service, disable NFS first (with enable (gbl-cifs, gbl-nfs)) and then use no nlm enable.

Samples

```
bstnA# show nlm client all

 NLM Clients
 ----------------
 rh1     (192.168.25.19)
```

finds that one client machine holds an NFS lock.

```
bstnA# show nlm client rh1

 Locks held by rh1     (192.168.25.19)

 SVID    Namespace/Path
-------------------------------------------------------------------------------
 17098   wwmed /acct/rvcIndex.html
```

searches for locks held by clients at the "rh1" machine. It finds a single locked file in the wwmed~/acct volume.

Related Commands show nlm file
show nlm statistics
clear nlm locks
no nlm enable

show nlm file

Purpose The NFS Lock Manager (NLM) manages file locks for NFS clients. Use the show nlm file command to show which clients have locks on a given file.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show nlm file`** ***`namespace file-path`***

namespace (1-30 characters) is the file's namespace.

file-path (1-4096 characters) identifies the file to check.

Guidelines An NFS client can lock a file or a region within a file, so that other NLM-compliant users cannot access the region until the lock is released. The ARX's NLM server manages all the locks for all NFS front-end services and volumes.

This command shows all locks for the file at *file-path*. These are locks that are currently held or have been held in the last five minutes. The following information is shown for each lock:

Client is the host name of a lock-holding client.

SVID (System V Interface Definition) is process ID of the program that holds the lock.

Locked Region shows which bytes are locked in the file.

Type is either "Exclusive" or "Shared." An exclusive lock is held by a single client, blocking access from all others. A shared lock can be shared by multiple clients, allowing them all to read the region but not change it.

If any clients are waiting for a lock to free up, an additional tables of "Waiters" appears before the "Locks" table. The "Waiters" table has the same columns as the "Locks" table.

To see all the locks held by a given client machine (or a list of all the client machines that have locks), use the show nlm client command. If a client is down, you can clear all of its locks with clear nlm locks. Use show nlm statistics to see statistics for all NLM activity.

To disable NLM for a given NFS service, disable NFS first (with no enable (gbl-cifs, gbl-nfs)) and then use no nlm enable.

This command fails if NLM is disabled for the all of the namespace's NFS services.

Samples

```
bstnA# show nlm file wwmed /acct/rvcIndex.html

 Locks on wwmed:/acct/rvcIndex.html

  Client                       SVID     Locked Region                     Type
-------------------------------------------------------------------------------
  rh1                          17098    0-512                          Exclusive
```

shows lock details for a file in the wwmed namespace, "/acct/rvcIndex.html."

```
bstnA# show nlm file EMC-NFS /NFS-EMC/lockfile

 Locks on EMC-NFS:/NFS-EMC/lockfile
  Client                        SVID     Locked Region                      Type
-------------------------------------------------------------------------------
  dartagnon                     4998     15-30                           Shared
  porthos                       234      15-30                           Shared
  aramis                        7898     15-30                           Shared
  athos                         11375    0-15                         Exclusive
  athos                         9372     4294967295-4294967310        Exclusive
  dartagnon                     765      15-30                           Shared
  porthos                       4095     15-30                           Shared
  client89                      19087    15-30                           Shared
  testrh8                       7702     15-30                           Shared
  aramis                        125      45-60                        Exclusive
  aramis                        259375   281474976710655-281474976710670 Exclusive

 Waiters on EMC-NFS:/NFS-EMC/lockfile
  Client                        SVID     Requested Region                   Type
-------------------------------------------------------------------------------
  dartagnon                     4095     15-30                        Exclusive
  caller6                       98223    15-30                        Exclusive
```

shows lock details for a file in the EMC-NFS namespace, "/NFS-EMC/lockfile." This file has several locks on it, and it has two clients waiting to put an exclusive lock on bytes 15-30.

Related Commands

show nlm client
show nlm statistics
clear nlm locks
no nlm enable

show nlm statistics

Purpose The NFS Lock Manager (NLM) manages file locks for NFS front-end service clients. Use the show nlm statistics command to display NFS-lock statistics.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show nlm statistics`

Guidelines An NFS client can lock a file or a region within a file, so that other NLM-compliant users cannot access the region until the lock is released. The ARX's NLM server manages all the locks for all NFS front-end services and volumes.

This command shows a table of system-wide NLM statistics. All of these statistics are reset to 0 (zero) at boot-up time.

Lock Requests is the total number of locks requested.

Locks Granted is the number of locks that the NLM server granted to clients.

Locks Denied counts the number of client requests that were rejected.

Locks Blocked is the number of client requests where the lock is unavailable but client is willing to block until the lock is free.

Unlock Requests is the number of requests to unlock a file region (NLM_UNLOCK requests).

Test Requests is the number of NLM_TEST requests, which a client sends to check for the availability of a lock.

Cancel Requests is the number of NLM_CANCEL requests, which a client sends cancel a blocked request.

Current Locks Held is the total locks that clients hold now, system-wide.

Cached Clients is the number of client machines with processes that hold (or recently held) locks. Lock-holding clients are cached for five minutes; this counts all the clients in that cache.

Cached Processes is the number of processes that have held locks over the past five minutes.

Cached File Handles is the number of NFS file handles that have been cached over the past five minutes. These identify files that have one or more locks on them.

Unsupported Lock Calls is the total number of DOS-file-sharing calls. The ARX does not support those NLM calls.

Unimplemented Lock Calls should always be zero; this indicates a malfunctioning client. If a client application is encountering NFS-locking errors (notably timeouts or hangs) then check this counter and look for NLM_UNIMPLEMENTED messages in the syslog. Those syslog messages specify the unsupported NLM function that was called.

To see all the client machines that hold locks for a given file, use the show nlm file command. For a list locks held by a given client machine (or a list of all the client machines that have locks), use the show nlm client command. If a client is down, you can clear all of its locks with clear nlm locks.

Guidelines (Cont.) To disable NLM for a given NFS service, disable NFS first (with enable (gbl-cifs, gbl-nfs)) and then use no nlm enable.

Sample See Figure 24.18 for sample output.

Related Commands show nlm file
show nlm client
clear nlm locks
no nlm enable

Figure 24.18 *Sample Output: show nlm statistics*

```
bstnA# show nlm statistics

 NLM Statistics
-------------------------------------------------------------------------------
 Lock Requests                  30
 Locks Granted                  15
 Locks Denied                   9
 Locks Blocked                  6
 Unlock Requests                21
 Test Requests                  9
 Cancel Requests                2

 Current Locks Held             3

 Cached Clients                 1
 Cached Processes               3
 Cached File Handles            1

 Unsupported Lock Calls         0
 Unimplemented Lock Calls       0
```

show virtual service

Purpose Use the show virtual service command to display the front-end services running on each virtual server.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show virtual service [hostname]`

hostname (optional; up to 128 characters) is the hostname for one ARX in the current redundant pair.

If no hostname is specified, this command lists the front-end services for both peers.

Guidelines This displays a table of front-end services, grouped by ARX.

Switch is the current ARX (or its redundant peer) where the virtual servers are running.

Global Server is the global server for which this virtual server is a host.

Virtual IP Address is the IP address where the front-end service is listening.

Service is the name of the front-end service:

- NFS or
- CIFS

State is the result of a health-check for the front-end service:

Disabled indicates that an administrator has disabled the front-end service, or that is was just enabled and is about to go into the Starting state. Use the enable (gbl-cifs, gbl-nfs) command to enable a disabled service.

Starting indicates that the NFS or CIFS service software is initializing or waiting for the exported volume(s) to become available. If (No Running Volumes) appears, too, the "Starting" state is frozen because of the service's namespace(s); you can use show cifs-service or show nfs-service to find the namespace name(s), then use show namespace to examine each namespace and its volumes.

Ready indicates that the NFS or CIFS service is ready to serve file data to external clients.

Stopping indicates a transition from Ready to Disabled. The service stops because someone has disabled it (no enable (gbl-cifs, gbl-nfs)), or because someone has removed it (no nfs or no cifs).

Failed indicates that the service has failed its health check. Call F5 Support if this state appears.

Suspended means that an nsck ... rebuild is re-creating one of the service's namespaces. When the rebuild finishes (or on reboot), the service goes back to the Starting state. Use show nsck to show the progress of all nsck jobs.

Unreachable indicates that the remote switch is unreachable over the RON, so the service state is unknown. Use show ron to examine the status of the RON.

Guidelines (Cont.)

These states should never appear, given that the CLI checks for these issues before it allows the virtual-server to be configured. Contact F5 if any of these results appear in the output:

- VIP In Use means that the VIP is the same as an internal IP address (such as a proxy IP) that is already in use.
- VMAC Error means that there are no virtual MAC (VMAC) addresses available to assign to the VIP (too many VIPs are already configured on this chassis).
- VIP Error is any other internal error.
- Unknown is an internal software error indicating that a virtual service has been configured but does not exist in the global service manager.
- Invalid is an internal software error indicating that the virtual service's internal numeric status value is invalid (values should vary from 1 to 10 only).

Samples

`prtlndA> **show virtual service**`

shows all virtual services on the "prtlndA" switch and its redundant peer. See Figure 24.19 for sample output.

`prtlndA> **show virtual service prtlndB**`

focuses on "prtlndB," the redundant peer.

Related Commands

show global service
show namespace
enable (gbl-cifs, gbl-nfs)
nsck ... rebuild
show nsck

Figure 24.19 *Sample Output: show virtual service*

```
prtlndA> show virtual service

Switch:  prtlndA
------------------------

  Global Server                              Virtual IP Address   Service    State
  ------------------------------------------------------------------------------------
  www.nemed.com                              192.168.74.91        NFS        Ready
  insurBkup.MEDARCH.ORG                      192.168.74.92        CIFS       Ready

Switch:  prtlndB
------------------------

  Global Server                              Virtual IP Address   Service    State
  ------------------------------------------------------------------------------------
```

signatures

Purpose

Use the signatures command to enable (or perhaps require) SMB signing between this CIFS front-end service and its clients. *SMB signing* is the process of placing a digital signature into each Server Message Block (SMB) exchanged between a CIFS server (this service) and any CIFS client. SMB signing prevents man-in-the-middle attacks at the cost of slower performance.

Use no signatures to refuse SMB signing from any CIFS clients. This prevents any new connections from clients that require SMB signing.

Mode

gbl-cifs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
signatures [required]
no signatures
```

required (optional) obligates all CIFS clients to use SMB signing in their communication with this service. If any of the clients refuse to support SMB signing, they cannot connect to the service.

Default(s)

`no signatures`

Guidelines

Without SMB signing, the default for the CIFS service, clients that require SMB signing refuse to start a CIFS session.

If you use this command to declare SMB signatures required, the service only accepts connections from clients that support SMB signing.

The least-restrictive setting is to enable SMB signing without requiring it (using the simple signatures syntax). The CIFS service can then accept any client connection, whether it requires or refuses SMB signing. If a client requests or requires SMB signatures, the CIFS service complies.

To control the SMB-signing policy between one of the service's backing namespaces and its back-end filers, you can use the cifs filer-signatures command in gbl-ns mode. To find the backing namespace(s) for this CIFS service, use show cifs-service *fqdn*, where *fqdn* is the name of the CIFS service. To see the number of filers and/or clients who have used SMB signing, along with some success and failure statistics, use the show fastpath cifs-signatures command.

Sample

```
bstnA(gbl-cifs[ac1.medarch.org])# signatures
```

allows SMB signing for any client of the "ac1.medarch.org" service. If a client requests or requires SMB signing, the CIFS service will comply. For clients that disallow SMB signing, the CIFS service does not send or expect any signatures.

Related Commands

cifs
enable (gbl-cifs, gbl-nfs)
cifs filer-signatures
show cifs-service
show fastpath cifs-signatures

smb2-allowed (gbl-cifs)

Purpose Use no smb2-allowed to disable SMB2 for the current CIFS service.

Use the smb2-allowed command to re-enable SMB2 between this CIFS front-end service and its clients.

Mode gbl-cifs

Security Role(s) storage-engineer or crypto-officer

Syntax

```
no smb2-allowed
smb2-allowed
```

Default(s)

```
smb2-allowed
```

Guidelines Use the no smb2-allowed command to disable SMB2 for the current CIFS service. The affirmative form of the command, smb2-allowed, re-enables SMB2.

Clients can use SMB2 when all of the following conditions are met:

- SMB2 must be enabled for the current service (with this command).
- SMB2 must be enabled for the current ARX. Use the smb2-allowed command to enable SMB2 for the entire system.
- All the filers behind this service must support SMB2. Use the show cifs-service filer-capabilities command to test all of the service's filers for SMB2 support.

This command affects all future CIFS connections, but does not affect currently-connected clients. If a client is already connected with SMB2 when you disallow it with this command, the client continues with SMB2 until the connection is cut, and all future connections use SMB.

Sample

```
bstnA(gbl-cifs[ac1.medarch.org])# no smb2-allowed
```

disables SMB2 for the "ac1.medarch.org" service.

Related Commands

cifs
smb2-allowed
show cifs-service filer-capabilities

sync cifs delegation

Purpose You can use the sync cifs delegation command to immediately synchronize an ARX cifs service with its delegation settings in the Active Directory (AD). With the proper *delegation* settings, a CIFS client can authenticate at the ARX once and then access any of the back-end servers without re-authenticating. This delegation feature is set in the AD, outside the ARX. When someone changes the delegation settings at a domain controller (DC), the ARX software can take up to 10 minutes to synchronize with the changes. Use this command to make the synchronization immediate.

Modes priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`sync cifs delegation {`*`fqdn`*` | all}`**

fqdn (1-255 characters) is the FQDN for a cifs service whose delegation settings changed in the AD.

all synchronizes the delegation settings for all CIFS services on the ARX.

Default(s) None.

Guidelines *Delegation* enables multiple CIFS servers to delegate client authentications to a front-end CIFS service, so that clients can authenticate with the front-end service once and remain authenticated for connections with the back-end servers. We recommend trusting your ARX CIFS service for delegation, to enable this single-sign-on feature. It is generally recommended that you also *constrain* this delegation to the particular CIFS servers behind this service. You set this trust (at a DC) in the CIFS service's machine account, or you can set it with the domain-join command if you have sufficient privileges.

If the service uses constrained delegation, you can use the probe delegate-to command to confirm that all the servers behind a particular CIFS service are properly assigned to its "delegate" list. Clients can only access the storage on a back-end server that is on that list.

Sample

```
stkbrgA# sync cifs delegation bgh.MEDARCH.ORG
  The cifs service 'bgh.MEDARCH.ORG' delegation type 'Constrained' is
synchronized with AD.
```

synchronizes all deletion and constraint settings for the "bgh.MEDARCH.ORG" service.

Related Commands

domain-join
show cifs-service
probe delegate-to

wins-name-encoding

Purpose

Character encoding is the mapping between binary numbers and written characters. Some character-encoding schemes use only a single byte for each character; these typically support alphabets without any Asian characters. Multi-byte encoding schemes encompass Asian character sets. Unicode (such as UTF-8) can also use more than one byte per character, and encompasses most character sets and languages. Use the wins-name-encoding command to set the CIFS service's character encoding for its NetBIOS names.

Use no wins-name-encoding to reset NetBIOS names back to the single-byte default.

Modes

gbl-cifs

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
wins-name-encoding {utf-8 | shift-jis | cp932 | euc-jp
  | ksc5601 | iso-8859-1}
no wins-name-encoding
```

utf-8 specifies UTF-8 (Unicode) character encoding.

shift-jis specifies Shift_JIS (Japanese) character encoding.

cp932 is Code Page 932, or Windows-31J (Japanese) character encoding. This is the Microsoft version of Shift_JIS.

euc-jp specifies EUC-JP (Extended Unix Code - Japanese) character encoding.

ksc5601 is KSC5601 (Korean) character encoding.

iso-8859-1 is ISO 8859-1 (Latin1, single-byte) character encoding.

Default(s)

ISO 8859-1 (Latin1)

Guidelines

This command determines the character encoding used to register the virtual server's NetBIOS computer name (called the wins-name) with its wins server. This should be the encoding supported by the network's WINS server. This also determines the encoding for all of the virtual server's wins-alias names, as well as the server's group (or domain) name.

The show global server *fqdn* command shows the full WINS configuration for the *fqdn* service.

Samples

```
bstnA(gbl-cifs[ac1.medarch.org])# wins-name-encoding shift-jis
```

sets the character encoding to Shift_JIS for the 'ac1.medarch.org' service. Whenever the virtual server registers its NetBIOS names with its wins server, it translates to Shift_JIS for the registration.

```
bstnA(gbl-cifs[ac1.medarch.org])# no wins-name-encoding
```

resets the character encoding for NetBIOS names to the default.

Related Commands

wins-name
wins-alias
show global server

25

Schedules

description (gbl-schedule)

Purpose

You can use the optional description command to set a descriptive string for the current schedule. This appears in show commands.

Use the no form of the command to delete the description.

Mode

gbl-schedule

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
description text
no description
```

text (1-256 characters) is your description. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s)

`no description`

Guidelines

The description appears in the output for show schedule.

Sample

```
bstnA(gbl-schedule[backupWindow])# description "regular backup times"
```

specifies a description for the current schedule.

Related Commands

schedule
show schedule

duration

Purpose

Use the duration command to limit the amount of time that a rule can run each time the schedule fires. (A *schedule* is applied to a policy rule, to determine when and how frequently the rule should fire.)

Use no duration to allow the rule to run indefinitely every time.

Mode

gbl-schedule

Security Role(s)

storage-engineer or crypto-officer

Syntax

`duration` *`HH:MM`*`[:00]`
`no duration`

***HH:MM*[:00]** is the duration allowed for any rule on this schedule (for example, 00:30:00 for 30 minutes). The only accepted value for seconds is 00 (zero).

Default(s)

no duration - no limit on the rule's duration

Guidelines

Use this command to impose a limit on all rules that use the current schedule. If (during any run) the rule exceeds this time limit, the rule stops processing. For example, if you set a 2-minute duration for a place-rule's schedule and one run of the rule exceeds the limit, that run stops. This stops any volume scan in progress, cancels all migrations, and causes the policy engine to start recording files or directories that clients change (if they now match the placement rule's fileset, they are queued for migration when the schedule next fires). New files and directories, which are directed to a share before they are created, are unaffected by an expired duration.

Samples

```
bstnA(gbl-schedule[fridays])# duration 00:15
```

sets a 15-minute duration for the "fridays" schedule.

```
bstnA(gbl-schedule[daily])# no duration
```

removes any time limit for rules that run on the "daily" schedule.

Related Commands

schedule

every

Purpose Use the every command to determine the current schedule's frequency.

Mode gbl-schedule

Security Role(s) storage-engineer or crypto-officer

Syntax

```
every count {minutes | hours | days | weeks | months | quarters
| years}
every {sunday | monday | tuesday | ... saturday}+
every first {sunday | monday | tuesday | ... saturday}
every last {sunday | monday | tuesday | ... saturday}
every day-list day-of-month[,day-of-month]+
every hour-list hour:minute[,hour:minute]+
```

count **{minutes | ... | years}** creates a simple schedule with a single time period between each time that the schedule fires. For example, **every 2 weeks** or **every 5 hours**.

> ***count*** (1-4294967295) is the number of minutes, hours, days, etc. between scheduled events.
>
> **minutes | ... | years** is a required choice.

sunday | ... | saturday creates a weekly schedule that fires on the given day or days of the week.

first sunday | ... | saturday makes a monthly schedule that fires on the first given day of the month.

last sunday | ... | saturday creates a monthly schedule that fires on the last given day of the month.

day-list results in a schedule that runs on one or more days of the month. For example, **day-list 10,20** creates a schedule that fires on the 10th and the 20th of each month. If you enter 31, the policy engine fires on the last day of the month, whether or not the month has 31 days.

hour-list makes a schedule that runs on one or more times of the day. Use a 24-hour clock to specify the hours. For example, **hour-list 04:00,12:00,23:00** fires three times per day: it fires at 4 AM, noon, and 11 PM.

Default(s) `every 1 weeks`

Guidelines To set a start time and/or date for the schedule, use the start command. If you choose an every command that does not include a time (such as every first tuesday), the start time determines the time of day that the schedule fires.

Samples

```
bstnA(gbl-schedule[hourly])# every 1 hours
```
sets an hourly interval.

```
prtlndA(gbl-schedule[payday])# every tuesday thursday
```
makes a weekly schedule that fires every Tuesday and Thursday.

```
prtlndA(gbl-schedule[quarterly])# every 3 months
```
sets a quarterly interval.

```
stkbrgA(gbl-schedule[threeDaily])# every hour-list 03:00,13:00,22:00
```
sets a schedule that fires three times a day: at 3 AM, 1 PM, and 10 PM.

```
bstnA(gbl-schedule[monthy_tues])# every first tuesday
```
makes a monthly schedule that fires on the first Tuesday of each month.

Related Commands

schedule
start

schedule

Purpose

Some policy rules, such as the shadow-copy rule, operate on a schedule. Use the schedule command to create a new schedule.

Use no schedule to delete a schedule.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
schedule name
no schedule name
```

name (1-64 characters) is the name you choose for the schedule.

Default(s)

None.

Guidelines

Some rules operate in discrete sessions instead of continuously, so they use a schedule to determine when to operate. For example, a shadow-copy rule copies its fileset from the source volume to the shadow volume; each shadow-copy session occurs on a schedule.

This command places you in gbl-schedule mode, where you define the period for the schedule (every 3 hours, every day, every 2 weeks, or some other period) with the every command. You can also define an optional start time and/or date; the default start time is the moment that you create the schedule. If you want to impose a time limit on rules each time they run, you can use the duration command. The optional description (gbl-schedule) command assigns a descriptive string to the schedule. To set an expiration date for the schedule, use the stop command.

Use the schedule (gbl-ns-vol-plc) command to assign a schedule to a file-placement rule. Use the schedule (gbl-ns-vol-shdwcp) command to assign a schedule to a shadow-copy rule. The schedule (gbl-ns-vol-...snap) command assigns a schedule to a snapshot rule. To assign a schedule to auto-diagnostics collection, which monitors usage trends on the chassis, you can use the schedule (gbl-auto-diag) command.

Samples

```
bstnA(gbl)# schedule hourly
bstnA(gbl-schedule[hourly])#
```

creates a schedule named “hourly.”

```
prtlndA(gbl)# no schedule every_2_months
```

deletes a schedule.

Related Commands

every
start
stop
duration
description (gbl-schedule)
schedule (gbl-ns-vol-plc)
schedule (gbl-ns-vol-shdwcp)
schedule (gbl-ns-vol-...snap)
schedule (gbl-auto-diag)

show schedule

Purpose Use the show schedule command to view all schedules.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax `show schedule [schedule-name]`

schedule-name (optional, 1-64 characters) focuses the command on a single schedule. If a name is not specified, all schedule names are displayed.

Guidelines The following fields appear for each schedule:

Schedule is the schedule's configured name, set by the schedule command.

Description only appears if set by the gbl-schedule description (gbl-schedule) command.

Start Time can be reset by the gbl-schedule start command. This is the date and time of day used for all time-based calculations: for example, if a schedule runs every Wednesday, this determines the time of day on Wednesdays when the schedule fires.

Stop Time only appears if set by the gbl-schedule stop command. This is the expiration time for the schedule, after which the schedule never fires.

Interval is set by the gbl-schedule every command. This is the time between scheduled runs.

Duration appears only if set by the gbl-schedule duration command. This is the amount of time allotted for migrations or replications before a scheduled rule goes into a *paused* state (similar to the one invoked by the policy pause (gbl-ns-vol) command.) A 0 (zero) means that there is no limit on the duration: each time a rule runs on this schedule, it can run indefinitely.

Status is "Paused" (the schedule is between run times), "Running," or "Waiting for start time" (the schedule has never been invoked).

Previous is a sub-table showing the Run time (when the run started) and End time of the most-recent run.

Current appears if the rule is currently running. This has the same fields as the Previous table, Run time and End time.

Next shows when the schedule's next run will begin.

Sample `bstnA# show schedule`

shows all schedules. See Figure 25.1 on page 25-8 for sample output.

Related Commands show policy
show policy filesets

Figure 25.1 *Sample Output: show schedule*

```
bstnA# show schedule

  Schedule:         hourly
```

```
  Start Time:       Wed Jun  8 01:11:00 2011
  Interval:         1 hours
  Status:           Running (runs in 00:13:17)

  Previous:
    Run Time:         Wed Jun  8 01:11:00 2011 (Schedule Start)
    End Time:         N/A

  Next:
    Run Time:         Wed Jun  8 02:11:00 2011
    End Time:         N/A


Schedule:         daily4am
  Description:      two hours between 4 and 6 AM
  Start Time:       Sun Sep  4 04:00:00 2005
  Stop Time:        Wed Jan  7 04:00:00 2015 (Expires in 1309 d 03:02:17)
  Interval:         1 days
  Duration:         02:00:00
  Status:           Paused (runs in 02:02:17)

  Previous:
    Run Time:         Tue Jun  7 04:00:00 2011
    End Time:         Tue Jun  7 06:00:00 2011

  Next:
    Run Time:         Wed Jun  8 04:00:00 2011
    End Time:         Wed Jun  8 06:00:00 2011


Schedule:         backupWindow
  Description:      regular backup times
  Start Time:       Sun Nov 12 13:00:00 2006
  Interval:         1 days
  Duration:         04:00:00
  Status:           Paused (runs in 11:02:17)

  Previous:
    Run Time:         Tue Jun  7 13:00:00 2011
    End Time:         Tue Jun  7 17:00:00 2011

  Next:
    Run Time:         Wed Jun  8 13:00:00 2011
    End Time:         Wed Jun  8 17:00:00 2011


Schedule:         weekly
  Start Time:       Sat May  6 02:00:00 1995
  Interval:         1 weeks
  Status:           Running (runs in 3 d 00:02:17)

  Previous:
    Run Time:         Sat Jun  4 02:00:00 2011
    End Time:         N/A

  Next:
    Run Time:         Sat Jun 11 02:00:00 2011
    End Time:         N/A
```

start

Purpose Use the start command to set a start time for the current schedule. A *schedule* is applied to a policy rule, to determine when and how frequently the rule should fire.

Use no start to eliminate the fixed start time from the schedule; this resets the start time to now.

Mode gbl-schedule

Security Role(s) storage-engineer or crypto-officer

Syntax

```
start [mm/dd/yyyy:HH:MM[:00]]
no start
```

***mm/dd/yyyy*:*HH:MM*[:00]** (optional) is the date and time to start the schedule (for example, **05/06/2005:06:00**). The only accepted value for seconds is 00 (zero). If you omit this, you reset the start time to now and start the schedule immediately.

Default(s) now

Guidelines This command sets a starting point for all scheduling calculations. For example, if you have a daily schedule, this determines the time of day. If you have a weekly schedule, this also determines the day of the week.

If you set a start date/time in the future, the schedule waits until that time.

A no start resets the date/time to now, but does not fire the schedule.

Samples

```
bstnA(gbl-schedule[daily])# start 11/12/2004:23:00
```
sets a start date of November 12, 2004, at 11PM.

```
prtlndA(gbl-schedule[weekly])# start
```
starts the weekly schedule now.

```
prtlndA(gbl-schedule[hourly])# no start
```
resets the hourly schedule's start time to now, but does not invoke any rules.

Related Commands schedule

stop

Purpose

The stop command establishes an optional end date for the current schedule. A *schedule* is applied to a policy rule, to determine when and how frequently the rule should run. The schedule stops functioning after this date and time; any rule on the schedule stops all future runs.

Use no stop to eliminate any end date for the schedule.

Mode

gbl-schedule

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
stop [mm/dd/yyyy:HH:MM[:00]]
no stop
```

***mm/dd/yyyy:HH:MM*[:00]** (optional) is the date and time to stop the schedule (for example, **10/24/2022:18:00**). The only accepted value for seconds is 00 (zero). If you omit this, you stop the schedule now, so that rules using the schedule will never run in the future.

Default(s)

None

Guidelines

This command sets an expiration date on the entire schedule. To limit the amount of time allotted to each scheduled run, use the duration command. You can use the start command to determine the first day and time that the schedule runs.

If you set a stop date/time in the past, the schedule never runs.

Sample

```
bstnA(gbl-schedule[daily4am])# stop 01/07/2011:04:00
```

sets a stop date of January 7, 2011, at 4 AM. The schedule will never fire again after that date and time.

Related Commands

schedule
start
duration

26

Filesets

This chapter contains an alphabetical list of commands for configuring a *fileset*, a group of files in a managed volume. You can use the commands in this chapter to create filesets based on file name, age, directory path, and various other criteria.

age-fileset

Purpose Use the age-fileset command to create a simplified age-based fileset for the current volume.

Mode gbl-ns-vol

Security Level storage-engineer or crypto-officer

Syntax

```
age-fileset name
no age-fileset name
```

name (1-1024 characters) is the name you choose for the fileset.

Default(s) None

Guidelines This command creates a fileset that is only usable in the current volume. To create a global fileset, usable from any namespace or volume, use policy-age-fileset in gbl mode.

The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.) This places you in gbl-ns-vol-fs-age mode, where you use several commands to define the fileset. You must use the select-files (...-fs-age) command to choose files older than (and/or newer than) 1 hour, 2 days, 20 minutes, or some other interval; this establishes the age. To determine whether the age is based on the last time the files were accessed or the last time they were modified, use the last command.

An age fileset selects its files each time a rule uses it. The rule's assigned schedule controls when it uses the fileset; each time the rule's schedule fires, it assesses every file's age from that moment. That is, "one hour old" or "one day old" is measured at the moment that the rule runs. The schedule (gbl-ns-vol-plc) command assigns a schedule to a file-placement rule, and the schedule (gbl-ns-vol-shdwcp) command assigns one to a shadow-copy rule.

Sample

```
bstnA(gbl-ns-vol[medarcv~/rcrds])# age-fileset 24hrs
This will create a new policy object.

Create object '24hrs'? [yes/no] yes
bstnA(gbl-ns-vol-fs-age[medarcv~/rcrds~24hrs])# ...
```

defines the '24hrs' fileset.

Related Commands

policy-age-fileset
select-files (...-fs-age)
last
schedule (gbl-ns-vol-plc)
schedule (gbl-ns-vol-shdwcp)

every (...-fs-age)

Purpose

** Deprecated **

Use the every command to determine the how frequently the current age fileset re-selects its files.

Use no every to select files only when a policy rule (such as a place-rule or a shadow-copy-rule) invokes the fileset.

Modes

gbl-fs-age

gbl-ns-vol-fs-age

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
every count {minutes|hours|days|weeks|months|quarters|years}
every {sunday|monday|tuesday| ... |saturday}
no every
```

count (1-2,147,483,647) is the number of minutes, hours, days, etc. between file selections. For example, **every 2 minutes** or **every 5 minutes**.

minutes| ... years is a required choice.

sunday| ... saturday causes the fileset select its files weekly, on the given day of the week.

Default(s)

None; select files whenever the fileset is used by a policy rule (such as a place-rule or a shadow-copy-rule).

Guidelines

This command is deprecated in favor of controlling age selection from the rule's schedule. Consider instead using the schedule (gbl-ns-vol-plc) command (for a file-placement rule) or the schedule (gbl-ns-vol-shdwcp) command (for a shadow-copy rule). These assign a schedule to the rule. Every time the assigned schedule fires, the rule reassesses the age-based fileset from the scheduled time. For example, if the schedule fires on 11/12, a month-old file is one that was modified before 10/12; if the schedule fires again in 11/13, a month-old file is one that was modified before 10/13. This command creates a different schedule for file assessment. To continue the example, all "month-old files" may remain the ones modified before 10/12 for a long time, as determined by the algorithm described below.

In most-all cases, it is more intuitive to let the rule's schedule determine all file selections.

Guidelines: File-Selection Behavior

This command determines how frequently the policy engine selects files for the current age fileset. At each interval, the policy engine re-selects its files based on the current time. For example, consider a fileset with files "older than 1 week" that selects its files every 1 day:

1. Saturday: select all files that were modified before last Saturday.
2. Sunday: select all files that were modified before last Sunday.
3. Monday: select files modified before last Monday.
4. ...

If the same fileset was evaluated every 2 days, the progression would look like this:

1. Saturday: select all files that were modified before last Saturday.
2. (Sunday: do nothing; keep the selection of all files that were modified before last Saturday.)
3. Monday: select files modified before last Monday.
4. ...

The time of day when the files are selected is the start time for the fileset. This is either the time of day when the fileset is invoked by a rule or the fixed time/date set with the start command.

The no every command resets this command and start at the same time: it makes the fileset select its files whenever a rule invokes the fileset.

Samples

```
bstnA(gbl-fs-age[hourOld])# every 1 hours
```

sets an hourly selection interval.

```
bstnA(gbl-fs-age[sixMonths])# no every
```

reverts to the default; select files whenever the rule runs (as controlled by its schedule). This is considered a best practice.

Related Commands

age-fileset
start

filename-fileset

Purpose Use the filename-fileset command to create a name-based fileset for the current namespace volume.

Use the no form of the command to delete a name-based fileset.

Mode gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
filename-fileset name
no filename-fileset name
```

name (1-1024 characters) is a required name that you assign to the fileset.

Default(s) None

Guidelines This command creates a fileset that is only usable in the current volume. To create a global fileset, usable from any namespace or volume, use policy-filename-fileset in gbl mode.

The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.) This places you in gbl-ns-vol-fs-name mode where you can define the fileset by choosing file names (with name), one or more directory paths within the volume, and/or an option to match files recursively (recurse). By default, a new file-name fileset matches all files in the volume's root directory, but not below it.

Samples

```
bstnA(gbl-ns-vol[ns~/])# filename-fileset logs
This will create a new policy object.

Create object 'logs'? [yes/no] yes
```

creates a name-based fileset, "logs."

```
bstnA(gbl-ns-vol[ns~/])# no filename-fileset logs
```

deletes a name-based fileset, "logs."

Related Commands

policy-filename-fileset
name
path
recurse

filesize-fileset

Purpose Use the filesize-fileset command to create a size-based fileset for the current volume.

Mode gbl-ns-vol

Security Level storage-engineer or crypto-officer

Syntax

```
filesize-fileset name
no filesize-fileset name
```

name (1-1024 characters) is the name you choose for the fileset.

Default(s) None

Guidelines This command creates a fileset that is only usable in the current volume. To create a global fileset, usable from any namespace or volume, use policy-filesize-fileset in gbl mode.

The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.) This places you in gbl-ns-vol-fs-filesize mode, where you use the select-files (...-fs-filesize) command to choose the size range for the fileset's files.

Sample

```
bstnA(gbl-ns-vol[medarcv~/rcrds])# filesize-fileset big
This will create a new policy object.

Create object 'big'? [yes/no] yes
bstnA(gbl-ns-vol-fs-filesize[medarcv~/rcrds~big])# ...
```

defines the 'big' fileset.

Related Commands policy-filesize-fileset
select-files (...-fs-filesize)

from fileset (gbl-...-fs-...)

Purpose Use the from fileset command to include a component fileset into the current union or intersection fileset.

Use the no from command to remove a component fileset from the current fileset.

Modes gbl-fs-isect

gbl-ns-vol-fs-isect

gbl-fs-union

gbl-ns-vol-fs-union

Security Role(s) storage-engineer or crypto-officer

Syntax

```
[no] from fileset fileset-name
no from [all]
```

fileset-name (1-1024 characters) is the name of an existing source fileset.

all (optional) removes all source filesets from the union or intersection fileset.

Default(s) None

Guidelines This command pulls other filesets into the current, compound fileset.

The no from all command removes all source filesets at once.

Samples

```
bstnA(gbl-fs-union[bulky])# from fileset fm_pdf
bstnA(gbl-fs-union[bulky])# from fileset veryLarge
```

adds two filesets to a union fileset named “bulky.” (The “bulky” fileset is defined in gbl-fs-union mode, so you can use it in any volume.)

```
prtlndA(gbl-ns-vol-fs-union[nemed~/acctShdw~test])# no from all
```

removes all source filesets from a union fileset named “test.” (The “test” fileset is defined within the /acctShdw volume, so it is limited to rules in that volume.)

Related Commands policy-intersection-fileset
intersection-fileset
policy-union-fileset
union-fileset

intersection-fileset

Purpose

Use the intersection-fileset command to create an intersection fileset in the current namespace volume. An intersection fileset contains all files that are common to all of its source filesets.

Use the no form of the command to delete the intersection fileset.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

`intersection-fileset` ***`name`***
`no intersection-fileset` ***`name`***

name (1-1024 characters) is the name you choose for the fileset.

Default(s)

None

Guidelines

This command creates a fileset that is only usable in the current volume. To create a global-intersection fileset, usable from any namespace or volume, use policy-intersection-fileset in gbl mode.

The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.)

An intersection fileset contains files that are common to two or more source filesets. This command places you in gbl-ns-vol-fs-isect mode; from there, use the from fileset (gbl-...-fs-...) command to identify each source fileset.

Samples

```
bstnA(gbl-ns-vol[ns~/vol])# intersection-fileset xmlViaFtp
This will create a new policy object.

Create object 'xmlViaFtp'? [yes/no] yes
bstnA(gbl-ns-vol-fs-isect[ns~/vol~xmlViaFtp])#
```

creates an intersection fileset.

```
bstnA(gbl-ns-vol[archives~/etc])# no intersection-fileset isectTst
```

removes an intersection fileset.

Related Commands

policy-intersection-fileset
from fileset (gbl-...-fs-...)

last

Purpose Use the last command to set the type of timestamp to use for age-based grouping: last-accessed time or last-modified time.

Modes gbl-fs-age

gbl-ns-vol-fs-age

Security Role(s) storage-engineer or crypto-officer

Syntax `last {accessed | modified}`

accessed | modified is a required choice:

accessed is the last-accessed (read) time.

modified is the last-modified (written) time.

Default(s) `last modified`

Guidelines The last-accessed time is the time that the file was last read. The last-modified time is the time that the file was last written.

Sample

```
bstnA(gbl-ns-vol-fs-age[ns~/vol~weeklyHttpLogs])# last accessed
```

sets last-accessed times.

Related Commands age-fileset

name

Purpose

Use the name command to select a file name or names for the current name-based fileset.

Use no name to erase the selection, thereby matching any file name.

Modes

gbl-fs-name

gbl-ns-vol-fs-name

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
name exact [not] file-name [ignore-case]

name match [not] wild-card-string [ignore-case]

name regexp [not] "regular-expression" [ignore-case]
no name
```

not (optional) negates the match; this causes the fileset to include all files in the fileset's path *except* any files that match.

exact *file-name* (1-1024 characters) is an exact file name to match.

ignore-case (optional) tells the rule to match files without paying attention to case (for example, **name exact index.html ignore-case** matches "index.html" and "Index.HTML").

match *wild-card-string* (1-1024 characters) uses a wild-card pattern to match multiple file names. See the *Guidelines* below for your wildcard options.

regexp "*regular-expression*" (1-1024 characters) uses a more-complex match pattern, based on the ICU standard for regular expressions. This also requires the quotation marks. See the *Guidelines* below for regular-expression syntax.

Default(s)

match every file in the fileset's path

Guidelines

You can use the path command (for example, path exact /var/log) to select a particular directory. You can optionally use the recurse to descend from a path into its subdirectories.

The no form of this command makes the fileset match all files in the fileset's path(s).

Guidelines: Wildcard-String Options

The ***wild-card-string*** uses the following Unix-like shell conventions:

***** is any string of 0 (zero) or more characters, including an empty string.

? is any single character, or no character.

[...] matches any one of the enclosed characters. For example, [xyz] matches x, y, or z.

[a-z] matches any character in the sorted range, a through z.

[^...] or **[!...]** matches any character that is *not* enclosed. For example, [!xyz] matches any character *except* x, y, or z.

Guidelines: Regular Expressions

The ***regular-expression*** syntax follows IBM's ICU conventions. You can use this expression to be extremely specific about your match criteria. At a high level, these are the conventions for ICU regular expressions:

.* matches any string, including the null string.

.? matches any single character, or no character.

[...] matches any one of the enclosed characters.

[a-z] matches any character in the sorted range, a through z.

[^...] matches any character that is *not* enclosed.

\d matches any digit, 1-9.

\s matches any white-space character.

**** matches the next character, even if that character is a special character (for example, **\?** is a literal "?"; it is not interpreted).

^... and **...$** match the beginnings and ends of lines, respectively.

expression1* | *expression2 matches both expressions.

For more details, refer to Regular Expression Syntax, on page 13-11 of the ARX® CLI Storage-Management Guide. Also, you can search the World-Wide Web for instructions on ICU regular expressions.

Samples

```
bstnA(gbl-fs-name[xmlFiles])# name match *.xml*
```

matches all files with ".xml" in their names (such as "addresses.xml" and "testData.xmllint.output").

```
bstnA(gbl-fs-name[xmlFiles])# name regexp "\.xml$"
```

matches all files whose names end with ".xml" (that is, "addresses.xml" matches but "testData.xmllint.output" does not).

```
bstnA(gbl-fs-name[noHTML])# name match not .htm?
```

matches all files except those with ".htm" or ".htm*x*" (where *x* could be any single character) in their names.

```
bstnA(gbl-fs-name[httpLogs])# no name
```

matches all files in this fileset's path.

Related Commands

policy-filename-fileset
filename-fileset
path
recurse

offline

Purpose Some back-end CIFS servers use Hierarchical Storage Management (HSM) systems to archive file data on a remote server. A small portion of the file remains on the CIFS server while the bulk of the file is stored remotely. The small, local file has an *offline* attribute set. Whenever a client (such as the ARX) accesses the file, the HSM restores the full file contents to the CIFS server.

A large-scale file migration could restore a large number of offline files all at once, possibly enough to run out of disk space on the back-end CIFS server. You can use the offline command to choose files based on the setting for their offline attribute. You can then exclude files with a raised offline attribute from a place-rule migration.

Mode gbl-fs-cifs-attr

Security Role(s) storage-engineer or crypto-officer

Syntax

```
offline {set | not-set | ignore}
no offline
```

set | not-set | ignore is a required choice:

set indicates that files in this fileset are all offline. We do not generally recommend migrating files (see place-rule) with the offline attribute set. The file's behavior is dependent on the back-end file server and the HSM; they may or may not restore the full file first. Additionally, assuming that the migration causes the file to be restored first, the change in the file may trigger a higher-priority rule and cancel the migration.

not-set matches only files that are online.

ignore matches all files, no matter how the offline attribute is set.

no offline is equivalent to **offline ignore**.

Default(s) `no offline`

Guidelines The offline attribute indicates whether or not the file has the bulk of its data stored on a remote server. Use this command in a policy-cifs-attributes-fileset to select all files in a volume based on this attribute setting.

Sample

```
bstnA(gbl-fs-cifs-attr[online])# offline not-set
```

selects files whose offline attribute is lowered. These files are at full size on the back-end CIFS server.

Related Commands policy-cifs-attributes-fileset

path

Purpose

Use the path command to choose one or more directories for the current named-based fileset.

This allows you to narrow the fileset scope to one or more directories inside the namespace volume. Only matching files from this directory are included in the fileset.

Use no path to revert to the default path.

Modes

gbl-fs-name

gbl-ns-vol-fs-name

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
path exact [not] dir-path [ignore-case]

path match [not] wild-card-string [ignore-case]

path regexp [not] "regular-expression" [ignore-case]
no path
```

not (optional) negates the match; this causes the fileset to search all paths in the volume *except* any paths that match.

dir-path (1-1024 characters) is an exact path to match.

ignore-case (optional) tells the rule to match paths without paying attention to case (for example, **path exact /src ignore-case** matches "/src" and "/SRC").

match *wild-card-string* (1-1024 characters) uses wild cards to match multiple paths. See the *Guidelines* below for your wildcard options.

regexp "*regular-expression*" (1-1024 characters) uses a more-complex match pattern, based on the ICU standard for regular expressions. This also requires the quotation marks. See the *Guidelines* below for regular-expression syntax.

Default(s)

/ (but no subdirectories) in the volume that uses this fileset.

Guidelines

This command specifies one or more directories for a name-based fileset. Only files in these directories are eligible for inclusion in this fileset. Use the name command to choose the files themselves.

By default, this does not recurse; use the recurse command to include all matching subdirectories in the fileset, too. Note that the match and regexp options often create a recursive effect: for example, the **path match "/var/*"** command matches "/var/," "/var/log/," "/var/log/httpd/," "/var/www/," and any other subdirectory under "/var/."

The **path exact *dir-path*** syntax identifies a single directory, and does not require a full volume scan. All other options (including **not** and **ignore-case**) may match multiple directories, so they require more time for a full scan.

Guidelines: Wildcard-String Options

The ***wild-card-string*** uses the following Unix-like shell conventions:

***** is any string of 0 (zero) or more characters, including an empty string.

? is any single character.

◆ Note

*The ***** and the **?** match any character, including a "/" (the Unix delimiter for directories). Therefore, **path match "/home/*/bin"** matches "/home/juser/bin/" as well as "/home/jrandom/misc/bin/."*

[...] matches any one of the enclosed characters. For example, [xyz] matches x, y, or z.

[a-z] matches any character in the sorted range, a through z.

[^...] or **[!...]** matches any character that is *not* enclosed. For example, [!xyz] matches any character *except* x, y, or z.

Guidelines: Regular Expressions

The ***regular-expression*** syntax follows IBM's ICU conventions. You can use this expression to be extremely specific about your match criteria. At a high level, these are the conventions for ICU regular expressions:

.* matches any string, including the null string.

.? matches any single character.

[...] matches any one of the enclosed characters.

[a-z] matches any character in the sorted range, a through z.

[^...] matches any character that is *not* enclosed.

\d matches any digit, 1-9.

\s matches any white-space character.

**** matches the next character, even if that character is a special character (for example, **\?** is a literal "?"; it is not interpreted).

^... and **...$** match the beginnings and ends of lines, respectively.

expression1 **|** ***expression2*** matches both expressions.

For more details, refer to Regular Expression Syntax, on page 13-11 of the ARX® CLI Storage-Management Guide. Also, you can search the World-Wide Web for instructions on ICU regular expressions.

Samples

```
bstnA(gbl-ns-vol-fs-name[ns~/vol~logs])# path exact /var/http/log
```

sets the current fileset to contain files in /var/http/log.

```
bstnA(gbl-fs-name[httpLogs])# path match /log ignore-case
```

includes any directories with "/log" in their paths, in any case, such as "/var/log" and "/home/juser/log," or "/LOG."

```
bstnA(gbl-fs-name[noHid])# path regexp not "/\.[^\.]"
```

includes all directories unless their names start with "." followed by some other character.

```
bstnA(gbl-fs-name[xslfiles])# no path
```

sets the current fileset to contain files in the volume's root, /.

Related Commands

policy-filename-fileset
filename-fileset
name
recurse

policy-age-fileset

Purpose Use the policy-age-fileset command to create an age-based fileset for any volume.

Mode gbl

Security Level storage-engineer or crypto-officer

Syntax

```
policy-age-fileset name
no policy-age-fileset name
```

name (1-1024 characters) is the name you choose for the fileset.

Default(s) None

Guidelines The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.) This places you in gbl-fs-age mode, where you use several commands to define the fileset. You must use the select-files (...-fs-age) command to choose files older than (and/or newer than) 1 hour, 2 days, 20 minutes, or some other interval; this establishes the age. To determine whether the age is based on the last time the files were accessed or the last time they were modified, use the last command.

An age fileset selects its files each time a rule uses it; you can alter this by setting a regular schedule for file selection. The optional every (...-fs-age) command sets the frequency of these file selections; by the time a rule invokes the fileset, its files are already selected based on an earlier, scheduled time. You can use the start command to set a start date (and time) for this fixed schedule (for example, start the "every 10 minutes" schedule at noon on January 1st). These commands are deprecated, however, in favor of letting the rule's schedule determine file-selection schedule.

This type of fileset is available to all namespaces and volumes. To create a similar fileset in a single volume, use age-fileset in gbl-ns-vol mode.

You cannot remove this global fileset if any volume or rule uses it.

Samples

```
bstnA(gbl)# policy-age-fileset weekOld
This will create a new policy object.

Create object 'weekOld'? [yes/no] yes
bstnA(gbl-fs-age[weekOld])# ...
```

defines the 'weekOld' fileset. Any volume can use this age-based fileset.

```
bstnA(gbl)# no policy-age-fileset lastMonth
bstnA(gbl)# ...
```

removes an age fileset, 'lastMonth.'

Related Commands

select-files (...-fs-age)
last
every (...-fs-age)
start
age-fileset

policy-cifs-attributes-fileset

Purpose Use the policy-cifs-attributes-fileset command to create a fileset based on CIFS attributes. Some CIFS attributes can be important for file migrations, such as the *offline* attribute. The offline attribute indicates that the file's full content is on a remote server, and would need to be fully restored before accessing it for migration. You can use this fileset to select files with particular CIFS attributes and carefully manage their migration.

Mode gbl

Security Level storage-engineer or crypto-officer

Syntax

```
policy-cifs-attributes-fileset name
no policy-cifs-attributes-fileset name
```

name (1-1024 characters) is the name you choose for the fileset.

Default(s) None

Guidelines The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.) This places you in gbl-fs-cifs-attr mode, where you use the offline command to choose files based on the setting for their CIFS *offline* attribute.

If a file has its CIFS offline bit set, the file has been reduced to a stub with most of its contents archived on another server. The CIFS filer first recovers the full contents of the file if the ARX migrates it. This can fill the CIFS filer's disk if it occurs for too many files at once. You can use this fileset to identify all files with this CIFS attribute and migrate them (or avoid migrating them) accordingly.

The offline setting is the only supported CIFS attribute in a CIFS-attribute fileset.

This type of fileset is available to all namespaces and volumes that support CIFS. This does not apply to an NFS-only namespace, and there is no analogous command in gbl-ns-vol mode.

You can combine this fileset with others using policy-union-fileset and policy-intersection-fileset. For example, to select all files with the offline attribute unset and in a particular directory, you can

- define the particular directory in a policy-filename-fileset,
- define a set of online files with this command (policy cifs-attributes-fileset) and offline not-set,
- join the two filesets in an policy-intersection-fileset, and
- use the intersection fileset in the from (gbl-ns-vol-plc) command of a place-rule.

You cannot remove this global fileset if any rule uses it.

Guidelines: Place Rules with this Fileset May Need Schedules The policy engine only finds the offline-file attribute during a scheduled scan of the volume's back-end filers. There are no inline (client-driven) events that indicate that the attribute changed. Therefore, any place-rule that uses this fileset only assesses the attribute once unless it has a schedule (gbl-ns-vol-plc).

Samples

```
bstnA(gbl)# policy-cifs-attributes-fileset online

This will create a new policy object.

Create object 'online'? [yes/no] yes
bstnA(gbl-fs-cifs-attr[online])# ...
```

defines the ‘online’ fileset. Any volume can use this attribute-based fileset.

```
bstnA(gbl)# no policy-filesize-fileset offline_not_set
bstnA(gbl)# ...
```

removes an attribute-based fileset, ‘offline_not_set.’

Related Commands offline

policy-filename-fileset

Purpose Use the policy-filename-fileset command to create a name-based fileset for any volume.

Use the no form of the command to delete a name-based fileset.

Mode gbl

Security Role(s) storage-engineer or crypto-officer

Syntax

```
policy-filename-fileset name
no policy-filename-fileset name
```

name (1-1024 characters) is a required name that you assign to the fileset.

Default(s) None

Guidelines The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.) This command places you in gbl-fs-name mode where you can define the fileset by choosing file names (with name), one or more directory paths within the volume, and/or an option to descend into matching directories recursively (recurse). By default, a new file-name fileset matches all files in the volume's root directory, but not below it.

This type of fileset is available to all namespaces and volumes. To create a similar fileset in a single volume, use filename-fileset in gbl-ns-vol mode.

You cannot remove this global fileset if any volume or rule uses it.

Samples

```
bstnA(gbl)# policy-filename-fileset logs
This will create a new policy object.

Create object 'logs'? [yes/no] yes
```

creates a name-based fileset, "logs."

```
bstnA(gbl)# no policy-filename-fileset logs
```

deletes a name-based fileset.

Related Commands

name
path
recurse
filename-fileset

policy-filesize-fileset

Purpose Use the policy-filesize-fileset command to create a size-based fileset for any volume.

Mode gbl

Security Level storage-engineer or crypto-officer

Syntax

```
policy-filesize-fileset name
no policy-filesize-fileset name
```

name (1-1024 characters) is the name you choose for the fileset.

Default(s) None

Guidelines The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.) This places you in gbl-fs-filesize mode, where you use the select-files (...-fs-filesize) command to choose files larger-than (and/or smaller-than) a size of your choosing.

This type of fileset is available to all namespaces and volumes. To create a similar fileset in a single volume, use filesize-fileset in gbl-ns-vol mode.

You cannot remove this global fileset if any volume or rule uses it.

Samples

```
bstnA(gbl)# policy-filesize-fileset midsize
This will create a new policy object.

Create object 'midsize'? [yes/no] yes
bstnA(gbl-fs-filesize[midsize])# ...
```

defines the ‘midsize’ fileset. Any volume can use this size-based fileset.

```
bstnA(gbl)# no policy-filesize-fileset 1_to_4_gig
bstnA(gbl)# ...
```

removes a filesize fileset, ‘1_to_4_gig.’

Related Commands select-files (...-fs-filesize)
filesize-fileset

policy-intersection-fileset

Purpose Use the policy-intersection-fileset command to create an intersection fileset. An intersection fileset contains all files that are common to all of its source filesets.

Use the no form of the command to delete the intersection fileset.

Mode gbl

Security Role(s) storage-engineer or crypto-officer

Syntax

```
policy-intersection-fileset name
no policy-intersection-fileset name
```

name (1-1024 characters) is the name you choose for the fileset.

Default(s) None

Guidelines The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.)

An intersection fileset contains files that are common to two or more source filesets. This command places you in gbl-fs-isect mode; from there, use the from fileset (gbl-...-fs-...) command to identify each source fileset.

This command creates a fileset that is available to all namespaces and volumes. To create a similar fileset in a single volume, use intersection-fileset in gbl-ns-vol mode.

You cannot remove this global fileset if any volume or rule uses it.

Samples

```
bstnA(gbl)# policy-intersection-fileset xmlViaFtp
This will create a new policy object.

Create object 'xmlViaFtp'? [yes/no] yes
bstnA(gbl-fs-isect[xmlViaFtp])#
```

creates an intersection fileset.

```
bstnA(gbl)# no policy-intersection-fileset isectTst
```

removes an intersection fileset.

Related Commands from fileset (gbl-...-fs-...)
intersection-fileset

policy-union-fileset

Purpose Use the policy-union-fileset command to create a union fileset. A union fileset is the union of two or more component filesets.

Use the no form of the command to remove the union fileset.

Mode gbl

Security Role(s) storage-engineer or crypto-officer

Syntax

```
policy-union-fileset name
no policy-union-fileset name
```

name (1-1024 characters) is the name you choose for the fileset.

Default(s) None

Guidelines A union fileset joins two or more source filesets into one. If a file belongs in any of the source filesets, it is also included in the union fileset.

The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.) This places you in gbl-fs-union mode; from there, use the from fileset (gbl-...-fs-...) command to identify each source fileset.

This command creates a fileset that is available to all namespaces and volumes. To create a similar fileset in a single volume, use union-fileset in gbl-ns-vol mode.

You cannot remove this global fileset if any volume or rule uses it.

Sample

```
bstnA(gbl)# policy-union-fileset allLogs
This will create a new policy object.

Create object 'allLogs'? [yes/no] yes
bstnA(gbl-fs-union[allLogs])#
```

creates a union fileset.

Related Commands from fileset (gbl-...-fs-...)
union-fileset

recurse

Purpose Use the recurse command to include all subdirectories in the current file-name fileset.

Use no recurse to include only files in the top directory.

Modes gbl-fs-name

gbl-ns-vol-fs-name

Security Role(s) storage-engineer or crypto-officer

Syntax

```
recurse
no recurse
```

Default(s)

```
no recurse
```

Guidelines Use this command to set or unset recursion. If recursion is set, the file-name fileset selects directories and/or files below the root directory. The path command chooses directory paths, and the name command chooses files.

If you use recurse in a fileset without any path or name specification, this selects all files in the volume.

Sample

```
bstnA(gbl-fs-name[logs])# recurse
```

includes files in all subdirectories.

Related Commands policy-filename-fileset
filename-fileset
name
path

select-files (...-fs-filesize)

Purpose A size-based fileset chooses a group of files based on their sizes. Use the select-files command to choose the sizes for the current fileset.

Modes gbl-fs-filesize

gbl-ns-vol-fs-filesize

Security Level storage-engineer or crypto-officer

Syntax

```
select-files {larger-than-or-equal-to | smaller-than}
size[k|M|G|T]
no select-files {larger-than-or-equal-to | smaller-than}
```

larger-than-or-equal-to | smaller-than is a required choice. This determines whether the selected files are larger or smaller than the size that follows.

size is the size in bytes, KiloBytes, MegaBytes, or whatever units you specify (see below).

k|M|G|T (optional) specifies units; **k** for KiloBytes (1024 bytes), **M** for MegaBytes (1024*1024 bytes), **G** for GigaBytes (1024*1024*1024), and **T** for TeraBytes (1024*1024*1024*1024). Do not put any space between the size number and this unit (for example, **100k** is valid but **100 k** is not). The default is bytes.

Default(s) None

Guidelines The select-files command is the only required step in configuring a size-based fileset.

You can use this command twice to define a size range (for example, to select files larger than 50 MegaBytes but smaller than 1 Gigabyte). The CLI rejects a select-files command that contradicts an existing select-files command: you cannot select files "smaller-than 1 MegaByte" if you have already selected files "larger-than or equal-to 2 GigaBytes."

Sample

```
bstnA(gbl-ns-vol-fs-filesize[midsize])# select-files
larger-than-or-equal-to 20M
bstnA(gbl-ns-vol-fs-filesize[midsize])# select-files smaller-than 50M
```

selects all files that are between 20 and 50 MegaBytes.

Related Commands policy-filesize-fileset
filesize-fileset

select-files (...-fs-age)

Purpose An age-based fileset chooses a single group of files based on time stamps: those newer or older than a certain interval (such as "older than 1 day") or between two intervals (for example, "newer than 2 weeks and older than 1 day"). Use the select-files command to set the selection criteria.

Modes gbl-fs-age

gbl-ns-vol-fs-age

Security Level storage-engineer or crypto-officer

Syntax

```
select-files {older-than | newer-than} interval {minutes |
hours | days | weeks | months | quarters | years}

no select-files {older-than | newer-than}
```

older-than | **newer-than** is a required choice. This determines whether the selected files are older or newer than the time interval that follows.

interval (1-2,147,483,647) is the number of minutes, hours, days, etc., that selected files are older than or newer than.

minutes | **... years** is a required choice.

Default(s) None

Guidelines The select-files command is the only required step in configuring an age-based fileset. Without any selection criteria, the fileset includes all files.

You can use this command twice to define an age range (for example, to select files older than 1 week but newer than 1 year). The CLI rejects a select-files command that contradicts an existing select-files command: you cannot select files "newer than 6 hours" if you have already selected files "older than 3 weeks."

Each file has two time stamps: last-accessed time and last-modified time. This command measures against the time stamp selected with the last command.

By default, this selects from every file in the current managed volume. To limit the files to those in some other fileset (for example, to choose only from a filename-fileset containing .wmv files), use the from fileset (gbl-...-fs-...) command.

By default, the policy engine runs this file selection whenever it is invoked by another rule (such as a place-rule or a shadow-copy-rule). For example, if the fileset is composed of files newer than 1 month and a file-placement rule uses it on September 30, it includes all files modified since August 31. If the file-placement rule runs again on October 5, the fileset re-selects its files: at this point, it holds all files modified since September 5, 1 month before October 5. You can use the every (...-fs-age) command to make the file re-selections happen at a fixed rate (for example, every 2 days).

Sample

```
bstnA(gbl-ns-vol-fs-age[ns~/vol~24hrs])# select-files older-than 24
hours
```

identifies all files that are older than 24 hours in the age-based fileset.

Related Commands policy-age-fileset
age-fileset

show policy filesets

Purpose Use the show policy filesets command to view fileset configurations.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show policy filesets
show policy filesets global-fileset fileset-name
show policy filesets namespace namespace volume volume
  fileset fileset-name
```

global-fileset ***fileset-name*** (optional, 1-1024 characters) focuses the command on a single, global fileset.

namespace ***namespace*** **volume** ***volume*** **fileset** ***fileset-name*** focuses on a single, volume-level fileset.

namespace (1-30 characters) is the fileset's namespace,

volume (1-1024 characters) is the fileset's volume path, and

fileset-name (1-1024 characters) is the fileset..

Guidelines If you do not enter a specific fileset, this command displays all filesets. The globally-defined filesets appear first, followed by a section for each namespace and volume. Each volume-level fileset (such as a age-fileset) appears with its volume.

For each fileset, a single Configuration table shows all of the administratively-set parameters for the fileset. You can create or edit each fileset type with a separate command:

- "Filename Fileset" with policy-filename-fileset,
- "Age Fileset" with policy-age-fileset,
- "Fileset Intersection" with policy-intersection-fileset, or
- "Fileset Union" with policy-union-fileset.

The remaining fields show the configuration parameters for the fileset, and the number of rules that use the fileset.

Samples

```
bstnA# show policy filesets
```

shows all filesets. See Figure 26.1 for sample output.

```
bstnA# show policy filesets global-fileset veryLarge
```

shows one globally-defined fileset, "veryLarge." See Figure 26.2 on page 26-31 for sample output.

```
bstnA# show policy filesets namespace insur volume /claims fileset
images
```

shows a single fileset in the "insur~/claims" volume.

Related Commands policy-filename-fileset
policy-age-fileset
policy-intersection-fileset
policy-union-fileset

Figure 26.1 *Sample Output: show policy filesets*

```
bstnA# show policy filesets
Global Policy:

  Filename Fileset:  website

    Configuration:
      Name Does Not Match Regular Expression:        \.(wmv|avi)$
      Path Is:                                       /www/xml/
      Recurse:                                       Yes
      Rules Using This Fileset:                      0


  Fileset Union:    bulky

    Configuration:
      From fileset:                                  fm_pdf
      From fileset:                                  veryLarge
      Rules Using This Fileset:                      1


  Filename Fileset:  hiddenFiles

    Configuration:
      Name Is:
      Path Matches Regular Expression:               /\.[^\.]
      Recurse:                                       No
      Rules Using This Fileset:                      0


  Filename Fileset:  fm_pdf

    Configuration:
      Name Matches Regular Expression:               \.(fm|pdf)$ (case ignored)
      Recurse:                                       Yes
      Rules Using This Fileset:                      1


  Filename Fileset:  xmlFiles

    Configuration:
      Name Matches Pattern:                          *.xml
      Recurse:                                       Yes
      Rules Using This Fileset:                      0


  File Size Fileset:  veryLarge

    Configuration:
      Select Files Larger Than:                      5.0 MB
      Rules Using This Fileset:                      1


  Age Fileset:      dayOld
```

```
    Configuration:
      Select files older than:                    1 days
      Mode:                                       Last Accessed
      Rules Using This Fileset:                   1


  Filename Fileset:  allDirs

    Configuration:
      Name Is:
      Recurse:                                    Yes
      Rules Using This Fileset:                   2


  Age Fileset:      modThisMonth

    Configuration:
      Select files newer than:                    1 months
      Mode:                                       Last Modified
      Rules Using This Fileset:                   1


  Age Fileset:      modEarlier

    Configuration:
      Select files older than:                    1 months
      Mode:                                       Last Modified
      Rules Using This Fileset:                   1


  CIFS Attributes Fileset:  online

    Configuration:
      Offline:                                    Not set
      Rules Using This Fileset:                   1


  Fileset Intersection:    onlineDayOld

    Configuration:
      From fileset:                               online
      From fileset:                               dayOld
      Rules Using This Fileset:                   1


  Age Fileset:      2mo

    Configuration:
      Select files newer than:                    2 months
      Mode:                                       Last Modified
      Rules Using This Fileset:                   0


Volume:               /vol

Volume:               /acct

Volume:               /rcrds

Volume:               /acopia$ns3
```

```
Volume:                /lab_equipment

Volume:                /test_results

Volume:                /claims

  Filename Fileset:  images

    Configuration:
      Name Is:
      Path Is:                                    /images/
      Recurse:                                    Yes
      Rules Using This Fileset:                   0


Volume:                /acopia$ns4
```

Figure 26.2 *Sample Output: show policy filesets global-fileset veryLarge*

```
bstnA# show policy filesets global-fileset veryLarge

Global Policy:

  File Size Fileset:  veryLarge

    Configuration:
      Rules Using This Fileset:                   1
      Select Files Larger Than:                   5.0 MB
```

start

Purpose

** Deprecated **

For age-based filesets, use the start command to set a start time for grouping files.

Use no start to select files each time a rule uses the fileset.

Modes

gbl-fs-age

gbl-ns-vol-fs-age

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
start [mm/dd/yyyy:HH:MM:SS]
no start
```

mm/dd/yyyy sets a start date (for example, **date 05/25/2003**).

HH:MM:SS sets a start time (for example, **06:00:00**).

Default(s)

Whenever the fileset is used by a policy rule (such as a place-rule or a shadow-copy-rule).

Guidelines

This command is deprecated in favor of controlling age selection from the rule's schedule. Consider instead using the schedule (gbl-ns-vol-plc) command (for a file-placement rule) or the schedule (gbl-ns-vol-shdwcp) command (for a shadow-copy rule). These assign a start time and a schedule to the rule. In most-all cases, it is more intuitive to let the rule's schedule determine start times and file selections.

Guidelines: Effects on Time Calculations

This command sets a starting point for all age-based calculations. For example, if you have a fileset that re-evaluates every 7 days, this determines the starting point and the time of day for each evaluation (Mondays at 2AM, Fridays at 5PM, or some other day and time).

The evaluation chooses files that belong in the fileset. By default, this occurs at the time the fileset is used by a rule.

If you set a start date/time in the future, the fileset waits until that time to start grouping files.

Sample

```
bstnA(gbl-fs-age[monthOld])# start 05/06/2003:23:00:00
```

sets a start date of May 6, 2003, at 11PM.

Related Commands

age-fileset
every (...-fs-age)

union-fileset

Purpose

Use the union-fileset command to create a union fileset in the current namespace volume. A union fileset is the union of two or more component filesets.

Use the no form of the command to remove the union fileset.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

`union-fileset` ***`name`***
`no union-fileset` ***`name`***

name (1-1024 characters) is the name you choose for the fileset.

Default(s)

None

Guidelines

A union fileset joins two or more source filesets into one.

This command creates a fileset that is only usable in the current volume. To create a global-union fileset, usable from any namespace or volume, use policy-union-fileset in gbl mode.

The CLI prompts for confirmation before creating a new fileset; enter **yes** to continue. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.) This places you in gbl-ns-vol-fs-union mode; from there, use the from fileset (gbl-...-fs-...) command to identify each source fileset.

Sample

```
bstnA(gbl-ns-vol[ns~/vol])# union-fileset allLogs
This will create a new policy object.

Create object 'allLogs'? [yes/no] yes
bstnA(gbl-ns-vol-fs-union[ns~/vol~allLogs])#
```

creates a union fileset.

Related Commands

policy-union-fileset
from fileset (gbl-...-fs-...)

27

Place Rules

Policy commands govern file migration to balance load and capacity among back-end filers.

enable (gbl-ns-vol-plc)

Purpose Use the enable command to enable the current rule. A new rule is disabled by default, and the policy engine ignores disabled rules.

Use no enable to disable the current rule.

Modes gbl-ns-vol-plc

Security Role(s) storage-engineer or crypto-officer

Syntax
```
enable
no enable
```

Default(s) `no enable`

Guidelines You must enable a rule for the policy engine to use it.

Sample
```
bstnA(gbl-ns-vol-plc[archives~/multimedia~PDFs])# enable
```
enables the current file-placement rule, “PDFs.”

Related Commands namespace -> volume -> place-rule

from (gbl-ns-vol-plc)

Purpose

A file-placement rule moves files onto chosen storage. Use the from command to select some source files to move.

Use no from to remove the source fileset, which effectively disables the rule.

Mode

gbl-ns-vol-plc

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
from fileset fileset-name [match files]
from fileset fileset-name match directories
[promote-directories]
from fileset fileset-name match all [promote-directories]
no from
```

fileset-name (1-1024 characters) is the name of an existing source fileset.

match files (optional) matches the fileset against files only, ignoring all directories. This is the default.

match directories (optional) matches the fileset against directories only, ignoring all files.

promote directories (optional) causes all migrated directories to be promoted at the target filer. Any new files or subdirectories in these directories will also go to the target filer (unless redirected by another rule).

match all (optional) matches both files and directories.

Default(s)

match ... defaults to matching files only.

Guidelines

Use this command to select particular files and/or directories to migrate. Use the source command to select only from a particular share or share farm. The placement rule migrates all matching files/directories to its target storage. Use the no form of the command to remove the fileset.

This command skips files with multiple Unix hard links. To migrate hard links off of a share, use the source command instead of this command. Alternatively, you can use the the source command and this command together with migrate hard-links.

If a single fileset is too restrictive, use the union-fileset command to join multiple filesets into one.

The from fileset syntax allows you to match the fileset against files, directories, or both.

Guidelines: Matching Files

By matching files only, you can redistribute files with certain names (such as *.mpg) to a share or a share farm. A files-only match can also implement tiering: an age-fileset can select files in a certain age group to migrate to a certain tier (that is, a certain share or share farm). The volume creates replica directories on the destination share, called directory *stripes*, to hold the files.

Guidelines: Matching Directories

Each directory has one *master* copy on the share where the volume first discovered it. A directory can have no stripes or up to one stripe per remaining share in the volume. As stated above, the volume typically creates stripes to hold migrated files. If you use the from fileset syntax to match against directories only or all (both files and directories), you can use the promote-directories flag to promote the stripe on the target share to master.

To make all new directory trees grow on the target filer, use the match directories option without promote-directories. This steers all ensuing client-created directories to the target share. These new directories will exist at the target share first, where they will be master. Directory trees under those new directories will follow them to the target share. Files under the pre-existing directories will tend to stay on their source filers, since their masters remain on the source filers.

You can make whole directory trees grow at the target filer with match directories and promote-directories. This creates master copies of all matching directories at the target filer. This does not migrate any existing files, but it does steer all new files and subdirectories to the target filer. This can be useful if a nsck ... rebuild operation has scattered a directory tree's mastership among multiple shares.

Important

Do not match and promote directories based on their last-accessed times in a CIFS or multi-protocol volume; a CIFS filer changes the directory's time stamp whenever the time stamp itself is read.

Guidelines: Matching Files and Directories

To move an existing directory tree to the target filer(s) and make it grow there, use match all and promote-directories.

Guidelines: no from

The no from command removes the only source fileset from the rule, effectively disabling it if there is no source to drain. The CLI prompts for confirmation before doing this; enter **yes** to proceed.

Sample

In gbl-ns-vol-plc mode, the following example selects the "doc_files" source fileset for the "copytoNas206" placement rule. It matches the fileset against files only:

```
bstnA(gbl-ns-vol[wwmed~/acct])# place-rule copytoNas206
bstnA(gbl-ns-vol-plc[wwmed~/acct~copytoNas206])# from fileset doc_files
bstnA(gbl-ns-vol-plc[wwmed~/acct~copytoNas206])# ...
```

Related Commands

place-rule
target

inline notify

Purpose

By default, a file-placement rule finds its source files by monitoring all new files as they are created, monitoring client changes as they happen (inline), *and* scanning the volume for existing files. Use the no inline notify command to disable inline-change notifications and work with new and scanned files only.

Use the inline notify command to re-enable inline notifications for the current rule.

Mode

gbl-ns-vol-plc

Security Level

storage-engineer or crypto-officer

Syntax

```
no inline notify
inline notify
```

Default(s)

```
inline notify
```

Guidelines

Inline notifications happen as clients change files. Whenever a file changes so that it conforms to the source fileset, the file-placement rule migrates the file as needed. The migration is queued immediately, not the next time the rule runs. This command can disable or re-instate inline notifications to the current rule.

Whenever inline notifications are disabled, changed files do not migrate until the next time that the rule scans the volume. If no future volume scan is scheduled, the files remain indefinitely. For this reason, no inline notify is only recommended for a rule that has an assigned schedule (gbl-ns-vol-plc).

A tiered-storage configuration has multiple place rules, typically one per tier, with different recommendations for this setting. The placement rule(s) that send files to Tier 1 should have inline notify enabled. The rule(s) that send files to other tiers should use no inline notify and have an assigned schedule. This recommendation ensures that files migrate to tier 1 as soon as they are qualified for that tier, but they do not move to a lesser tier until a scheduled run of a file-placement rule. This reduces the computation load on the policy engine.

This command has no effect on placement of new files and/or directories; the rule causes these entries to be created on the target storage, so no migrations are necessary.

Volume scans are controlled by the volume-scan command.

Sample

```
bstnA(gbl-ns-vol-plc[ns3~/vol~new2shr5])# no inline notify
```

disables inline notifications to the “new2shr5” rule.

Related Commands

place-rule
schedule (gbl-ns-vol-plc)
volume-scan

inline report

Purpose An *inline* event is one that occurs between scheduled runs of the file-placement rule. Clients can make inline changes to files that make them eligible for migration by the file-placement rule. Use this command to enable regular inline-migration reports for the current file-placement rule, to track any inline migrations that may occur.

Use no inline report to prevent inline-migration reports.

Mode gbl-ns-vol-plc

Security Role(s) storage-engineer or crypto-officer

Syntax

```
inline report {hourly|daily} prefix [verbose]
  [delete-empty|error-only]

no inline report
```

prefix (1-1024 characters) sets a prefix for all inline-migration reports. Each report has a unique name in the following format:

prefix_YearMonthDayHourMinute.rpt

For example, "tier1-inline_200903031200.rpt" could be the name for one report with the "tier1-inline" prefix.

verbose (optional) enables verbose data in the reports.

delete-empty | error-only (optional) are mutually exclusive.

delete-empty causes the rule to delete any reports that contain no migrated files or errors.

error-only causes the rule to delete any reports that contain no errors.

Default(s) `no inline report`

Guidelines This command creates hourly or daily reports to track the current rule's inline migrations. These reports can consume a great deal of internal disk space over time; you can use the delete-empty or error-only flags to conserve on this space.

Use show reports for a list of reports, or show reports *file-name* to show the contents of one report.

Samples

```
bstnA(gbl-ns-vol-plc[wwmed~/acct~docs2das8])# inline report hourly
docsPlc verbose
```

enables hourly inline-migration reports for a file-placement rule, "docs2das8." See Figure 27.1 on page 27-8 for sample output.

```
bstnA(gbl-ns-vol-plc[ns3~/usr~mp3Off])# no inline report
```

disables inline-migration reports for the "mp3Off" rule.

Related Commands place-rule
show reports

Figure 27.1 *Sample Report: Inline Migration*

```
bstnA# show reports docsPlc_201002270055.rpt
**** Inline File Placement Report: Started at Sat Feb 27 00:55:15 2010 ****
**** Software Version: 5.02.000.12541 (Feb 23 2010 20:12:44) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

  Place Rule:       docs2das8

    Configuration:
      Namespace:                                    wwmed
      Volume Name:                                  /acct
      From fileset:                                 bulky (files only)
      Target share:                                 bills
      Report:                                       docsPlc
      Report Verbose:                               Enabled
      Report Delete Empty:                          Disabled
      Report Errors Only:                           Disabled
      Inline Report:                                docsPlc
      Inline Report Interval:                       hourly
      Inline Report Verbose:                        Enabled
      Inline Report Delete Empty:                   Disabled
      Inline Report Errors Only:                    Disabled
      Migrate limit:                                50G
      Volume Scan:                                  Enabled
      Inline Notifications:                         Enabled
      Promote Directories:                          Disabled
      Auto-Close Files:                             Disabled

      Tentative:                                    No
      State:                                        Enabled


  Date                        Source Share     Target Share
      File
      Result
  ------------------------------------------------------------------------------
  Sat Feb 27 01:18:21 2010   budget           bills
      /acct/receipts/jun2005.csv
      Complete
  Sat Feb 27 01:19:16 2010   budget           bills
      /acct/spreadsheets/cleanup.pdf
      Complete
  Sat Feb 27 01:19:16 2010   budget           bills
      /acct/images/legal/classA_ITE_label.pdf
      Complete
  Sat Feb 27 01:19:16 2010   budget           bills
      /acct/images/sampleNetAllIPs.fm
      Complete
  Sat Feb 27 01:21:44 2010   bills2           bills
      /acct/wksheets/ADMIN.HTML
      Complete

    Inline Statistics:
      Inline Interval Start:                        Sat Feb 27 00:55:15 2010
      Inline Interval End:                          Sat Feb 27 01:55:15 2010
      Inline Files Migrated:                        5
      Inline Failed Migrations:                     0
      Inline Size of Failed Migrations:             0 B
      Inline Files Renamed:                         3
      Inline Files Placed:                          39
      Inline Directories Renamed:                   0
```

```
      Inline Directories Placed:                    0
      Inline Directories Promoted:                  0
      Inline Failed Directory Promotes:             0
      Inline Suppressed Error Messages:             0

Total processed:              5
Elapsed time:          01:00:00
**** Inline File Placement Report: DONE at Sat Feb 27 01:55:15 2010 ****
```

limit-migrate

Purpose

A file-placement rule migrates files from one share to another within a managed volume. Use the limit-migrate command to limit the amount of data moved in this migration. If the placement rule runs on a schedule, this limit applies to each run of the rule.

Use no limit-migrate to remove any limit on the size of migrations.

Mode

gbl-ns-vol-plc

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
limit-migrate size[k|M|G|T]
no limit-migrate
```

size (1-18,446,744,073,709,551,615) is the size.

k|M|G|T (optional) is the units:

- **k** is kilobytes (1024 bytes),
- **M** is Megabytes (1024*1024 bytes),
- **G** is Gigabytes (1024*1024*1024 bytes), and
- **T** is Terabytes (1024*1024*1024*1024 bytes).

Default(s)

`no limit-migrate`

size is in bytes.

Guidelines

The policy engine migrates as many files as possible until it reaches one that would reach or exceed this limit. If the placement rule is on a schedule (gbl-ns-vol-plc), the next run of the rule continues migrating any leftover files until it either finishes or reaches the limit again.

If the rule's source is a share (as opposed to a fileset), you can use this limit in conjunction with a schedule in order to gradually drain the share. For example, suppose you want to migrate 1 terabyte of files from a share, but you do not want to slow down switch processing during business hours. A schedule can run the placement rule only during off-hours, and the limit-migrate command can ensure that the placement rule stops after migrating a reasonable amount of storage.

Samples

```
bstnA(gbl-ns-vol-plc[wwmed~/acct~copytoNas206])# limit-migrate 35G
```

sets a migration limit of 35 Gigabytes for the "copytoNas206" rule.

```
bstnA(gbl-ns-vol-plc[archives~/etc~empty])# no limit-migrate
```

removes any migration limit from the "empty" rule. This rule continues to migrate files until they are all on the target.

Related Commands

place-rule
schedule (gbl-ns-vol-plc)

migrate close-file

Purpose

A file-placement rule blocks all CIFS access from a file before it migrates it; this is impossible if another CIFS client already has the file open. You can use the migrate close-file command to permit the rule to automatically close any file opened through CIFS, and hold it closed until it has finished migrating.

Use no migrate close-file to disable the auto-close feature for the current file-placement rule.

Mode

gbl-ns-vol-plc

Security Level

storage-engineer or crypto-officer

Syntax

```
migrate close-file [exclude fileset]
no migrate close-file
```

fileset (optional, 1-1024 characters) is a fileset to exclude from automatic closure. If a file in this fileset is open through CIFS, the rule places it on a retry queue instead of automatically closing it.

Default(s)

```
no migrate close-file
```

Guidelines

This command is only relevant in a managed volume that supports CIFS. It has no effect on files opened through NFS.

If a CIFS client attempts to open a file while it is held closed, the client receives a sharing violation.

You can use the show policy files-closed command to view any files that are currently closed by file-placement rules.

If this feature is disabled, the placement rule cannot migrate any open files. You have two commands that you can use to monitor open files and close them manually: the show cifs-service open-files command shows all such files, and close cifs file closes one from the CLI.

Samples

```
bstnA(gbl-ns-vol-plc[insur~/claims~2old])# migrate close-file
```

allows the "2old" rule to close any open files that it needs to migrate.

```
bstnA(gbl-ns-vol-plc[ns3~/vol~new2shr5])# migrate close-file exclude
homedirs
```

allows the "new2shr5" rule to close any open files *except* files in the "homedirs" fileset.

Related Commands

place-rule
show policy files-closed
show cifs-service open-files
close cifs file

migrate hard-links

Purpose Some Unix files have more than one *hard link*, where each hard link is a different file name (often in a different directory) pointing to the same inode. A file-placement rule typically skips all files with multiple hard links. If your site has many such files and you are migrating a fileset off of a single source share, you can use the migrate hard-links command to migrate all hard links. If one of a file's hard links matches the fileset, this command causes the rule to migrate all of the file's hard links off of the source share.

Use no migrate hard-links to stop migrating any files with multiple hard links.

Mode gbl-ns-vol-plc

Security Level storage-engineer or crypto-officer

Syntax

```
migrate hard-links
no migrate hard-links
```

Default(s)

```
no migrate hard-links
```

Guidelines This only applies to a namespace that supports NFS, where the back-end filers have Unix file systems. This has no effect in a CIFS-only namespace.

The file-placement rule must have a single source share and a fileset (identified with from (gbl-ns-vol-plc)) for this command to function. Additionally, no other place-rule can use the same source share. These rules minimize contention between conflicting file-placement rules, and are enforced by the CLI.

An inline event (see inline notify) is not guaranteed to trigger migration of a multi-hard-link file. To guarantee that all matching multi-link files migrate off of the source share, turn off inline notifications (with no inline notify), assign a schedule (gbl-ns-vol-plc) to the rule, and wait for the next scheduled run of the rule.

If any hard link is outside the source share on the back-end filer, the file-placement rule cannot possibly find that link. The rule skips all such external hard links.

A file-placement rule already migrates all hard links, regardless of this setting, if it is designed to drain a share or share farm. Such a rule has a setting for source but no fileset designated with from (gbl-ns-vol-plc).

Sample

```
bstnA(gbl-ns-vol-plc[insur~/claims~old2new])# migrate hard-links
```

allows the "old2new" rule to migrate all of a file's hard links if any of them match.

Related Commands

place-rule
source
from (gbl-ns-vol-plc)
inline notify

migrate retain-files

Purpose Policy rules can migrate files off of the share; to retain a copy of all migrated files in a hidden directory, use the migrate retain-files command.

Use the no form of the command to stop retaining copies of migrated files.

Modes gbl-ns-vol-shr

Security Role(s) storage-engineer or crypto-officer

Syntax

```
migrate retain-files
no migrate retain-files
```

Default(s) None

Guidelines If this command is enabled, all files migrate off of the current share are also copied into the "~acopia_msnap" directory in the root of the share.

We recommend that you only use this before using remove-share migrate or a place-rule to migrate all files off of the share. You can use the directory tree as a backup in case you are dissatisfied with the results of the migration. To restore the files after a migration, you must access the filer directly (the directory is named so that the volume cannot import it).

Important

Do not use this in a share-farm with a balance or auto-migrate rule enabled. The policy engine may attempt to migrate files off of the share, and this command would prevent file migration from ever changing the free space on the share. This leads to the policy engine continually migrating files to other shares in the share farm.

This feature cannot function if the backing share contains two or more NetApp qtrees.

Sample

```
bstnA(gbl-ns-vol-shr[wwmed~/acct~bills])# migrate retain-files
```

keeps a copy of any file that policy migrates off of the "bills" share.

Related Commands namespace -> volume -> share
show namespace

place-rule

Purpose

A *placement* rule determines where files are stored on back-end storage devices. Use this command to start configuring a file-placement rule.

Use the no form of the command to remove the placement rule.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
place-rule name
no place-rule name
```

name (1-1024 characters) is a name you choose for the placement rule.

Default(s)

None

Guidelines

When you create a new place-rule, the CLI prompts for confirmation. Enter **yes** to create the rule. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.)

Guidelines: Simple Placement

This command places you in gbl-ns-vol-plc mode, where you choose the source files and the destination storage. This rule dynamically places chosen files and/or directories onto chosen storage devices. Use the from (gbl-ns-vol-plc) command to select a source fileset; you have options to determine whether to match the fileset against files, directories, or both, and you can determine whether any matching directories should be promoted to "master" on the target storage. Use the target command to select a destination share farm or share for the files.

To start migrating the files, or to start a "tentative" simulation (described below), use the enable (gbl-ns-vol-plc) command.

This scans all files in the volume (or a subset of source shares), determines which files or directories belong in the fileset, and immediately migrates them to the target share or share farm. It only needs to run once if it places files based on their names or their sizes. Newly-created files are assessed at creation time and placed according to this rule. Files changed "inline" by clients are assessed when they change; if they now match the fileset, they migrate immediately after the change.

You can use the show statistics migration command to track all migration activity.

Guidelines: Scheduling for an Age-Based Fileset (Tiering)

If the placement rule uses a age-fileset, you need a schedule (gbl-ns-vol-plc) to re-assess the fileset as time goes on. Each time the schedule fires, the fileset gathers a new set of files based on that scheduled time. For example, consider a fileset where files are 2 weeks old; on 9/30, those are files modified before 9/16, but on 10/1 the set includes all files modified before 9/17. A daily schedule would expand the set by one day as each day passes. A weekly schedule would expand the set by a full week as each week passes: on 9/30, the "2-week-old" set is anything modified before 9/16, on 10/1 it is still "anything modified before 9/16," and the set remains that way until the schedule fires again on 10/6. On 10/6 the schedule fires and the fileset becomes "anything modified before 9/23" (2 weeks before 10/6).

Use this for creating tiers of storage in your volume. For example, if you want files older than 3 months to move to your Tier-2 shares, use an age-based fileset and a schedule to determine how frequently to re-assess the file ages and move them to the second tier.

Guidelines: Scheduling for a CIFS-Attribute Fileset

Some placement rules use a policy-cifs-attributes-fileset, which checks the setting of the *offline* attribute on each file. If the attribute is set, the file is actually a stub with most of its content archived on another server. The attribute is governed by a Hierarchical Storage Management (HSM) system on the back-end filer, not by client action. The policy engine must scan the back-end filers to detect this attribute setting, and it must re-scan to see if the attribute changed. To make the rule periodically re-scan and discover any changes in this type of fileset, use a schedule (gbl-ns-vol-plc).

Guidelines: Scheduling for Other Filesets

The optional schedule (gbl-ns-vol-plc) command has limited uses for non age-based and non-CIFS-attribute filesets. As mentioned above, these filesets are static, and a file-placement rule migrates them to their target share(s) continuously, as they are created or modified.

You typically need a schedule only for very-large filesets of this type, filesets so large that they require days for the initial migration. For these filesets, you can use the limit-migrate command to limit the size of each migration or the duration command to limit the time.

One other occasion for a schedule is a file-placement rule with a no inline notify setting. A client creates an *inline* change by editing a file or directory. The rule software receives a notification of the change by default, and immediately migrates the changed file if its new name or size places it into the fileset. With inline notifications disabled, the volume may accumulate changed files that break the current placement rule. You can use a schedule to periodically migrate these changed files to their desired share(s).

Guidelines: Configuration Changes in a Scheduled Rule

If you make a configuration change in a rule that is running with a schedule, the configuration change is ineffective until the next time the rule runs. This includes changes to the rule's behavior between scheduled runs, such as a change to inline notify behavior. You can make a change effective immediately by running no enable (gbl-ns-vol-plc) followed by enable on the rule.

Guidelines: Draining a Share or Share Farm

You can drain all files and directories from a share or share farm by specifying a source share or share farm without using the from (gbl-ns-vol-plc) command to specify particular files or directories. After draining all of the files off of a share, you can wait for all of the volume's snapshot rules (snapshot rule) to age out, and then you can remove the shares from the managed volume (remove-share nomigrate or remove-share migrate).

Guidelines: Migrating a Closed File in a CIFS Volume

The migrate close-file command enables the rule to close any file that is opened through CIFS. A rule with this setting can hold the file closed until the migration is finished. Without this setting, the rule cannot migrate an open file if it remains open for the duration of the overall migration run.

Guidelines: Migrating Files with Multiple Hard Links

A file-placement rule typically skips files with multiple Unix hard links. To migrate hard links off of a share, use migrate hard-links together with source and from (gbl-ns-vol-plc). Alternatively, you can use the source command alone (without from (gbl-ns-vol-plc) or migrate hard-links) to drain all files off of the share, including files with multiple hard links.

Guidelines: Other Options

To enable report-generation for each file-placement session (recommended), use the report (gbl-ns-vol-plc) command.

You have the option to set migrate retain-files for the source share(s) before you enable the rule. This keeps copies of all the migrated files in a hidden directory at the root of the share(s). You can use these copies to recover from a failed migration.

You can make the rule simulate a migration with the tentative command. This causes the rule to log all migrations to the syslog as "tentative," without actually migrating any files. This is useful for gauging the effects of a potential file migration.

To wait for a particular migration to finish (for example, as part of a CLI script), you can use the wait-for migration command.

Guidelines: Minimum Free Space on Shares

You can use the policy freespace command to determine the amount of free space to maintain on a given share. If the place rule encounters a file that would break this restriction, the rule pauses until the share's free space rises back up to a *resume-migrate* threshold (also set with the policy freespace command). The volume software probes the back-end share for available free space every 15 seconds. If some other rule migrates files away from the share until its free space rises to the *resume* level, the place rule continues migrating files onto the share.

If the rule encounters a file big enough to fill the share while the share is still above the *resume* level, the rule skips the file and tries another file. If the next file fits, the place rule migrates the file. Otherwise, the place rule skips that file, too. This continues until the share's free space drops below the *resume* level, or until the rule runs out of files to migrate. Once the free space is below the *resume* level, the migrations proceed as described above: a large enough file causes the rule to pause and wait for the free space to rise back up to the *resume* level.

Guidelines: Multi-Protocol Restrictions

A directory in a multi-protocol (NFS and CIFS) share cannot migrate if it has an NFS-only name; the volume cannot access the directory through CIFS to replicate its CIFS attributes. The nsck ... report inconsistencies ... multi-protocol command finds all of the NFS-only names in a multi-protocol volume. You can access the volume through an NFS export and rename the NFS-only directory, or you can temporarily turn off strict-attribute-consistency at the destination share(s).

Guidelines: Rule Order for Fileset Placement

It is possible to configure multiple fileset-placement rules that, to some extent, contradict one another. For example, one rule could place all .xml files onto Share A, while another places all files that have not been accessed in over a month onto share B. If a .xml file exists that has not been accessed in over a month, both placement rules are in contention for that file. By default, the first-configured rule places the file. You can view and change this rule order: use the show policy namespace command to see the current rule order, and use the policy order-rule command to change it.

Sample

```
bstnA(gbl-ns-vol[wwmed~/acct])# place-rule copytoNas206
This will create a new policy object.

Create object 'copytoNas206'? [yes/no] yes
bstnA(gbl-ns-vol-plc[wwmed~/acct~copytoNas206])# ...
```

instantiates a new placement rule, “copytoNas206,” in the current namespace volume.

Related Commands

filename-fileset
from (gbl-ns-vol-plc)
source
target
migrate close-file
migrate hard-links
schedule (gbl-ns-vol-plc)
limit-migrate
report (gbl-ns-vol-plc)
migrate retain-files
tentative
enable (gbl-ns-vol-plc)
strict-attribute-consistency
policy order-rule
show statistics migration

policy freespace

Purpose Whenever a rule migrates a file to a share, it checks the share's free space against two thresholds. Each share has a threshold of free space to *maintain*, as well as another (typically larger) threshold of free space called a *resume* threshold. If the file would reduce the share's free space below the *maintain* level, the rule waits until the free space rises to the *resume* level before continuing. Use the policy freespace command to set these free-space thresholds for the current share.

Use no policy freespace to return to the defaults.

Mode gbl-ns-vol-shr

Security Role(s) storage-engineer or crypto-officer

Syntax

```
policy freespace maintain{k|M|G|T} resume-migrate
resume{k|M|G|T}

policy freespace percent maintain-pct resume-migrate resume-pct

no policy freespace
```

maintain is the amount of free space to maintain on this share.

> **k|M|G|T** chooses the unit of measure: **k**ilobytes, **M**egabytes, **G**igabytes, or **T**erabytes. A kilobyte is 1,024 bytes, a megabyte is 1,024 kilobytes (1,048,576 bytes), and so on..

resume is a free-space level for resuming migrations to this share. The policy engine uses this value if the share has reached it's *maintain* value; no rule can migrate to the share until its free space rises back to the *resume* level.

percent (optional) indicates that you are using disk-space percentages for these values instead of specific size measures.

maintain-pct (1-100) expresses the *maintain* value as a percentage of the overall share size.

resume-pct (1-100) is the *resume* threshold, expressed as a percentage of the share's total space.

Default(s) ***maintain*** - 1G

resume - 2G

Guidelines

A rule that is migrating files to this share (such as a place-rule, or a share farm's auto-migrate directive) checks each file before migrating it. If the file would reduce the share's free space below the level that the share is supposed to maintain, the rule pauses and sets its status to "Target Full." The rule determines its next steps based on the share's current level of free space.

- If the share's space is below the *resume* level, the rule stops migrating. Instead, the rule waits for the share's free space to rise back up to its resume-migrate level. Some other rule must migrate files off of the share to achieve this. The volume software probes the back-end share for free-space every 15 seconds.
- If the share's space is already above the *resume* level, the rule does not wait for the share's free space to rise; the share's free space is already high enough to resume migrating to it. The rule skips the file and moves on to the next one. If the next file fits onto the share while maintaining the required free space, the rule migrates the file. Otherwise, the rule skips the file and moves on to another file. The rule continues migrating and/or skipping files until the share's free space is below its resume level, or until it runs out of files to migrate. If the share's space falls below the resume level, it follows the process above.

Detailed output from show policy shows the number of files skipped along with their total disk space. If the migrating rule is a place-rule with verbose reports enabled (report (gbl-ns-vol-plc)), the report lists all of the skipped files.

A share farm also follows the share's free-space thresholds. If a share is at or below the free space that it should maintain, it cannot receive new files from a balance rule.

To run this command on every share in a namespace, volume, or share farm, you can use the macro version of this command in one of their CLI modes. See the documentation for policy freespace (gbl-ns, gbl-ns-vol) or policy freespace (gbl-ns-vol-sfarm).

Samples

```
bstnA(gbl-ns-vol-shr[medarcv~/rcrds~rx])# policy freespace 4G
resume-migrate 5G
```

makes the policy engine maintain at least 4 Gigabytes of free space on the "rx" share. If that limit is reached by clients or approached by a migration rule, the share is ineligible for any more migrations until its free space rises back up to 5 Gigabytes.

```
bstnA(gbl-ns-vol-shr[medarcv~/rcrds~bulk])# no policy freespace
```

reverts the "bulk" share to the default thresholds for free space.

Related Commands

namespace -> volume -> share
place-rule
show policy
policy freespace (gbl-ns, gbl-ns-vol)
policy freespace (gbl-ns-vol-sfarm)

policy freespace (gbl-ns, gbl-ns-vol)

Purpose Whenever any rule migrates a file to a share, it checks the share's free space against two thresholds. Each share has a minimum free space to *maintain*, as well as another (typically larger) level of free space called a *resume* threshold. If the file would reduce the share's free space below the *maintain* level, the rule waits until the free space rises to the *resume* level before migrating any more files to the share. The policy freespace command sets these free-space thresholds for an ARX share. Use this policy freespace command, from gbl-ns or gbl-ns-vol mode, to set the same thresholds for every share in the current namespace or volume.

Use no policy freespace to return to the defaults for every share in the current namespace or volume.

Modes gbl-ns

gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
policy freespace maintain{k|M|G|T} resume-migrate resume{k|M|G|T}

policy freespace percent maintain-pct resume-migrate resume-pct

no policy freespace
```

maintain is the amount of free space to maintain on these shares.

> **k|M|G|T** chooses the unit of measure: **k**ilobytes, **M**egabytes, **G**igabytes, or **T**erabytes. A kilobyte is 1,024 bytes, a megabyte is 1,024 kilobytes (1,048,576 bytes), and so on..

resume is a free-space level for resuming migrations to any of these shares. The policy engine uses this value if one of the shares has reached it's *maintain* value; no rule can migrate to such a share until its free space rises back to the *resume* level.

percent (optional) indicates that you are using disk-space percentages for these values instead of specific size measures.

maintain-pct (1-100) expresses the *maintain* value as a percentage of the overall share size.

resume-pct (1-100) is the *resume* threshold, expressed as a percentage of the share's total space.

Default(s) ***maintain*** - 1G

resume - 2G

Guidelines This is a macro command for the policy freespace command in gbl-ns-vol-shr mode. This command invokes the individual policy freespace command for every share in the current namespace or volume. The output from show global-config shows the individual share-level policy freespace commands, not this macro command.

This is similar to the policy freespace (gbl-ns-vol-sfarm) command, a macro command for a share farm.

Samples

```
bstnA(gbl-ns-vol[medarcv~/rcrds])# policy freespace percent 5
resume-migrate 6
```

makes the policy engine maintain at least 5-percent free space on every share in the "medarcv~/rcrds" volume. If that limit is reached by clients or approached by a migration rule on any of these shares, the share is ineligible for any more migrations until its free space rises back up to 6 percent.

```
bstnA(gbl-ns[insur])# no policy freespace
```

reverts all shares in the "insur" namespace to the default free-space thresholds.

Related Commands

namespace -> volume -> share
policy freespace
policy freespace (gbl-ns-vol-sfarm)
place-rule
show policy

policy migrate-attempts

Purpose The policy engine tries to migrate each file a limited number of times. Use the policy migrate-attempts command to set the migration-attempt limit for this namespace.

Use no policy migrate-attempts to return to the default.

Modes gbl-ns

Security Role(s) storage-engineer or crypto-officer

Syntax

```
policy migrate-attempts {count | unlimited}
no policy migrate-attempts
```

count (1-1000) is the total number of attempts for a failed migration before declaring that the migration failed.

unlimited (optional) causes the policy engine to infinitely retry its failed migrations.

Default(s) 32

Guidelines The default is adequate for most installations. Change this only on the advice of F5 personnel.

Each file migration fails individually; if one file exhausts its migrate-attempts, its failure does not affect any other files in the same fileset or on the same source share. If the migration runs on a schedule, the migration may succeed on the next scheduled run.

Use the policy migrate-retry-delay command to set the time between retries. If any file is waiting for the policy engine to retry its migration, it is kept in a queue until it successfully migrates; use show policy queue to see all the files in this queue.

Samples

```
bstnA(gbl-ns[medarcv])# policy migrate-attempts 100
```

allows the policy engine to attempt file migrations 100 times in the “medarcv” namespace.

```
bstnA(gbl-ns[wwmed])# no policy migrate-attempts
```

allows 32 attempts for file migrations in the “wwmed” namespace.

```
bstnA(gbl-ns[wwmed])# policy migrate-attempts unlimited
```

allows an infinite number of migration attempts in the “wwmed” namespace.

Related Commands

namespace
policy migrate-retry-delay
show policy queue

policy migrate-delay

Purpose

The policy engine does not migrate a file that has been very-recently modified. This is because clients and client applications tend to perform writes in batches; if a write occurred very recently, a new write is likely to occur very soon. You can set the amount of time since the last write, called the *migration delay*, that a file requires to be eligible for migration. Use the policy migrate-delay command to set the migration delay for this namespace.

Use no policy migrate-delay to return to the default.

Modes

gbl-ns

Security Role(s)

storage-engineer or crypto-officer

Syntax

`policy migrate-delay` ***`seconds`***
`no policy migrate-delay`

seconds (0-1000) is the delay required after the last write. Before this time elapses, the file is ineligible for migration.

Default(s)

30 (seconds)

Guidelines

The default is adequate for most installations. Change this only on the advice of F5 personnel.

If the migration fails because the delay time has not yet elapsed for the file, the policy engine periodically retries. It retries every *n* seconds, where *n* is the number that you set with this command.

Samples

`bstnA(gbl-ns[insur])# policy migrate-delay 60`

stops the policy engine from migrating any file that has been modified 60 seconds ago (or less than 60 seconds ago).

`bstnA(gbl-ns[medarcv])# no policy migrate-delay`

reverts to the default migration-delay time in the “medarcv” namespace.

Related Commands

namespace

policy migrate-method

Purpose The policy engine migrates a file by transferring it directly from the source share to the target directory on the target share. You can instead transfer the file to a hidden staging area on the target share, and then move the file to its final destination after the network transfer is complete. This method allows large files to migrate successfully during one or more ARX-snapshot operations. It also creates a small performance penalty, one that is more noticeable when migrating a large number of small files.

You can use the policy migrate-method direct (or no policy migrate-method) command to choose the direct method, which allows migrations to be severely disrupted by snapshots but offers better performance in some deployment scenarios.

Modes gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
policy migrate-method {direct | staged}
no policy migrate-method
```

direct | staged is a required choice.

direct makes the policy engine migrate all of its files directly to their destinations. If a snapshot occurs in the middle of a direct migration, the migration is cancelled and must be restarted from the beginning on any later migration attempt. If the file is large enough to require a very long migration time, regular snapshots could prevent the file from ever fully migrating. However, direct migrations are sometimes faster than staged migrations, especially in a volume that migrates large numbers of small files.

staged makes the policy engine migrate each file to a hidden staging area at the destination share, and then move the file to its final name and location. This method succeeds while the volume is taking snapshots, with a minor performance penalty.

Default(s) `direct`

Guidelines If you use the direct method, the default, snapshots severely disrupt file migrations. (Snapshots are configured with a snapshot rule in the volume.)

The direct method also creates a larger problem for the following scenario:

- a file is in the middle of migrating, with the original file on the source share and a partial copy of the file on the destination share, and
- the volume has a catastrophic failure, necessitating a full rebuild (similar to a nsck ... rebuild).

After this failure, the file collides with its partial copy during the subsequent re-import. The collision either causes the re-import to fail or it causes one of the file's instances to be renamed (as determined by the reimport-modify setting).

The staged method can create a small performance penalty in a system that migrates large numbers of small files, and it introduces a very small possibility of a client silently losing his or her changes to the ACL on the file's parent directory. The ACL changes can only be lost in a rare set of circumstances.

Sample

```
bstnA(gbl-ns-vol[medarcv~/rcrds])# policy migrate-method staged
```

returns the “medarcv~/rcrds” volume to the default method for migrations, staged.

Related Commands

namespace -> volume
snapshot rule

policy migrate-retry-delay

Purpose

When the policy engine tries and fails to migrate a file (for example, because the file is open on the back-end CIFS share), the engine waits for a short time before it tries the migration again. Use the policy migrate-retry-delay command to set the delay between retries.

Use no policy migrate-retry-delay to return to the default delay.

Modes

gbl-ns

Security Role(s)

storage-engineer or crypto-officer

Syntax

policy migrate-retry-delay *seconds*
no policy migrate-retry-delay

seconds (0-1000) is the number of seconds between migration retries.

Default(s)

900 (seconds, or 15 minutes)

Guidelines

The default is adequate for most installations. Change this only on the advice of F5 personnel.

Use the policy migrate-delay command to set the maximum number of retries. If any file is waiting for the policy engine to retry its migration, it is kept in a queue until it successfully migrates; use show policy queue to see all the files in this queue.

Samples

```
bstnA(gbl-ns[medarcv])# policy migrate-retry-delay 300
```

allows 300 seconds (5 minutes) between migration retries in the “medarcv” namespace.

```
bstnA(gbl-ns[insur])# no policy migrate-retry-delay
```

returns to the default migration-retry delay in the “insur” namespace.

Related Commands

namespace
policy migrate-delay
show policy queue

policy order-rule

Purpose All rules are given a rule priority as they are created, depending on the rule's type and creation time. If two file-placement rules contradict for a given file, the rule with the higher priority is enforced and the other rule is ignored. Use the policy order-rule command to change a rule's priority.

Modes gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
policy order-rule rule1 {before | after} rule2
policy order-rule rule1 {first | last}
```

rule1 (1-1024 characters) is the name of the rule to move.

before | after is a required choice; this determines if *rule1* is ordered before or after *rule2*.

rule2 (1-1024 characters) does not move. This is used to set the new priority for *rule1*.

first | last sets an absolute priority for *rule1*.

Default(s) None

Guidelines The order/priority for most rules is set at creation time, according to the following order:

1. All shadow-copy-rules.
2. All place-rules that use a share (without any fileset) as their source.
3. All place-rules that use a fileset as their source.
4. All share-farms.

Within each group of rules, the individual rules are ordered by their creation times.

You can only reorder the rules in the third category, file-placement rules that use a fileset as their source. The other rules have their order permanently set. Use show policy namespace to see the current rule order.

The moving rule (***rule1***) must go to the other side of the stationary rule (***rule2***), or the order does not change. For example, if you want to move ruleA to be just before ruleD, you must phrase the command this way: **policy order-rule ruleA after ruleC**. Nothing happens if you say **policy order-rule ruleA before ruleD** because ruleA is already before ruleD.

Samples

```
bstnA(gbl-ns-vol[ns~/])# policy order-rule mytest after archiveOldFiles
```

moves the "mytest" file-placement rule after the "archiveOldFiles" rule.

```
bstnA(gbl-ns-vol[ns~/])# policy order-rule wmv2nas3 last
```

moves the "wmv2nas3" file-placement rule to the end of the list.

Related Commands

namespace -> volume
place-rule
shadow-copy-rule
share-farm
show policy

policy pause

Purpose

You can use the policy pause command to stop all volume scans and migrations in a managed volume. This can be useful during a storage maintenance procedure, such as a volume backup.

Use the no form of the command to resume all scans and migrations.

Modes

priv-exec

Security Role(s)

storage-engineer or crypto-officer

Syntax

policy pause namespace *namespace* volume *volume*
no policy pause namespace *namespace* volume *volume*

namespace (1-30 characters) identifies the namespace.

volume (1-1024 characters) is the volume where you want to pause all rules.

Default(s)

None

Guidelines

This command suspends all migrations and filer scans for a volume. While policy is paused, clients may change files and/or directories so that they should be migrated according to the volume's rules; the volume software notes these changes and performs the migrations after you use no policy pause.

A file-placement rule causes all *new* files and directories to be created on their configured back-end shares, without performing any migrations. New files and directories therefore continue to be placed according to your rules while the volume has policy paused.

As an alternative, you can set a schedule for pausing policy in a volume. This is useful for suspending migrations during schedule backup windows, for example. Use the policy pause (gbl-ns-vol) command for scheduled pauses.

Samples

```
bstnA# policy pause namespace insur volume /claims
```

pauses all rules in the "insur~/claims" volume.

```
bstnA# no policy pause namespace insur volume /claims
```

resumes policy processing for the same volume.

Related Commands

policy pause (gbl-ns-vol)

policy pause (gbl-ns-vol)

Purpose

You can use this policy pause command to regularly pause all policy processing in the current volume, according to a fixed schedule.

Use the no form of the command to stop pausing policy on a scheduled basis.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

`policy pause` ***`schedule-name`***
`no policy pause`

schedule-name (1-64 characters) is a schedule to use for pausing all rules.

Default(s)

None

Guidelines

Before you use this command, someone must create a schedule for pausing policy in the current volume. This schedule must have a duration so that policy is not paused indefinitely.

This command assigns a schedule for pausing policy in the current volume. You can use this to suspend all rules on a regular basis, perhaps so that they run during off-hours.

When policy is paused for the volume, the volume suspends all migrations and filer scans. During the pause, clients may change files and/or directories so that they should be migrated according to the volume's rules; the volume software notes these changes and performs the migrations after the duration expires in the schedule.

A file-placement rule causes all *new* files and directories to be created on their configured back-end shares, without performing any migrations. New files and directories therefore continue to be placed according to your rules while the volume has policy paused.

As an alternative, you can use the priv-exec policy pause command to pause policy manually.

Samples

`bstnA(gbl-ns-vol[medarcv~/rcrds])# policy pause midday`
assigns the "midday" schedule to control policy pauses in the "medarcv~/rcrds" volume.

`bstnA(gbl-ns-vol[ns2~/movies])# no policy pause`
disables scheduled pauses for the rules in the "ns2~/movies" volume.

Related Commands

namespace -> volume
schedule
duration
policy pause

policy treewalk-threads

Purpose

A *treewalk* is a full examination of all files and directories in a back-end share. For performance reasons, each of the namespace's volume groups uses a small pool of software threads to perform all of its treewalks. The number of simultaneous treewalks in each volume group is equal to the number of threads in the pool; this can create a bottleneck in volume groups with many managed volumes. Use the policy treewalk-threads command to change the number of threads for each of the namespace's volume groups.

Use the negative form of the command, no policy treewalk-threads, to revert to the default number of threads.

Modes

gbl-ns

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
policy treewalk-threads thread-count
no policy treewalk-threads
```

thread-count (1-10) is the number of threads to use for treewalk operations. This is the thread count for each volume group in the namespace.

Default(s)

4

Guidelines

The default is adequate for most installations. Change this only on the advice of F5 personnel.

Many rules and filesets require treewalks; to preserve CPU cycles, a small pool of threads performs all treewalks for each volume group.

Each of the namespace's volume groups gets this pool of threads, to be shared among all volumes in that group. (You can use the volume-group command to assign a volume to a particular group.)

Filesets that are based on file timestamps (such as the age-fileset) require a treewalk every time a rule invokes one of them. If more than four volumes use these filesets at the same time on the same volume group, only four volumes at a time can perform their treewalks. A fifth volume's treewalk cannot begin until one of the first four treewalks is finished. Two filesets in the same volume can share the results of a single treewalk, but two filesets from different volumes cannot.

You can use this command to expand the thread pools for a namespace with many age filesets, minimizing this performance issue. This change only takes effect the next time you enable the namespace. Use enable (gbl-ns, gbl-ns-vol) to enable the namespace.

Sample

```
bstnA(gbl-ns[wwmed])# policy treewalk-threads 5
```

uses 5 threads for treewalking in the "wwmed" namespace.

Related Commands

namespace
enable (gbl-ns, gbl-ns-vol)
policy-age-fileset
age-fileset

remove namespace ... policy-only

Purpose Use the remove namespace ... policy-only command to remove all policy objects from a namespace or volume.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`remove namespace`** ***`name`*** **`[volume`** ***`volume`*****`] policy-only [sync]`**

name (1-30 characters) identifies the namespace.

volume (optional, 1-1024 characters) focuses on a single volume.

policy-only causes the command to remove rules, filesets, and all other policy objects from the namespace or volume. The volume and namespace configurations remain.

sync (optional) shows the operation's progress at the command line. With this option, the CLI prompt does not return until all policy components have been removed.

Default(s) None.

Guidelines The CLI prompts for confirmation before removing any policy objects; enter **yes** to proceed.

By default, this command generates a report to show all of the actions it takes to remove the volume(s), in order. The CLI shows the report name after you issue the command, and then returns. You can enter CLI commands as the namespace software removes the objects in the background. Use tail to follow the report as it is written. Use show reports file-name to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands. Use the sync option to send the status to the command line instead; the command does not generate a report if you use the sync option.

Use remove namespace to remove an entire namespace or volume, including all policy objects. To remove a namespace and all other configuration objects dedicated to the namespace (including global servers and external filers), use remove service. To remove a share from a volume, use remove-share migrate or remove-share nomigrate. The remove namespace ... volume ... exports-only command finds all front-end exports for a volume and removes them, leaving the volume itself intact.

Sample

```
prtlndA# remove namespace insur_bkup policy-only sync

Remove policy components from namespace 'insur_bkup'? [yes/no] yes
% INFO: Removing service configuration for namespace insur_bkup
% INFO: Removing CIFS browsing for namespace insur_bkup
% INFO: Removing volume policies for namespace insur_bkup
% INFO: destroy policy insur_bkup /insurShdw
```

removes all policy objects from the "insur_bkup" namespace. This uses the sync option, so the progress report appears on the command line instead of a file.

Related Commands

remove namespace
remove service
remove namespace ... volume ... exports-only
remove-share migrate
remove-share nomigrate

report (gbl-ns-vol-plc)

Purpose Use this command to enable progress reports for the current file-placement rule.

Use no report to prevent progress reports.

Mode gbl-ns-vol-plc

Security Role(s) storage-engineer or crypto-officer

Syntax

```
report file-prefix [verbose] [delete-empty|error-only]
no report
```

file-prefix (1-1024 characters) sets a prefix for all file-placement reports. Each report has a unique name in the following format:

prefixYearMonthDayHourMinute.rpt

For example, wmv2sol_200403031200.rpt could be the name for one report with the "wmv2sol_" prefix.

verbose (optional) enables verbose data in the reports.

delete-empty | error-only (optional) are mutually exclusive.

delete-empty causes the rule to delete any reports that contain no migrated files or errors.

error-only causes the rule to delete any reports that contain no errors.

Default(s) no report

Guidelines Every time the file-placement rule fires and migrates files to the target storage, the rule can generate a report to show the details of the migration. The reporting feature is disabled by default; we recommend that you enable reporting for all file-placement rules, to keep track of file migrations between your back-end filers.

The rule only generates a report if at least one existing file is migrated. It does not report on placement of new files, created by clients through the VIP.

Use show reports for a list of reports, or show reports file-name to show the contents of one report.

Samples

```
bstnA(gbl-ns-vol-plc[medarcv~/rcrds~dailyArchive])# report
daily_archive
```

enables reports for the file-placement rule, "dailyArchive." See Figure 27.2 on page 27-35 for sample output.

```
bstnA(gbl-ns-vol-plc[archives~/etc~wmv2sfarm1])# report wmv2sf1 verbose
```

enables verbose reports for the file-placement rule, "wmv2sfarm1."

```
bstnA(gbl-ns-vol-plc[ns3~/usr~mp3Off])# no report
```

disables reports for the "mp3Off" rule.

Related Commands place-rule
show reports

Figure 27.2 *Sample Report: File Placement*

```
bstnA# show reports daily_archive_201206200048.rpt
**** File Placement Report: Started at 06/20/2012 00:48:30 -0400 ****
**** Software Version: 6.03.000.14755 (Jun 15 2012 20:15:33) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

  Place Rule:       dailyArchive

    Configuration:
      Namespace:                                    medarcv
      Volume Name:                                  /rcrds
      From fileset:                                 onlineDayOld (files only)
      Target share:                                 bulk
      Schedule:                                     hourly
      Report:                                       daily_archive
      Report Verbose:                               Disabled
      Report Delete Empty:                          Disabled
      Report Errors Only:                           Disabled
      Migrate limit:                                0
      Volume Scan:                                  Enabled
      Inline Notifications:                         Disabled
      Promote Directories:                          Disabled
      Auto-Close Files:                             Enabled
      Migrate Hard Links:                           Disabled

      Tentative:                                    No
      State:                                        Enabled

    Scan Statistics:
      Scan Started:                                 06/20/2012 00:48:40 -0400
      Scan Completed:                               06/20/2012 00:48:56 -0400
      Elapsed Time:                                 00:00:16
      Number of Times Paused:                       0
      Total Time Paused:                            00:00:00
      Number of Times Stopped by Low Space:         1
      Time Waiting for Free Space:                  00:00:11
      Files Scanned:                                166
      Directories Scanned:                          29
      Files in Fileset:                             160
      Files Migrated:                               160
      Size of Files Migrated:                       11 MB
      Directories Promoted:                         0
      Failed Migrations:                            0
      Size of Failed Migrations:                    0 B
      Failed Directory Promotes:                    0
      Files Forced Closed:                          0
      Suppressed Error Messages:                    0

Total processed:            160
Elapsed time:          00:00:26
**** File Placement Report: DONE at 06/20/2012 00:48:56 -0400 ****
```

schedule (gbl-ns-vol-plc)

Purpose Use this schedule command to assign a schedule to the current file-placement rule.

Use no schedule to remove the rule's schedule.

Mode gbl-ns-vol-plc

Security Role(s) storage-engineer or crypto-officer

Syntax

```
schedule name
no schedule
```

name (1-64 characters) identifies the schedule. Use show policy for a list of configured schedules.

Default(s) None.

Guidelines A file-placement rule migrates a source fileset to a target share or share farm; each file-placement session occurs on a schedule. This command determines which schedule to use.

To create a schedule, use the gbl-mode schedule command.

Samples

```
bstnA(gbl-ns-vol-plc[ns3~/logs~distFiles])# schedule hourly
```

assigns the "hourly" schedule to the placement rule, "distFiles."

Related Commands

place-rule
schedule
show policy

show policy

Purpose Use the show policy command to view policy configurations.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax

```
show policy [details]
show policy namespace namespace [details]
show policy namespace namespace volume volume [details]
show policy namespace namespace volume volume rule rule-name
```

details (optional) changes the output into a detailed view of each rule and share farm. If you omit this, the output is one line per rule or share farm.

namespace (1-30 characters) focuses the command on a single namespace.

volume (1-1024 characters) narrows the scope of the command to one volume in the namespace.

rule-name (1-1024 characters) narrows the scope even further, to one rule or share farm in the volume. This expands the output to show detailed information and statistics about the rule or share farm.

Guidelines This shows all namespace policies. Use show policy to show all globally-defined filesets, and use show schedule to show all globally-defined schedules. To see the history of policy-related events for a volume, rule, or share farm, use show policy history. The show policy queue command shows all files currently waiting to be migrated, if any.

Guidelines: Output The simplest syntax, show policy, outputs tables of the rules and share farms in each volume. Each volume has its own table with the following labels:

Namespace and

Volume.

Each rule or share farm appears on a separate line in the volume's table, with the following fields:

Rule (the name of the rule or share farm),

Type (Place for a place-rule, Share Farm for a share-farm, Snapshot for a snapshot rule, Replica Snapshot for a snapshot replica-snap-rule, Shadow Copy for a shadow-copy-rule, AutoDiagnostics for an auto-diagnostics rule, Config Replication for a config-replication rule, and Notification for a notification rule), and

Status for the rule. For most rules, this is either Enabled or Disabled. For Place, Share Farm, and Shadow Copy, this shows the volume-scan status. For Place and Share Farm, this also shows the current file-migration status.

If you use show policy details, the output shows the full details of every rule and share farm on the system. These details are described below, in the sections about rules, share farms, and volume-level filesets.

Guidelines: Namespaces

The show policy namespace command shows only the rules and share farms in the given namespace. This output contains a Namespace Migration Configuration table, followed by tables of the namespace's rules.

The Namespace Migration Configuration table contains the following fields:

Migrate-Attempts is the number of times that a rule attempts to migrate a file before it declares a failed migration. You can change this with the policy migrate-attempts command.

Migrate-Delay is the number of seconds that a rule waits after a file has changed before the rule attempts to migrate the file. Changes often occur in batches, so this delay prevents repetitive migrations for a file that is being rewritten. You can use the policy migrate-delay command to reset this delay.

Migrate-Retry-Delay is the delay after a failed migration before the rule retries. The policy migrate-retry-delay command controls this setting.

Each volume has a separate table with all of its rules. These tables are similar to the one in the summary output, with the addition of the Rule Priority field. If two rules contradict for a given file, the higher-priority rule is enforced and the other rule is ignored. For example, if Rule 3 migrates a file to share A and Rule 6 migrates the same file to share B, Rule 3 is enforced and the file migrates to share A.

Rules are prioritized in groups as follows: shadow-copy rules (shadow-copy-rule) are the highest priority, followed by file-placement rules (place-rule) that only use a share or share farm as their source (source), followed by file-placement rules that use a source fileset (from (gbl-ns-vol-plc)), followed by share farms (share-farm) at the lowest priority. The only rules that can possibly contradict one another are the third group: file-placement rules that move a fileset (as opposed to draining a share). To change the order of these rules, you can use policy order-rule.

A "drain_share_rule" is created as the by-product of the remove-share migrate command. This is prioritized in the second group (file-placement rules that use a source share) while it is running. The policy engine demotes it to the third group (lower-priority file-placement rules) after the share is removed.

Within their categories, rules are prioritized in the order that they are created.

The details view shows the full details of every rule and share farm in this namespace.

Guidelines: Volumes

The show policy namespace ... volume command focuses on the snapshot configuration, free space, share farms, and rules in a single volume. This is a series of summary tables unless you use the details flag.

Guidelines: Volume Snapshot Configuration

The show policy namespace ... volume command shows the volume-level configuration for snapshots. These settings are only relevant in a volume that has at least one snapshot rule configured. This section contains the following fields:

Point-in-Time Consistency is "Enabled" or "Disabled," depending on whether or not the volume uses snapshot fencing. The VIP fence, if enabled, blocks all client access to the volume while the filers take their coordinated snapshots. Use the snapshot consistency command to allow or disallow this fence.

Management Command Timeout is the number of seconds that the volume software waits for any filer to respond to a command. If this time expires, the command times out. This is almost always 80 seconds.

CIFS Directory Name only appears if the volume supports CIFS. This is the pseudo directory that well-informed CIFS clients (administrators) can use to access their snapshots. You can use the snapshot directory cifs-name command to change this name.

NFS Directory Name only appears if the volume supports NFS. This is the pseudo directory that well-informed NFS clients (administrators) can use to access their snapshots. You can use the snapshot directory nfs-name command to change this name.

Directory Display is "All Exports" (clients see the ~snapshot/.snapshot directory in any front-end CIFS share or NFS export) or "Volume Root" (clients only see the directory only in a front-end share of the volume's root directory), or "None." You can use the snapshot directory display command to change this.

Hidden File Attribute only appears if the volume supports CIFS. This is "Yes" if the special ~snapshot directory has its "hidden" DOS attribute raised. Use an optional argument in the snapshot directory display command to control this setting. This has no effect on NFS clients.

Restricted Access Configured also only appears for a volume that supports CIFS. This is "Yes" if a Windows-Management Authorization (WMA) group restricts the CIFS clients that can access snapshots. Use the snapshot privileged-access command to control this setting. As above, this has no effect on NFS clients.

VSS Mode only appears for a volume that supports CIFS. This field indicates the client-machine versions for which the volume supports the Volume Shadowing Service (*VSS*). VSS is an intuitive interface that clients can use to access their snapshots. This is "Windows XP" (the volume supports VSS for Windows-XP and newer client machines), "Pre-Windows XP" (the volume also supports VSS for Windows-2000 clients), or "None." Direct volumes do not support VSS. For managed volumes, use the snapshot vss-mode command to change this setting. This does not affect NFS clients.

Guidelines: Share Free Space

The show policy namespace ... volume command also shows the free-space thresholds and status from each of its shares. This is a table with the following columns:

Share Name identifies the share in the volume.

Free Space Thresholds shows the two thresholds that you can set with the policy freespace command:

- Maintain is the amount of free space to preserve on the share. The policy engine prevents any migrations from reducing the free space below this value.
- Resume is the amount of free space that is required on the share before it is eligible for migrations. The policy engine uses this value when the share has reached its Maintain value; it pauses all migrations to the share until the share's free space rises back up to this value.

Free Space Status displays the current state of free space on each share:

- Free is the current amount of free space,
- Size is the total size of the back-end share, and
- Pct Free is the percentage of free space remaining on the share.

Guidelines: Share Farms

A section for each share farm appears if you focus on the volume (show policy namespace ... volume) or focus on the share farm itself (show policy namespace ... volume ... rule *share-farm-name*). The output has a table of shares for a share farm, followed by a "New File Placement" rule for the share farm. The first table of shares has the following columns:

Placement Frequency shows the percentage of new files that the balance rule sends to this share. This also shows the numbers used to calculate that percentage.

Freespace Status shows the free space at the share.

The second section describes a "New File Placement" rule with the same name as the share farm. This is an implicit rule for every share farm. Its output is described below.

Guidelines: New File Placement Rules and Place Rules

A section for each file-placement rule appears if you focus on the volume (show policy namespace ... volume) or focus on the rule itself (show policy namespace ... volume ... rule). The output for file-placement rules and new-file-placement rules contains the following tables of settings and statistics:

Configuration shows all of the administrative settings for this rule or share farm.

Status displays the current status of volume scans and file migrations, which happen periodically. This also shows the status of new-file placement, which occurs as clients create new files.

Cumulative Statistics are the numbers of files migrated, migration failures, and migration retries since the share farm was created. This only appears if the rule or share farm is configured to perform migrations. You can use the policy migrate-attempts, policy migrate-delay, and policy migrate-retry-delay commands to control the migration-retry behavior. One field refers to "Inline Overflow" errors, which indicate that the placement rule received more inline-notification events (see inline notify) than it could record in its database; contact F5 Support if any such errors appear.

Queue Statistics counts the files where an initial attempt at migration failed (perhaps because the file was locked), so the file was placed into a queue for a later migration attempt.

Last Scan Statistics describes the last full scan of the volume. This does not appear until the first scan is complete.

Current Scan Statistics describes the current scan of the volume, which may not have started yet or may be currently running. This does not appear while the rule is idle.

Guidelines: Snapshot Rules

A section for each snapshot rule appears if you focus on the volume (show policy namespace ... volume) or focus on the rule itself (show policy namespace ... volume ... rule). Each section describes one snapshot rule. The output contains the following tables:

Configuration shows all of the administrative settings for this snapshot rule.

Archive Configuration appears for a snapshot rule that is recording its volume's configuration (and, typically, its metadata) in a file-history archive. If the snapshot rule regularly records this data, you can query the archive later to find the back-end location of any file at any given time. This is useful for backups.

Cumulative Statistics shows the number of snapshots attempted, successful snapshot runs, failed runs, and the overall success rate. If the snapshot sends its data to a file-history archive, there are also statistics for metadata and volume configurations that were archived.

Last Snapshot Statistics describes the results of the rule's most-recent snapshot. If the snapshot rule sends data to a file-history archive, this also includes archiving results.

Snapshots is a table of the rule's currently-retained snapshots. These are the snapshots that are accessible to the volume's clients. These snapshots are all managed by this snapshot rule, and do not include any snapshots invoked at the filer itself.

Guidelines: Replica Snapshot Rules

A section for each replica-snapshot rule appears if you focus on the volume (show policy namespace ... volume) or focus on the rule itself (show policy namespace ... volume ... rule). Each section describes one snapshot replica-snap-rule. The output contains the following tables:

Configuration shows all of the administrative settings for this replica-snapshot rule.

Cumulative Statistics,

Last Snapshot Statistics, and

Snapshots are the same tables that also appear for standard snapshots. These are described above.

Guidelines: Notification Rules

A section for each notification rule appears if you focus on the volume (show policy namespace ... volume) or focus on the rule itself (show policy namespace ... volume ... rule). Each section describes one notification rule, which takes regular snapshots to support the ARX API. The output contains the following tables:

Configuration shows all of the administrative settings for this notification rule.

Cumulative Statistics,

Last Snapshot Statistics, and

Snapshots are the same tables that also appear for standard snapshots. These are described above.

Guidelines: Shadow-Copy Rules

Each shadow-copy rule has a section in the detailed output, too. This contains the following tables:

Configuration shows all of the administrative settings for this shadow-copy rule.

Status shows the overall status of the most-recent shadow-copy run. A status of "Complete" indicates that the volume copied all files and directories successfully.

Guidelines: Volume Filesets

For each fileset in the volume, a single Configuration table shows all of the administratively-set parameters for the fileset. To see all filesets defined in gbl mode, use show policy.

Samples

`bstnA# show policy`

shows all policy information. See Figure 27.3 on page 27-43 for sample output.

`prtlndA# show policy namespace nemed`

on a different switch, shows policy information for the namespace named "nemed." See Figure 27.4 on page 27-44 for sample output.

`prtlndA# show policy namespace nemed details`

shows details for every rule and share farm in "nemed." See Figure 27.5 on page 27-44 for sample output.

`prtlndA# show policy namespace nemed volume /acctShdw`

shows policy information for the volume named "nemed~/acctShdw." See Figure 27.6 on page 27-45 for sample output.

`bstnA# show policy namespace wwmed volume /acct rule docs2das8`

shows the details for one rule. See Figure 27.7 on page 27-45 for sample output.

Related Commands

show schedule
show policy
show policy history
show policy queue
policy migrate-attempts
policy migrate-delay
policy migrate-retry-delay

Figure 27.3 *Sample Output: show policy*

```
bstnA# show policy

Namespace:             medco
Volume:                /vol

Rule                      Type             Status
------------------------  ---------------  ------------------------------------------

Namespace:             wwmed
Volume:                /acct

Rule                      Type             Status
------------------------  ---------------  ------------------------------------------
docs2das8                 Place            Vol. Scan: Complete      Migration: Complete
fm1                       Share Farm       Vol. Scan: Complete      Migration: Complete

Namespace:             medarcv
Volume:                /rcrds

Rule                      Type             Status
------------------------  ---------------  ------------------------------------------
rcrdsArchive              Snapshot         Enabled
dailyArchive              Place            Vol. Scan: Complete      Migration: Complete
masterDirs2Rx             Place            Vol. Scan: Complete      Migration: Complete
medFm                     Share Farm       Vol. Scan: Complete      Migration: Complete
```

```
Namespace:             medarcv
Volume:                /lab_equipment

Rule                      Type              Status
------------------------  ----------------  ------------------------------------------
hourlySnap                Snapshot          Enabled
dailySnap                 Snapshot          Enabled
mirrorSnap                Replica Snapshot  Enabled
labArchive                Snapshot          Enabled
busy2tier1                Place             Vol. Scan: Paused       Migration: Paused
nonbusy2tier2             Place             Vol. Scan: Paused       Migration: Paused
masterDirs2Tier1          Place             Vol. Scan: Complete     Migration: Complete
tier1                     Share Farm        Vol. Scan: Complete     Migration: Complete
tier2                     Share Farm        Vol. Scan: Complete     Migration: Complete
```

Figure 27.4 *Sample Output: show policy namespace*

```
prtlndA# show policy namespace nemed

Namespace:             nemed

  Namespace Migration Configuration:
    Migrate-Attempts:                          32
    Migrate-Delay:                             30 Seconds
    Migrate-Retry-Delay:                       900 Seconds

Volume:                /acctShdw

  Rule
Priority   Rule                      Type              Status
---------  ------------------------  ----------------  --------------------------------------
    1      farm1                     Share Farm        Vol. Scan: Complete   Migration: Complete
```

Figure 27.5 *Sample Output: show policy namespace ... details*

```
prtlndA# show policy namespace nemed details

Namespace:             nemed

  Namespace Migration Configuration:
    Migrate-Attempts:                          32
    Migrate-Delay:                             30 Seconds
    Migrate-Retry-Delay:                       900 Seconds

Volume:                /acctShdw

  Volume Snapshot Configuration:
    Point-in-Time Consistency:                 Disabled
    Management Command Timeout:                80 seconds
    NFS Directory Name:                        .snapshot
    Directory Display:                         None

  Share Farm:       farm1

                  Placement Frequency   Freespace Status
     Share Name   Count/Total           Free   Size   Pct Free
     -----------  --------------------  ----------------------
     back1        20  / 30       66 %   18 GB   19 GB   92 %
     back2        10  / 30       33 %   111 GB  113 GB  98 %


  New File Placement Rule:  farm1
```

```
    Configuration:
      Constrain Files:                          No
      Constrain Directories:                    No
      Balance Mode:                             Round-Robin
      Auto Migrate:                             Disabled

      State:                                    Enabled

    Status:
      New File Placement Status:                Enabled
```

Figure 27.6 *Sample Output: show policy namespace ... volume*

```
prtlndA# show policy namespace nemed volume /acctShdw

Namespace:             nemed

  Namespace Migration Configuration:
    Migrate-Attempts:                           32
    Migrate-Delay:                              30 Seconds
    Migrate-Retry-Delay:                        900 Seconds

Volume:                /acctShdw

  Volume Snapshot Configuration:
    Point-in-Time Consistency:                  Disabled
    Management Command Timeout:                 80 seconds
    NFS Directory Name:                         .snapshot
    Directory Display:                          None

  Share Freespace:

                 Free Space Thresholds   Freespace Status
    Share Name   Maintain    Resume      Free   Size    Pct Free
    -----------  ----------------------  ----------------------
    back1             2 %         4 %      18 GB   19 GB   92%
    back2             2 %         4 %     111 GB  113 GB   98%

  Rule
Priority   Rule                      Type              Status
---------  ------------------------  ----------------  ---------------------------------------
    1      farm1                     Share Farm        Vol. Scan: Complete   Migration: Complete
```

Figure 27.7 *Sample Output: show policy namespace ... volume ... rule*

```
bstnA# show policy namespace wwmed volume /acct rule docs2das8

Namespace:             wwmed

  Namespace Migration Configuration:
    Migrate-Attempts:                           32
    Migrate-Delay:                              30 Seconds
    Migrate-Retry-Delay:                        900 Seconds

Volume:                /acct

  Volume Snapshot Configuration:
    Point-in-Time Consistency:                  Disabled
    Management Command Timeout:                 80 seconds
    NFS Directory Name:                         .snapshot
    Directory Display:                          None
```

```
Place Rule:       docs2das8

  Configuration:
    From fileset:                               bulky (files only)
    Target share:                               bills
    Report:                                     docsPlc
    Report Verbose:                             Enabled
    Report Delete Empty:                        Disabled
    Report Errors Only:                         Disabled
    Inline Report:                              docsPlc
    Inline Report Interval:                     hourly
    Inline Report Verbose:                      Enabled
    Inline Report Delete Empty:                 Disabled
    Inline Report Errors Only:                  Disabled
    Migrate limit:                              50G
    Volume Scan:                                Enabled
    Inline Notifications:                       Enabled
    Promote Directories:                        Disabled
    Auto-Close Files:                           Disabled
    Migrate Hard Links:                         Disabled

    Tentative:                                  No
    State:                                      Enabled

  Status:
    Volume Scan Status:                         Complete
    File Migration Status:                      Complete
    New File Placement Status:                  Enabled

  Inline Statistics:
    Inline Interval Start:                      03/07/2012 00:43:21 -0500
    Inline Interval Elapsed:                    00:53:01
    Inline Files Migrated:                      7
    Inline Failed Migrations:                   0
    Inline Size of Failed Migrations:           0 B
    Inline Files Renamed:                       3
    Inline Files Placed:                        33
    Inline Directories Renamed:                 0
    Inline Directories Placed:                  0
    Inline Directories Promoted:                0
    Inline Failed Directory Promotes:           0

  Cumulative Statistics:
    Total Files Migrated From Scan:             68
    Total Files Migrated From Inline Activity:  7
    Total Directories Promoted:                 0
    Total Failed Migrations:                    0
    Total Size of Failed Migrations:            0 B
    Total Failed Directory Promotes:            0
    Total Files Forced Closed:                  0
    Total Retried Migrations:                   0
    Total Canceled Migrations:                  0
    Total Hard Links Skipped:                   0
    Total Files Placed Inline:                  33
    Total File Renames Processed Inline:        3
    Total Directories Placed Inline:            0
    Total Directory Renames Processed Inline:   0
    Number of Inline Overflow Errors:           0
    Number of Scans Performed:                  1

  Queue Statistics:
    First-time Migrates:                        0
```

```
    Requeued Migrates:                          0
    Queued Directory Promotes:                  0

  Last Scan Statistics:
    Scan Started:                               03/07/2012 00:43:31 -0500
    Scan Completed:                             03/07/2012 00:43:48 -0500
    Elapsed Time:                               00:00:17
    Scan Report:                                docsPlc_20120307004321.rpt
    Number of Times Paused:                     0
    Total Time Paused:                          00:00:00
    Number of Times Stopped by Low Space:       0
    Time Waiting for Free Space:                00:00:00
    Files Scanned:                              1211
    Directories Scanned:                        165
    Directories Rescanned After Failure:        0
    Directories That Could Not Be Scanned:      0
    Files in Fileset:                           73
    Files Migrated:                             68
    Size of Files Migrated:                     89 MB
    Directories Promoted:                       0
    Failed Migrations:                          0
    Size of Failed Migrations:                  0 B
    Failed Directory Promotes:                  0
    Files Forced Closed:                        0
```

show policy files-closed

Purpose A file-placement rule can automatically close an open file and hold it closed until the rule finishes migrating the file. (This only applies to files opened by CIFS clients.) Use the show policy files-closed command to view all files that have been auto-closed by a particular volume.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show policy files-closed namespace`** ***`namespace`*** **`volume`** ***`vol-path`*** **`[rule`** ***`rule-name`*** **`]`**

namespace (1-30 characters) is the CIFS-supporting namespace,

vol-path (1-1024 characters) identifies a managed volume by its path name, and

rule-name (optional, 1-1024 characters) is a particular file-placement rule. You can use show policy *namespace vol-path* for a list of all rules in a volume. If you omit this, the output shows the files closed by all file-placement rules in the volume.

Guidelines The migrate close-file command makes a file-placement rule close open files automatically. This command shows files that have been automatically closed by a rule, not files that are closed by CIFS clients.

The output shows the Namespace and Volume in its top two fields. The Rule field shows the name of the rule. If the rule is a file-placement rule, a table appears below it with one row for every auto-closed file. Each row shows the exact time the file was closed (in UTC, not local time) and the virtual path to the file. The virtual path starts at the root of the managed volume.

Sample `bstnA# show policy files-closed namespace medarcv volume /rcrds`
shows all files that have been automatically closed by the "medarcv~/rcrds" rule. See Figure 27.8 for sample output.

Related Commands migrate close-file
show policy

Figure 27.8 *Sample Output: show policy files-closed*

```
bstnA# show policy files-closed namespace medarcv volume /rcrds

Namespace:        medarcv
Volume:           /rcrds
Rule
---------------------------------------
dailyArchive

Time                 Filename
-------------------  ------------------------------------------------------------
2009-06-25T05:43:26  /recoveryStats/stdProcedure.doc
2009-06-25T05:43:26  /recoveryStats/memo.doc
2009-06-25T05:43:26  /recoveryStats/recoveryStats.xls
2009-06-25T05:43:26  /2005/recoveryStats/recoveryStats.xls
2009-06-25T05:43:26  /recoveryStats/labBudget.xls
```

```
2009-06-25T05:43:26   /recoveryStats/paid_holidays.doc
2009-06-25T05:43:26   /VIP_wing/recoveryStats/examSchedule.doc
Rule
----------------------------------------
masterDirs2Rx

Time                  Filename
--------------------  ------------------------------------------------------------
Rule
----------------------------------------
rcrdsArchive

Time                  Filename
--------------------  ------------------------------------------------------------
```

source

Purpose A file-placement rule moves files onto chosen storage. You can also use a file-placement rule to drain all files off of a share or shares. Use the source command to select a source share or a source share farm for a file-placement rule.

Mode gbl-ns-vol-plc

Security Role(s) storage-engineer or crypto-officer

Syntax

```
source share share-name
source share-farm share-farm-name
no source
```

share-name (1-64 characters) is the name of a source share in the current volume. This causes the placement rule to select its files from a single share.

share-farm-name (1-64 characters) is a share farm in the current volume. This causes the placement rule to select its files from the chosen share farm.

Default(s) no source

Guidelines This command restricts a file-placement rule to a particular source share or share farm. It is designed for draining all files from a share or share farm, usually to prepare for removing the share(s) from the volume. You can drain the shares, wait for all of the volume's snapshot rules to age out, and then use remove-share migrate to remove them.

This is *not* recommended for a volume with Tiered storage, where file-placement rules migrate new files to Tier-1 shares and older files to Tier-2 shares. Those file-placement rules should select their files from every share in the managed volume.

Samples The following example selects the "archives~/etc~rh1" as a source share:

```
bstnA(gbl-ns-vol[archives~/etc])# place-rule empty
bstnA(gbl-ns-vol-plc[archives~/etc~empty])# source share rh1
bstnA(gbl-ns-vol-plc[archives~/etc~empty])# ...
```

This example selects an entire share farm:

```
bstnA(gbl-ns-vol[insur~/claims])# place-rule drainFm
bstnA(gbl-ns-vol-plc[insur~/claims~drainFm])# source share-farm archival
bstnA(gbl-ns-vol-plc[insur~/claims~drainFm])# ...
```

Related Commands place-rule
snapshot rule
remove-share migrate

target

Purpose Use the target command to choose a storage target for the current placement rule. A placement rule puts chosen files onto selected storage. From gbl-ns-vol-plc mode, specify a share or a share farm as the storage target.

Mode gbl-ns-vol-plc

Security Role(s) storage-engineer or crypto-officer

Syntax

```
target share share-name
target share-farm share-farm-name
```

share-name (1-64 characters) is a share from the current volume. Use the show namespace command to see the shares in each volume.

share-farm-name (1-64 characters) is a share farm within the current volume. Use the show namespace command to see the share farms in the namespace.

Default(s) None

Guidelines This command sets a storage target for the current placement rule. You can set one storage target.

Samples

```
bstnA(gbl-ns-vol-plc[wwmed~/acct~toNas23])# target share nas23
```

selects a share, “nas23,” as the home for all files in the rule’s source filesets.

```
bstnA(gbl-ns-vol-plc[ns3~/logs~distFiles])# target share-farm fm3
```

selects a share farm, “fm3,” as the target for a different placement rule:

Related Commands place-rule

tentative

Purpose

You can set a file-placement rule to run simulations instead of actually migrating files. This can be useful for projecting the effects of a file-placement rule on your back-end filers. Use the tentative command to disable actual migrations in the current rule.

Use no tentative to activate migrations in the current file-placement rule.

Mode

gbl-ns-vol-plc

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
tentative
no tentative
```

Default(s)

disabled

Guidelines

A tentative rule logs its proposed migrations in the syslog without actually migrating any files. This allows you to see the results of the rule without actually committing to it. The log component, POLICY_ACTION, creates the syslog messages. Use the show logs syslog (or grep logs pattern syslog) command to see the results that would have occurred had the rule been fully enabled.

Samples

```
bstnA(gbl-ns-vol-plc[archives~/docs~2bkup])# tentative
```

makes the current file-placement rule, “2bkup,” a tentative rule. The next time it runs, it does not perform any migrations.

```
bstnA(gbl-ns-vol-plc[archives~/docs~2bkup])# no tentative
```

takes the “2bkup” rule out of the tentative state. The next time the rule runs, it performs actual file migrations.

Related Commands

namespace -> volume -> place-rule
show logs syslog
grep syslog

volume-scan

Purpose

By default, a file-placement rule finds its source files by monitoring all new files as they are created, monitoring client changes as they happen (inline), *and* scanning the volume for existing files. Use the no volume-scan command to disable volume scans and work with new and changed files only.

Use the volume-scan command to re-enable volume scans for the current rule.

Mode

gbl-ns-vol-plc

Security Level

storage-engineer or crypto-officer

Syntax

```
no volume-scan
volume-scan
```

Default(s)

```
volume-scan
```

Guidelines

This can be useful for situations where shares already have current files distributed correctly. Note that these situations are somewhat rare. With volume scans disabled, all new and edited files are placed appropriately and no time is wasted examining back-end shares.

You can also stop the rule from receiving client-change notifications; use no inline notify command.

Sample

```
bstnA(gbl-ns-vol-plc[wwmed~/acct~toNas23])# no volume-scan
```

disables volume scans for the current rule on the '/acct' volume.

Related Commands

place-rule
inline notify

wait-for migration

Purpose Use the wait-for migration command to wait for a file-placement rule to finish migrating files from a source to a target.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax **`wait-for migration namespace`** ***`name`*** **`volume`** ***`vol-path`*** **`rule`** ***`rule`*** **`[timeout`** ***`timeout`*****`]`**

name (1-30 characters) is the name of the namespace.

vol-path (1-1024 characters) identifies the volume.

rule (1-2096 characters) is the name of the rule that is migrating files (such as a place rule).

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - none, wait indefinitely

Guidelines A file-placement rule (created with the place-rule command) may take a long time to migrate a very-large fileset. You can use the wait-for migration command to wait for the operation to finish.

If you set a ***timeout*** and it expires before all files have finished migrating, the command exits with a warning. To interrupt the wait-for migration command, press **<Ctrl-C>**.

This can be useful for CLI scripts, which you can copy onto the switch (with copy ftp, copy scp, copy {nfs|cifs}, or copy tftp), and run.

Sample

```
bstnA# wait-for migration namespace medarcv volume /rcrds rule
dailyArchive timeout 30
```

waits up to 30 seconds for the “dailyArchive” rule to finish a file migration.

Related Commands place-rule

28

Share Farms

auto-migrate

Purpose Use the auto-migrate command to allow automatic migrations of files off of any share where the free space is too low.

Modes gbl-ns-vol-sfarm

Security Role(s) storage-engineer or crypto-officer

Syntax
```
auto-migrate
no auto-migrate
```

Default(s)
```
no auto-migrate
```

Guidelines The auto-migrate command allows the share farm to migrate files away from any of its shares who get too low on free space. The policy freespace command sets the minimum free space to maintain at a given share. The policy engine migrates files off of the share to other shares in the same share farm, until the share's free space reaches or exceeds the *resume-migrate* threshold you set with the same policy freespace command.

This command migrates files that already exist in the share farm. To balance the distribution of newly-created files, use the balance command. This command only functions with balance capacity, which takes a share's free space into consideration when it distributes new files. This command cannot function in a share farm where constrain-files or constrain-directories are active; the constraints make auto migration impossible.

The share farm cannot migrate a file that is open through CIFS and holding an exclusive-write lock. The show cifs-service open-files command shows all such files, and close cifs file closes one from the CLI.

Use the show statistics migration command to monitor the file-migration activity in this share farm.

If an automatic migration fails, the policy engine waits 300 seconds (five minutes) and retries the migration.

Sample
```
bstnA(gbl-ns-vol-sfarm[wwmed~/acct~fm1])# auto-migrate
```
starts auto migration away from any share in the "fm1" farm whose free space drops below the policy freespace setting.

Related Commands
namespace -> volume -> share-farm
policy freespace
balance
show cifs-service open-files
close cifs file
show statistics migration

balance

Purpose Use the balance command to balance all newly-created files amongst the shares in a share farm.

Modes gbl-ns-vol-sfarm

Security Role(s) storage-engineer or crypto-officer

Syntax **`balance {capacity | latency | round-robin}`**

capacity distributes files to each share based on the share's current free space (capacity). If share A has twice as much free space as share B, it gets twice as many new files.

latency places more files in shares with low latency than shares with higher latency. Share *latency* is the average round-trip time from the ARX to the share; low latency is a sign of more available bandwidth between the switch and the share.

round-robin uses manually-configured weights to guide the distribution. If share A has a weight that is twice as high as share B's weight, share A gets twice as many new files. You assign weight to a share when you add it to the share farm with share (gbl-ns-vol-sfarm).

Default(s) **`balance round-robin`**

Guidelines Any unavailable shares are removed from the new-file balancing rotation automatically. A share also drops out of the rotation if its free space drops too far; use the policy freespace command to set the minimum free space to maintain on a share.

To stop the distribution of new files, you can use constrain-files to constrain any new file to the same share as its parent directory. Additionally, you can use constrain-directories to keep directories with their parents. These constrain commands disable the effects of any balance command.

To activate balancing for the current share farm, use the enable (gbl-ns-vol-sfarm) command at the end of the command sequence. To stop balancing for the current share farm, use the no enable (gbl-ns-vol-sfarm) command, or invoke one of the constrain commands discussed above.

The auto-migrate command moves existing files off of a share if its free space drops too far. The auto-migrate command only functions if balance capacity is set, and you cannot set any other balance option if auto-migrate is enabled in this share farm.

Samples

```
bstnA(gbl-ns-vol-sfarm[wwmed~/acct~fm1])# balance capacity
```

balances new files for the '/acct~fml' share farm based on share capacity.

```
bstnA(gbl-ns-vol-sfarm[ns12~/usr~fm5])# share sh15 weight 40
bstnA(gbl-ns-vol-sfarm[ns12~/usr~fm5])# share sh16 weight 40
bstnA(gbl-ns-vol-sfarm[ns12~/usr~fm5])# balance round-robin
```

configures round-robin distribution in the "/usr~fm5" share farm. Two shares in the farm have equal weights, so the policy engine assigns them the same number of new files.

Related Commands

namespace -> volume -> share-farm
auto-migrate
policy freespace
constrain-files
constrain-directories
enable (gbl-ns-vol-sfarm)

constrain-directories

Purpose

Use the constrain-directories command to keep any newly-created directory on the same share as its parent. That is, the share that holds the parent's master directory, for situations where the parent is striped on multiple shares in the volume.

Modes

gbl-ns-vol-sfarm

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
constrain-directories [below depth]
no constrain-directories
```

depth (optional, integer 0-100) is the maximum depth of new directories that the share farm is allowed to balance. Below this directory depth, any new directory stays on the same share as its parent. This tends to preserve cohesive directory trees on back-end filers. For example, consider a volume that contains three directory trees, /a, /b, and /c: a share-farm with a depth of 1 could distribute the three directories to three separate shares, then keep the lower-level directories with their parents.

Default(s)

```
no constrain-directories
```

Guidelines

This limits or entirely disables the balance rule, which balances the placement of new files.

Use the constrain-files command to keep new files with their parent directories, too. If new files are not constrained and a client creates multiple files in the same directory, the balance rule is likely to distributes the files to multiple shares; to hold those files, it must replicate the directory in each share. This is called directory *striping*. In this situation, the constrain-directories command is choosing the share with the first (master) copy of the directory, but more copies are likely to exist.

Guidelines: Share-Farm Mirroring in a Tiered Volume

Many installations use place-rules and age-filesets to implement tiered storage in a managed volume. For example, files modified in the past quarter are placed onto tier 1, and files that have gone unmodified for more than a quarter migrate down to tier 2. The tiers can be individual shares or they can be share farms. If they are share farms where constrain-files or constrain-directories is active, the volume uses an algorithm called *share-farm mirroring* to determine which share in the destination share farm gets the migrated subdirectory.

The share-farm mirroring technique finds a share for a directory when its parent's master directory is on another share farm. For example, suppose the master directory for "/home/jsmith" is on share farm A, the tier 1 share farm, but one of its subdirectories should migrate to share farm B. The mirroring technique finds the right share to receive the constrained subdirectory. The volume mirrors the share farms by mapping the first shares in each farm to each other, the second shares in each farm to each other, and so forth. If the master resides on share 2 in share farm A, all other share farms constrain their stripes to share 2. If "/home/jsmith" is on share 3 in the tier 1 share farm, the stripes for "/home/jsmith" reside on share 3 of the tier 2 farm and share 3 of the tier 3 farm. The constraints keep the directory on the same share in each farm, together with its constrained subdirectories.

Share farm mirroring requires all tiers to be share farms; do not mix a tier 1 share with a tier 2 share farm. Also, the mirroring algorithm works best if all share farms have the same number of shares.

Samples

```
bstnA(gbl-ns-vol-sfarm[ns2~/usr~fm4])# constrain-directories
```

forces all new directories in the ns2~/usr~fm4 share farm to remain in their parent directories in the share farm.

```
bstnA(gbl-ns-vol-sfarm[ns2~/usr~fm4])# constrain-directories below 1
```

forces all new directories in the ns2~/usr~fm4 share farm to be constrained to their parent directories when the new directories are below the first subdirectory level.

Related Commands

namespace -> volume -> share-farm
balance
constrain-files

constrain-files

Purpose

Use the constrain files command to specify that new files be placed in the same shares as their parent directories. This reduces directory striping in the share farm; the share that has the master directory is typically the only share with any instance of the directory at all.

Use the no form of the command to remove the restriction and allow the balance rule to distribute files evenly throughout the current share farm.

Modes

gbl-ns-vol-sfarm

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
constrain-files
no constrain-files
```

Default(s)

```
no constrain-files
```

Guidelines

This command enables you to restrict new-file placement to the same share as the parent directory. It disables the balance command with respect to files. This has no effect on the auto-migration of existing files, which is done to relieve a share that is running low on free space (see auto-migrate). Note that an auto migration has the side effect of creating stripe directories on other shares, but only when free space is running low.

You can also use the constrain-directories command to keep directories on the same share as their parents.

Guidelines: Share-Farm Mirroring in a Tiered Volume

Many installations use place-rules and age-filesets to implement tiered storage in a managed volume. For example, files modified in the past week are placed onto tier 1, files modified last week are on tier 2, and files that have not been modified any time in the past 2 weeks are placed on tier 3. The tiers can be individual shares or they can be share farms. If they are share farms where constrain-files or constrain-directories is active, the volume uses an algorithm called *share-farm mirroring* to determine which share gets a file.

The share-farm mirroring technique finds a share for a file when its parent's master directory is on another share farm. For example, suppose the master directory for "/home/jdoe" is on share farm A, the tier 1 share farm, but one of its files should migrate to share farm B. The mirroring technique finds the right share to receive the constrained file. The volume mirrors the share farms by mapping the first shares in each farm to each other, the second shares in each farm to each other, and so forth. If the master resides on share 4 in share farm A, all other share farms constrain their stripes to share 4. If "/home/jdoe" is on share 2 in the tier 1 share farm, the stripes for "/home/jdoe" reside on share 2 of the tier 2 farm and share 2 of the tier 3 farm. The constraints keep the directory on the same share in each farm, together with its constrained files.

Share farm mirroring requires all tiers to be share farms; do not make a stand-alone share into tier 1 and a share farm into tier 2. Also, we recommend that all share farms have the same number of shares.

Sample

```
bstnA(gbl-ns-vol-sfarm[ns~/vol~fm])# constrain-files
```

causes all new files in the "ns~/vol~fm" share farm to remain in the same shares as their parent directories.

Related Commands namespace -> volume -> share-farm
balance
constrain-directories

enable (gbl-ns-vol-sfarm)

Purpose

Use the enable command to enable the current rule. A new rule is disabled by default, and the policy engine ignores disabled rules.

Use no enable to disable the current rule.

Modes

gbl-ns-vol-sfarm

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
enable
no enable
```

Default(s)

```
no enable
```

Guidelines

You must enable a share farm for the policy engine to enforce any of its policies (such as those set with balance or auto-migrate).

Sample

```
bstnA(gbl-ns-vol-sfarm[archives~/multimedia~minSpace])# no enable
```

disables all rules in the "minSpace" share farm.

Related Commands

namespace -> volume -> share-farm
auto-migrate
balance

policy freespace (gbl-ns-vol-sfarm)

Purpose

Whenever any rule migrates a file to a share, it checks the share's free space against two thresholds. Each share has a minimum free space to *maintain*, as well as another (typically larger) level of free space called a *resume* threshold. If the file would reduce the share's free space below the *maintain* level, the rule waits until the free space rises to the *resume* level before migrating any more files to the share. The policy freespace command sets these free-space thresholds for an ARX share. Use this policy freespace command, from gbl-ns-vol-sfarm mode, to set the same thresholds for every share in the current share farm.

Use no policy freespace to return to the defaults for every share in the current share farm.

Mode

gbl-ns-vol-sfarm

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
policy freespace maintain{k|M|G|T} resume-migrate resume{k|M|G|T}

policy freespace percent maintain-pct resume-migrate resume-pct

no policy freespace
```

maintain is the amount of free space to maintain on these shares.

> **k|M|G|T** chooses the unit of measure: **k**ilobytes, **M**egabytes, **G**igabytes, or **T**erabytes. A kilobyte is 1,024 bytes, a megabyte is 1,024 kilobytes (1,048,576 bytes), and so on..

resume is a free-space level for resuming migrations to any of these shares. The policy engine uses this value if one of the shares has reached it's *maintain* value; no rule can migrate to such a share until its free space rises back to the *resume* level.

percent (optional) indicates that you are using disk-space percentages for these values instead of specific size measures.

maintain-pct (1-100) expresses the *maintain* value as a percentage of the overall share size.

resume-pct (1-100) is the *resume* threshold, expressed as a percentage of the share's total space.

Default(s)

maintain - 1G

resume - 2G

Guidelines

This is a macro command for the policy freespace command in gbl-ns-vol-shr mode. This command invokes the individual policy freespace command for every share in the current share farm. The output from show global-config shows the individual share-level policy freespace commands, not this macro command.

This is similar to the policy freespace (gbl-ns, gbl-ns-vol) command, a macro command for namespaces and volumes.

Samples

```
bstnA(gbl-ns-vol-sfarm[wwmed~/acct~fm1])# policy freespace 3G
resume-migrate 4G
```

makes the policy engine maintain at least 3 Gigabytes of free space on every share in the “fm1” share farm. If that limit is reached by clients or approached by a migration rule on any of these shares, the share is ineligible for any more migrations until its free space rises back up to 4 Gigabytes.

```
bstnA(gbl-ns-vol-sfarm[wwmed~/acct~fm1])# no policy freespace
```

reverts all shares in the above share farm to the default free-space thresholds.

Related Commands

namespace -> volume -> share-farm
policy freespace
policy freespace (gbl-ns, gbl-ns-vol)
place-rule
show policy

share (gbl-ns-vol-sfarm)

Purpose Use the share command from gbl-ns-vol-sfarm mode to add a managed volume share to the current share farm.

Use the no form of the command to delete a share from the share farm.

Mode gbl-ns-vol-sfarm

Security Role(s) storage-engineer or crypto-officer

Syntax

```
share name [weight weight]
no share name
```

name (1-64 characters) is the name of a share in the current managed volume. Use the show global-config namespace command for a list of configured shares.

weight (0 - 100) is the new-file-placement weight for the share. This number is proportional to other shares in the share farm; if one share has a 60 weight and another share has 30, the first share gets twice as many new files as the second. This weight is used with the balance round-robin rule; it is ignored for any other balance setting.

Default(s) None.

Guidelines You can add as many shares as desired. A share can only belong to one share farm.

The no share command from gbl-ns-vol-sfarm mode does not delete the share's configuration; it only deletes the share from the share farm.

The order in which you invoke this command determines the ID of the share within the share farm. The first-added share has an ID of 1, the second has an ID of 2, and so forth. The share IDs are used for *share-farm mirroring* in a tiered volume; if a file migrates to a tier 3 share farm, its parent's master directory is in the tier 1 share farm, and it is constrained to the same share as its parent's master directory, the file goes to the share with the same ID as the parent's master directory. The documentation for constrain-files and constrain-directories provides further details on share-farm mirroring.

If you remove the last share from the share farm with no share, other rules (such as a place-rule) can no longer use the share farm as a target for file migrations. The CLI prompts for confirmation if you attempt to remove the last share from the farm; enter **yes** to proceed.

Samples

```
bstnA(gbl-ns-vol-sfarm[archives~/multimedia~gold])# share lun77
```

adds the lun77 share to the gold share farm.

```
bstnA(gbl-ns-vol-sfarm[archives~/multimedia~gold])# no share lun70
```

deletes the lun70 share from the "gold" share farm.

Related Commands

namespace -> volume -> share-farm
show global-config namespace
constrain-files
constrain-directories

share-farm

Purpose

A share farm is a group of shares in a managed volume. You can balance the new files that are created in a share farm, and you can auto-migrate files off of a share that is low on free space. Use the share-farm command to create a new share farm, or to edit an existing one.

Use the no form of the command to delete a share farm.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
share-farm name
no share-farm name
```

name (1-1024 characters) is the name you choose for the share farm.

Default(s)

None

Guidelines

When you create a new share farm, the CLI prompts for confirmation. Enter **yes** to create the share farm. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.)

This places you in gbl-ns-vol-sfarm mode, where you must use the share (gbl-ns-vol-sfarm) command to add at least one share to the farm. You can optionally use auto-migrate to migrate files off of any share that falls below a minimum amount of free space. The balance command distributes new files amongst the shares in the farm; you can weight this distribution based on current free space at each share (shares with more space get more new files), current latency at each share (shares with more bandwidth get more new files), or according to the weights you set with the share (gbl-ns-vol-sfarm) command. The latter form of the balance rule is active by default. To stop balancing new files, use the constrain-files and/or constrain-directories commands to keep any new file and/or directory on the same share as its parent directory.

Use show policy for details about a share farm and its policies. Use the show statistics migration command to monitor all migrations in the share farm.

Deleting a share farm has no effect on the farm's shares, other than to invalidate any of the share farm's policies.

Samples

```
bstnA(gbl-ns-vol[archives~/multimedia])# share-farm gold

This will create a new share farm.

Create share farm 'gold'? [yes/no] yes
bstnA(gbl-ns-vol-sfarm[archives~/multimedia~gold])#
```

creates a share farm named "gold" in the current volume.

```
bstnA(gbl-ns-vol[archives~/multimedia])# no share-farm platinum
```

deletes the "platinum" farm.

Related Commands

namespace -> volume
share (gbl-ns-vol-sfarm)
auto-migrate
balance
constrain-files
constrain-directories
show policy
show statistics migration

29

Shadow Volume

bandwidth-limit

Purpose

You can use the bandwidth-limit command to set a bandwidth ceiling for the current shadow-copy rule. Each run of the rule abides by this limit.

Use no bandwidth-limit to remove any limit on bandwidth.

Modes

gbl-ns-vol-shdwcp

Security Level

storage-engineer or crypto-officer

Syntax

`bandwidth-limit` *`rate`*`[K|M|G|T]`
`no bandwidth-limit`

rate (100,000-4,000,000,000,000, or 1M-4T) is maximum bandwidth that you choose for the current shadow-copy rule.

K|M|G|T (optional) chooses the units. These are Kbps, Mbps, Gbps, or Tbps. They are 10-based: 1 Kbps is 1,000 bps, 1 Mbps is 1,000,000 bps, and so forth. Do not include any space between the ***rate*** and this unit: **100G** is correct, but **100 G** is not. If you omit this, the units are bits-per-second (bps).

Default(s)

`no bandwidth-limit`

Guidelines

Each run of the current shadow-copy rule abides by this bandwidth limit.

Sample

```
bstnA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])# bandwidth-limit 750000
```

allows the 'SVrule' rule to use up to 750,000 BPS for its file transfers.

Related Commands

shadow-copy-rule

cifs-8dot3-resolution

Purpose File systems that support CIFS also support an *alternate name* for its files and directories. A file or directory whose name does not fit the old-style "8.3" pattern gets an alternate name from its back-end filer. If a file on the source volume has a primary name that matches another file's alternate name, the shadow-copy rule cannot copy the file to the target volume. To enable the rule to safely perform this copy operation, use the cifs-8dot3-resolution command.

Use no cifs-8dot3-resolution to return to the default; this is not recommended.

Modes gbl-ns-vol-shdwcp

Security Level storage-engineer or crypto-officer

Syntax

```
cifs-8dot3-resolution
no cifs-8dot3-resolution
```

Default(s)

```
no cifs-8dot3-resolution
```

Guidelines The no form of this command implies that it is impossible for a client to create an 8.3 name that matches the alternate name of a different file. While this event is rare, it is possible at any site that supports CIFS. We recommend enabling this feature for all shadow-copy rules in all CIFS-supporting volumes.

Sample

```
bstnA(gbl-ns-vol-shdwcp[insur~/claims~insurSV])# cifs-8dot3-resolution
```

sets the 'insurSV' rule to support copies of 8.3 names.

Related Commands shadow-copy-rule

database-location

Purpose

The shadow-copy database is critical for the maintenance of the shadow volume, so the shadow-copy rule should never run out of space for it. You can place this database on the shadow volume's metadata share or on its storage shares. The metadata share is the recommended location for this critical database. For rare situations when a shadow volume's storage shares have more space than its metadata share, you can use the database-location volume command to place the shadow volume's database on its storage shares.

Use database-location metadata-share to place the shadow-copy database onto the metadata share, as generally recommended.

Modes

gbl-ns-vol-shdwcp

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
database-location volume
database-location metadata-share
```

Default(s)

```
database-location metadata-share
```

Guidelines

You must rebuild the shadow volume (with nsck ... rebuild) before you can use the database-location volume command. This is not required for moving the database to the recommended metadata share.

If the shadow-copy rule runs out of disk space for its database, the database may be corrupted or the rule could hang indefinitely. A corrupted database causes a full scan of the source volume. The re-scan occurs on the next run of the rule, and it is time-consuming if the source volume is large. We recommend storing this on the metadata share for the volume. The volume's storage shares can expand unpredictably, according to the changing habits of the source volume's clients, and may unexpectedly run low on space. Volume metadata is generally smaller and grows at a regular pace. Furthermore, the ARX carefully monitors the disk space on a metadata share, issuing SNMP traps whenever space runs low.

Sample

```
bstnA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])# database-location
metadata-share
```

places the shadow-copy database for "SVrule" onto its target volume's metadata share.

Related Commands

shadow-copy-rule

delta-threshold

Purpose

When a source file changes, a shadow-copy rule transfers only the file's changed portion (delta) to the shadow volume. Files that are considered "too small" are transferred in their entirety. Use delta-threshold to set the minimum file size to trigger a delta transfer, so that a file is always copied whole if it is smaller than the threshold.

Use no delta-threshold to revert to the default delta threshold.

Modes

gbl-ns-vol-shdwcp

Security Role(s)

storage-engineer or crypto-officer

Syntax

`delta-threshold` ***`file-size`*****`[k|M|G|T]`**
`no delta-threshold`

file-size (1-18,446,744,073,709,551,615) is the minimum size of a file that is eligible for delta transfers.

k|M|G|T (optional) specifies units; **k** for KiloBytes (1024 bytes), **M** for MegaBytes (1024*1024 bytes), **G** for GigaBytes (1024*1024*1024), and **T** for TeraBytes (1024*1024*1024*1024). Do not put any space between the number and this unit (for example, **100k** is valid but **100 k** is not). The default is bytes.

Default(s)

100M

Guidelines

A delta-transfer uses lower bandwidth than a full transfer, but requires more computation time on both the source volume and the shadow volume. This threshold saves the computation time for small files, which would not require a great deal of bandwidth if the rule always transfers the whole file. This is intended for networks with higher bandwidth, or for installations where the source and shadow volumes are on the same switch.

Samples

```
bstnA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])# delta-transfer 1G
```

allows delta transfers only for files larger than 1 GigaByte.

```
bstnA(gbl-ns-vol-shdwcp[archives~/etc~shdw])# no delta-transfer
```

returns to the default threshold.

Related Commands

shadow-copy-rule

enable (gbl-ns-vol-shdwcp)

Purpose Use the enable command to enable the current shadow-copy rule.

Use no enable to disable the rule.

Modes gbl-ns-vol-shdwcp

Security Role(s) storage-engineer or crypto-officer

Syntax

```
enable
no enable
```

Default(s)

```
no enable
```

Guidelines You must enable the shadow-copy rule for the policy engine to use it. Once enabled, the shadow-copy rule follows its assigned schedule and copies its fileset from the source volume to the shadow volume.

Samples

```
bstnA(gbl-ns-vol-shdwcp[archives~/etc~shdw])# enable
```

enables the “shdw” rule.

```
prtlndA(gbl-ns-vol-shdwcp[nemed~/acct~SV2])# no enable
```

disables the “SV2” shadow-copy rule.

Related Commands shadow-copy-rule

from fileset (gbl-ns-vol-shdwcp)

Purpose Use the from fileset command to select the source fileset for the current shadow-copy rule. The shadow-copy rule periodically copies the fileset from the source volume to the shadow volume.

Mode gbl-ns-vol-shdwcp

Security Role(s) storage-engineer or crypto-officer

Syntax `from fileset` ***`fileset-name`*** `[match {files | all}]`

fileset-name (1-1024 characters) identifies the fileset.

match {files | all} (optional) determines whether the fileset matches files only or files and directories.

Default(s) **match** - **files** only (and the directories that contain the matching files)

Guidelines You can use a single fileset as your source. All files in this fileset will be copied to the shadow volume, which you configure with the target (gbl-ns-vol-shdwcp) command.

If a single fileset is too restrictive, use the union-fileset command to join multiple filesets into one.

The policy-cifs-attributes-fileset is not compatible with shadow-copy rules. This command returns an error if you use a policy-cifs-attributes-fileset or a container fileset that includes one. (Container filesets are intersection-filesets and union-filesets).

Sample The following example selects the "worthSaving" source fileset for the shadow-copy rule, "SVrule:"

```
prtlndA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])# from fileset worthSaving
```

Related Commands shadow-copy-rule
target (gbl-ns-vol-shdwcp)

inline-notify

Purpose A shadow-copy rule keeps track of all source-volume directories where at least one change occurred. Each time the shadow-copy rule runs, it only examines and copies from the set of changed directories. Use no inline-notify to stop tracking source-volume changes, thus forcing a full tree-walk of the source volume for every shadow-copy run.

Use the inline-notify command track all source-volume changes.

Modes gbl-ns-vol-shdwcp

Security Role(s) storage-engineer or crypto-officer

Syntax

```
no inline-notify
inline-notify
```

Default(s) `inline-notify`

Guidelines The default improves performance by keeping track of changed directories since the previous shadow copy. The policy engine performs a tree-walk (a complete scan) of the directories that changed, not wasting any cycles on directories that remained the same. You should only turn off inline notifications on advice from F5 Support.

Sample

```
bstnA(gbl-ns-vol-shdwcp[ns3~/vol~betaSV])# no inline-notify
```

disables inline notification for “betaSV.”

Related Commands shadow-copy-rule

prune-target

Purpose

By default, the first run of a shadow-copy rule ends by pruning all extraneous files and directories from the shadow volume. These are files and/or directories that were imported into the shadow volume but that have no counterparts in the source volume. The rule prunes these files and directories to ensure that the source and shadow are identical from the start. If the two volumes are already known to be identical, you can use no prune-target to avoid the unnecessary pruning phase.

Use the prune-target command to enable initial target pruning on the shadow (target) volume.

Modes

gbl-ns-vol-shdwcp

Security Level

storage-engineer or crypto-officer

Syntax

```
no prune-target
prune-target
```

Default(s)

```
prune-target
```

Guidelines

The prune-target command allows enabling or disabling of initial pruning on the target volume.

Even if the source and target volumes are identical, millions of files can consume time as the rule compares the two directory trees. The no form of this command is useful when performance is an issue.

This command has no effect on the rule's ability to synchronize deletions. If a client deletes a file or directory on the source volume, the shadow-copy rule always deletes its counterpart on the shadow volume. This command only affects files and directories that are unique to the shadow volume.

Sample

```
bstnA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])# no prune-target
```

disables pruning of shadow volume target volumes for the 'SVrule' rule.

Related Commands

shadow-copy-rule

publish

Purpose

A shadow-copy rule copies each file individually, publishing the file in the client-visible directory tree immediately after a successful transfer to a staging area. Some applications require all files to successfully transfer as a group. Use publish group to publish all of the files as a single group, after all files transfer without errors.

Use publish individual, to return to the default: publish each file after a successful transfer.

Modes

gbl-ns-vol-shdwcp

Security Role(s)

storage-engineer or crypto-officer

Syntax

`publish {group | individual}`

group | **individual** is a required choice:

group publishes all files as a group, after they have all successfully transferred to the shadow volume's staging area.

individual publishes each file individually, as it is transferred.

Default(s)

`publish individual`

Guidelines

The group publishing phase is designed for applications that require all shadow files to be present and consistent, such as a CAD project. If any part of the file-transfer fails, none of the changed files are published.

The publish individual method is intended for shadow volumes where file consistency is less important. For example, an installation with /home directories may prefer to have all the latest files as soon as possible, not allowing one failed transfer to prevent the publication of thousands of other files.

Some CIFS applications block all access to a file when the file is open, making it impossible for the shadow volume to copy the file. To prevent open CIFS files from stopping the entire transfer, use publish individual.

Samples

`bstnA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])# publish group`

reinstates the group-publishing phase for "SVrule."

`prtlndA(gbl-ns-vol-shdwcp[nemed~/acct~SV2])# publish individual`

removes the group-publishing phase; each file is published directly after a successful transfer.

Related Commands

shadow-copy-rule

report (gbl-ns-vol-shdwcp)

Purpose Use this command to enable progress reports for the current shadow-copy rule.

Use no report to prevent progress reports.

Mode gbl-ns-vol-shdwcp

Security Role(s) storage-engineer or crypto-officer

Syntax **`report`** ***`file-prefix`*** **`[verbose] [list-identical] [delete-empty]`**
`no report`

file-prefix (1-1024 characters) sets a prefix for all shadow-copy reports. Each report has a unique name in the following format:

prefixYearMonthDayHourMinute.rpt

For example, home_backup200403031200.rpt could be the name for one report with the "home_backup" prefix.

verbose (optional) enables verbose data in the reports.

list-identical (optional) includes all files that are identical on both the source and shadow volumes. These files are not copied, so they are unlisted by default.

delete-empty (optional) causes the rule to delete any reports that contain no data. Empty reports are created when the source fileset did not change since the previous shadow copy.

Default(s) **`no report`**

Guidelines Every time the shadow-copy rule fires and duplicates the source fileset, the rule can generate a report to show the details of the copy session. The reporting feature is disabled by default; we recommend that you enable reporting for all shadow-copy rules, to keep a log of all shadow-copy sessions.

Use show reports for a list of reports, or show reports file-name to show the contents of one report.

Samples

```
prtlndA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])# report SVetc
```

enables reports for the shadow-copy rule, "SVrule." For a sample report, see Figure 29.1.

```
bstnA(gbl-ns-vol-shdwcp[archives~/etc~shdw])# report etc verbose
list-identical
```

enables verbose reports for the shadow-copy rule, "shdw." The reports also include any files that were found to be identical on both the source and shadow volumes.

```
bstnA(gbl-ns-vol-shdwcp[ns3~/usr~cdn])# no report
```

disables reports for the "cdn" rule.

Related Commands shadow-copy-rule

Figure 29.1 *Sample Report: Shadow-Copy Report*

```
bstnA# show reports SVetc_201002270226.rpt
**** Shadow Copy Report: Started at Sat Feb 27 02:26:41 2010 ****
**** Software Version: 5.02.000.12541 (Feb 23 2010 20:12:44) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

Shadow Copy Status
==================

    Namespace                                 :   wwmed
    Source Volume                             :   /acct/
    Shadow Rule                               :   SVrule
    Fileset                                   :   worthSaving
    Publishing Mode                           :   Individual
    Directory Copy                            :   All
    Shared Access Allowed                     :   Yes
    Inline Notifications                      :   Yes
    Prune Target                              :   Yes
    Bandwidth Limit                           :   5.0 Mb/s (625.0 kB/s)
    Retry Attempts                            :   5
    Retry Delay                               :   300 seconds
    Database Location                         :   metadata-share

    ================================================================================


Shadow Copy Status
==================

    Namespace                                 :   wwmed
    Source Volume                             :   /acct
    Shadow Rule                               :   SVrule
    Fileset                                   :   worthSaving (Files and Directories)
    Delta Threshold                           :   100 MB
    Publishing Mode                           :   Individual
    CIFS SID Translation                      :   Ignore Errors
    Shared Access Allowed                     :   Yes
    Inline Notifications                      :   Yes
    Bandwidth Limit                           :   5.0 Mb/s (625.0 kB/s)
    Retry Attempts                            :   5
    Retry Delay                               :   300 seconds
    Database Location                         :   metadata-share

    ================================================================================
      Processing Started                      :   Feb 27 02:26
      Operating Mode                          :   Full tree walk
      Tree Walk Reason                        :   Initial run

      Target Information
      ------------------
        prtlndA: nemed:/acctShdw/

Shadow Copy File Details
========================

    Filename                                                                    Size
Status
    ------------------------------------------------------------------------------
--------------  ------------------------------
    /layer3.fm.lck
prtlndA: nemed:/acctShdw/: Full update (212 bytes sent)
```

```
    /layer3.fm.lck                                                                    84
New(1)
    /finance_sites.html
prtlndA: nemed:/acctShdw/: Full update (1,428 bytes sent)
    /finance_sites.html                                                            1,299
New(1)
    /layer2.fm.lck
prtlndA: nemed:/acctShdw/: Full update (212 bytes sent)
    /layer2.fm.lck                                                                    84
New(1)
...
Shadow Copy Status
==================

    Namespace                                 :   wwmed
    Source Volume                             :   /acct
    Shadow Rule                               :   SVrule
    Fileset                                   :   worthSaving (Files and Directories)
    Delta Threshold                           :   100 MB
    Publishing Mode                           :   Individual
    CIFS SID Translation                      :   Ignore Errors
    Shared Access Allowed                     :   Yes
    Inline Notifications                      :   Yes
    Bandwidth Limit                           :   5.0 Mb/s (625.0 kB/s)
    Retry Attempts                            :   5
    Retry Delay                               :   300 seconds
    Database Location                         :   metadata-share

    ================================================================================
      Processing Started                      :   Feb 27 02:26
      Processing Completed                    :   Feb 27 02:36
      Elapsed Time                            :   00:10:01
      Operating Mode                          :   Full tree walk
      Tree Walk Reason                        :   Initial run
      Current Phase                           :   Completed

      Target Information
      ------------------
        prtlndA: nemed:/acctShdw/             :   Successful

      Copy Phase Information
      ----------------------
        Phase Started                         :   Feb 27 02:26
        Phase Completed                       :   Feb 27 02:36
        Elapsed Time                          :   00:09:59
        Average Transmission Rate             :   4.7 Mb/s (599.4 kB/s)
        Total Files/Directories               :   4,848
        Files/Directories Scanned             :   4,848
        Files/Directories Skipped             :   0
        Files/Directories Processed           :   4,848
            Identical Files/Directories       :   0
            New Files/Directories             :   4,848
            Updated Files/Directories         :   0
        Full Update Bytes Sent                :   355,762,776 (339 MB)
        Delta Update Bytes Sent               :   0 (0 B)
            Effective Data Bytes              :   0 (0 B)

      Publishing Phase Information
      ----------------------------
        Phase Started                         :   Feb 27 02:36
        Phase Completed                       :   Feb 27 02:36
        Elapsed Time                          :   00:00:00
```

```
        Database Records Scanned                 :   4,848
        Files/Directories Published              :   5,318
            New Files/Directories                :   4,848
            Renamed Files/Directories            :   0
            Updated Files/Directories            :   470
            Removed Files/Directories            :   0

Total processed:           4,848
Elapsed time:           00:10:01
**** Shadow Copy Report: DONE at Sat Feb 27 02:36:42 2010 ****
```

retry attempts

Purpose If a shadow-copy rule fails to transfer any files to the shadow volume, the rule makes multiple attempts to transfer the failed files. This retry mechanism can help a shadow-copy operation to succeed during a transient network failure, or during a failover at the target-volume site. Use the retry attempts command to set the retry limit for this rule.

Modes gbl-ns-vol-shdwcp

Security Role(s) storage-engineer or crypto-officer

Syntax **`retry attempts`** ***`count`***

count (0-2,147,483,647) is the total number of retries for a failed shadow copy before waiting for the next rule run.

Default(s) 5

Guidelines The shadow-copy rule keeps a list of all directories with failed files as it runs. After the run is complete, the rule checks the list and retries any directories that are in it. The rule repeats this process as often as you dictate with this command. If there are still failed transfers after the final retry attempt, the rule declares the run a failure.

The retry delay command sets the delay between these retry attempts. You can use these commands to compensate for transient network issues between your source and target sites.

Each retry attempt creates a separate shadow-copy report. The report shows the results of the retry operation. (Use the report (gbl-ns-vol-shdwcp) command to make the shadow-copy rule generate reports.)

Sample `bstnA(gbl-ns-vol-shdwcp[insur~/claims~insurSV])# retry attempts 10`

sets the shadow-copy rule, “SVrule,” to retry up to 10 times if it encounters any file-transfer failures.

Related Commands shadow-copy-rule
retry delay

retry delay

Purpose When a shadow-copy run fails to transfer one or more of its files, the rule pauses for a period of time before retrying the failed file(s). You can use the retry delay command to set the length of this pause.

Modes gbl-ns-vol-shdwcp

Security Role(s) storage-engineer or crypto-officer

Syntax **`retry delay`** ***`seconds`***

seconds (60-7200) is the number of seconds between retries.

Default(s) 300 (seconds)

Guidelines The retry attempts command sets the number of retry attempts. Use these commands to compensate for transient network issues between your source and target sites.

After the connection is re-established, the rule resumes at the last-completed directory.

Sample

```
bstnA(gbl-ns-vol-shdwcp[insur~/claims~insurSV])# retry delay 120
```

causes the shadow-copy rule, "insurSV," to pause for two minutes (120 seconds) between each retry attempt.

Related Commands shadow-copy-rule
retry attempts

schedule (gbl-ns-vol-shdwcp)

Purpose Use this schedule command to assign a schedule to the current shadow-copy rule.

Use no schedule to remove the rule's schedule.

Mode gbl-ns-vol-shdwcp

Security Role(s) storage-engineer or crypto-officer

Syntax

```
schedule name
no schedule
```

name (1-64 characters) identifies the schedule. Use show schedule for a list of configured schedules.

Default(s) None.

Guidelines A shadow-copy rule copies its fileset from the source volume to the shadow volume; each shadow-copy session occurs on a schedule. This command determines which schedule to use.

To create a schedule, use the gbl-mode schedule command.

You cannot use a schedule with a fixed duration; a shadow-copy rule must always run to completion in order to succeed.

Sample

```
prtlndA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])# schedule hourly
```

assigns the "hourly" schedule to the shadow-copy rule, "SVrule."

Related Commands shadow-copy-rule
schedule
show schedule

shadow

Purpose

Use the shadow command to transform a standard volume into a read-only shadow volume. A *shadow volume* is a target for a shadow-copy.

Use no shadow to make the volume ineligible as a shadow-copy target, and to make it writable.

Mode

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
shadow
no shadow
```

Default(s)

```
no shadow
```

Guidelines

The volume must be disabled (no enable (gbl-ns, gbl-ns-vol)) before you can use either shadow or no shadow.

A Shadow volume is typically created on a different switch from its source volume(s). To add a volume at a different switch, log into the CLI at that switch and add the volume from there.

A shadow volume does not keep its imported files; it uses its shares as repositories for the shadowed fileset(s). Therefore, the shadow volume deletes all extraneous files on its shares as soon as the shadow-copy rule first runs. The CLI prompts for confirmation before allowing this; to create the shadow volume, answer **yes** at the prompt.

The **no shadow** command issues a similar warning, indicating that shadow copies in progress are aborted if you continue. This means removal of some interim files and databases, possibly affecting the shadow volume. The source volume is unaffected. Again, answer **yes** to continue.

Clients can only have read access to the shadow volume. (Writes would corrupt the shadow copy.)

Guidelines: Configuration Rules

A volume cannot be used as a shadow volume if it contains any imported NetApp share with multiple qtrees. This is not an issue if each NetApp share is a single qtree, or if none of the NetApp shares contain any qtrees.

The shadow-copy-rule must write to the volume's metadata when it copies files; use modify to allow this.

Choose a shadow volume with at least as much disk capacity as its source volume(s), or some shadow-copy operations may fail.

A direct volume (also known as a presentation volume) cannot be used as a shadow volume.

CIFS subshares are not supported on the filers behind a shadow volume. Some filers do not allow the shadow-copy rule to delete a directory under a subshare, and this is necessary if the same directory was deleted at the source volume. For more information on CIFS subshares and how a managed volume supports them, see the documentation for the filer-subshares command.

A shadow volume must support all of the protocols for its source volume(s). For example, an NFSv3 and CIFS volume can be a shadow for another volume with both of those protocols, or for a CIFS-only volume, or for an NFSv3 volume. It cannot be the target for an NFSv2 volume.

Samples

```
prtlndA(gbl-ns-vol[archives~/cdn])# shadow

This will cause all the shares in the volume to be erased.

Are you sure you want to remove all data? [yes/no] yes
```

changes the "/cdn" volume into a shadow volume.

```
prtlndA(gbl-ns-vol[nemed~/testvol])# no enable
prtlndA(gbl-ns-vol[nemed~/testvol])# no shadow

This will cause any un-published files and shadow databases to be
erased.

Are you sure you want to proceed? [yes/no] yes
```

reverts the "/testvol" volume back to a managed volume.

Related Commands

namespace -> volume
no enable (gbl-ns, gbl-ns-vol)
modify

shadow-copy-rule

Purpose Use the shadow-copy-rule command to instantiate a shadow-copy rule in the current volume. A *shadow-copy rule* regularly copies a fileset from the current volume to a read-only *shadow volume*.

Use the no form of the command to delete the rule.

Modes gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
shadow-copy-rule name
no shadow-copy-rule name
```

name (1-1024 characters) is the name you choose for the rule.

Default(s) None

Guidelines When you first create a shadow-copy rule, a prompt asks for confirmation; enter **yes** to continue.

This command puts you into gbl-ns-vol-shdwcp mode, where you have various configuration commands for configuring the shadow-copy rule. Use the from fileset (gbl-ns-vol-shdwcp) command to select a fileset to copy. Use the target (gbl-ns-vol-shdwcp) command to choose a shadow volume as a destination for the file copies. The shadow volume must reside in the current RON; use show ron for a list of ARXes in the current RON. The schedule (gbl-ns-vol-shdwcp) command determines the schedule for making the shadow copies. Use the enable (gbl-ns-vol-shdwcp) command to start the rule.

If there are any network disconnects during a rule run, the rule pauses and retries to establish the connection. You can use the retry attempts and retry delay commands to manipulate the number of attempts and the length of time between them.

If a CIFS client has a file open on the source volume, the client must have opened the file with “shared read” access for the file to be shadow copied. Otherwise, the shadow copy fails for the current run of the rule. No application (including the shadow-copy rule) can open a file that is already opened, unless the first application opened with “shared read” access.

A volume cannot support a shadow-copy rule if it contains any imported NetApp share with multiple qtrees. This is not an issue if each NetApp share is a single qtree, or if none of the NetApp shares contain any qtrees.

Guidelines: CIFS Options (Local Groups and "8.3" Names)

If the filers behind the source volume use any Windows Local Groups or Local users, you may want to translate CIFS Security IDs (SIDs) between the source and shadow volumes. The sid-translation (gbl-ns-vol-shdwcp) command tells the rule to translate SIDs. This only applies to source/shadow volumes that support CIFS.

File systems that support CIFS also support an *alternate name* for any file or directory with a long file or directory name. The alternate-name support is for older Windows applications that only support "8.3" names. An 8.3 name is one with up to 8 characters and an optional extension with up to 3 characters. If a source-volume file has a primary name that matches a target-volume file's alternate name, the source-volume file could overwrite the target-volume file. This causes the shadow copy to fail for the given file or directory, unless you set cifs-8dot3-resolution for this rule. This is a recommended setting.

Guidelines: Options

The gbl-ns-vol-shdwcp mode contains some additional, optional commands:

- The report (gbl-ns-vol-shdwcp) command enables the shadow-copy rule to generate a report every time a shadow copy occurs. Reports are recommended, to keep a running log of all shadow copies.
- The optional delta-threshold command changes the minimum size of a file to be *delta-transferred* (as opposed to entirely transferred) to the shadow volume.
- You can use the publish group command to back out any shadow copy where any of its files fails to transfer properly.
- The optional bandwidth-limit command sets a limit on the amount of bandwidth used when the rule runs.
- After the initial run of the rule, it goes through a pruning phase to find extra files and directories at the target and remove them. If you are sure that the target has no extra files or directories, you can use the no prune-target command to avoid this extra processing.

Guidelines: Shadow-Volume Targets

To create a shadow volume, use the gbl-ns-vol shadow command at the target switch. This command changes a standard volume into a shadow volume.

To copy a source volume to multiple shadow-volume targets, configure one shadow-copy-rule for each target.

Guidelines: Schedule

To create a schedule to be applied to the shadow-copy rule, use the gbl schedule command at the source switch. The rule fires as dictated by its schedule: if the schedule has a start time that is earlier than now, the first shadow copy begins as soon as you enable the rule.

If you make a configuration change when the rule is running on a schedule, the configuration change is ineffective until the next time the rule runs. This includes changes to the rule's behavior between scheduled runs, such as a change to inline-notify behavior. You can make the change take effect immediately by using no enable (gbl-ns-vol-shdwcp) followed by enable on the rule.

Guidelines: Reports

As mentioned above, we recommend that you use the shadow-copy report feature to keep a detailed log of all shadow copies. The reports appear on the source switch: the show reports command shows all reports on the switch, including shadow-copy reports. You can use the standard file-management commands with these reports: delete, rename, show reports file-name, tail, and/or grep.

Sample

```
prtlndA(gbl-ns-vol[wwmed~/acct])# shadow-copy-rule SVrule
This will create a new policy object.

Create object 'SVrule'? [yes/no] yes
prtlndA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])#
```

instantiates a new shadow-copy rule, “SVrule,” for the current volume.

Related Commands

from fileset (gbl-ns-vol-shdwcp)
target (gbl-ns-vol-shdwcp)
retry attempts
retry delay
schedule (gbl-ns-vol-shdwcp)
report (gbl-ns-vol-shdwcp)
sid-translation (gbl-ns-vol-shdwcp)
delta-threshold
publish
bandwidth-limit
prune-target
inline-notify
enable (gbl-ns-vol-shdwcp)
schedule
show reports

show shadow

Purpose Use the show shadow command to see the current status of one or more shadow-copy rules.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show shadow
show shadow namespace
show shadow namespace vol-path
show shadow namespace vol-path rule-name
```

namespace (1-30 characters) identifies the namespace with the source volume.

vol-path (1-1024 characters) focuses on one source volume in the namespace.

rule-name (1-1024 characters) focuses on one rule.

Guidelines Each rule starts with the following fields, for identification:

Namespace

Source Volume

Shadow Rule

Report File, where the prefix is chosen by the report (gbl-ns-vol-shdwcp) command.

Fileset, which is chosen by the from fileset (gbl-ns-vol-shdwcp) command.

Delta Threshold shows the minimum size for a file to qualify for delta-transfers. A delta transfer is a copy of only the changed portions of a file. This is set with the delta-threshold command.

Publishing Mode indicates whether or not all of the files are published together as a "group" if (and only if) all of them are successfully transferred. This can be "group" or "individual," as set by the publish command.

CIFS SID Translation is "Fail on Errors," "Ignore Errors," or "Disabled." If either of the first two options, the shadow-copy rule translates CIFS Security IDs (SIDs) from source-volume filers to their equivalent SIDs at the shadow-volume filers. This enables local-group support at both filers. "Fail on Errors" indicates that the rule refuses to copy any file or directory where the SID translation fails. You control this with the sid-translation (gbl-ns-vol-shdwcp) command.

Shared Access Allowed indicates that the shadow-copy rule allows other applications to write to a file during the copy operation. This makes it possible for the rule to copy an open file.

Inline Notifications appears if the shadow-copy rule monitors the source volume for inline changes. An *inline change* is one that occurs due to a client operation, such as a file rename. If the rule tracks these changes, it can avoid a full scan of the source volume the next time the rule runs. Use the inline-notify command to set this.

Prune Target only appears if the rule prunes empty directories from the target volume after a successful shadow copy. Use the prune-target command to determine this.

Guidelines (Cont.)

Bandwidth Limit only appears if it was set with the bandwidth-limit command.

Retry Attempts shows the number of retries the rule makes if a shadow-copy run is interrupted by a network disconnect. This only applies to a rule with a target volume on another ARX. Use the retry attempts command to change this number.

Retry Delay shows the number of seconds between retry attempts. Use the retry delay command to change this time.

Database Location is either "metadata-share" or "volume." This indicates where the shadow-copy rule keeps its database. A setting of "metadata-share" is recommended. Use the database-location command to change this setting.

CIFS 8dot3 Resolution appears if the rule can copy files with "8.3" names that collide with other files at the shadow volume. This is recommended for all shadow-copy rules in volumes that support CIFS. Use the cifs-8dot3-resolution command to set this.

These are followed by three fields to show the latest run time and status for the rule:

Processing Started states the date and time.

Processing Completed only appears when the shadow copy has concluded.

Elapsed Time is stated in hours, minutes, seconds.

The next several fields describe the shadow-copy run at a high level:

Operating Mode is "Full tree walk" (examine the full source-volume tree with every shadow-copy session) or "Inline notification" (keep track of directories that change between shadow copies, and only examine the changed directories on the next shadow copy). Use the inline-notify command to change the mode.

Tree Walk Reason is present if the shadow volume rule is performing a full tree walk of the source volume (instead of using inline notification of changes) and indicates why the tree walk is being performed. Possible reasons include:

- Initial run means that it is the first run of the rule, so a tree walk is necessary.
- Initial run (resumed) is caused by the ARX rebooting before the first tree-walk completes.
- Rule started indicates the rule was disabled and then re-enabled.
- Rule modified indicates the rule parameters changed.
- Publishing database rebuild is caused by an nsck ... rebuild of the source or target volume requiring the target's publishing database to be rebuilt.
- Inline notification overflow occurs if a large number of notification events were generated and the database used to store these events overflowed.
- Previous run failed indicates the previous run of the rule encountered an error.
- Inline notification disabled indicates that someone used no inline-notify to disable inline notification.

Source Status appears if there is a problem with the source volume. The problem is summarized here.

Current Phase shows the phase of the current shadow-copy run. See Guidelines: Shadow-Copy Phases for a list of all possible settings in this field.

Guidelines (Cont.)

Target Information lists the target for the current shadow-copy rule. This shows the status of the most-recent (or current) shadow copy to the target volume. See Guidelines: Target Status for a list of all possible status settings.

Copy Phase Information shows details about the copy phase. This is when the rule copies all files from the source volume to the shadow volume. If publish group is in effect, all files must copy successfully before going into "Publishing" phase, below.

Publishing Phase Information shows the files that have moved to the final locations. If publish group is in effect, this phase does not start until the Copy Phase ends. If the rule is set to publish individual, this table updates in tandem with the Copy Phase table.

Pruning Phase Information appears only for the first run of the shadow-copy rule, if at all. In this phase, the shadow volume scans its directory tree and prunes away all files and directories that are not on the source volume, too. For cases where the source and shadow volumes are known to be identical from the outset, you can disable this phase with no prune-target.

Cleaning Phase Information also appears very rarely. This means removing empty directories and files that the shadow-copy rule created in the staging area. This is only necessary in a limited number of cases.

Guidelines: Shadow-Copy Phases

Current Phase goes through the following settings, in order:

- Inactive - The rule is enabled, but has not yet run.
- Initializing - The rule is starting.
- Connecting - The source switch is connecting to the target switch.
- Opening Database - The rule is opening a database at the target switch.
- Reindexing Database only occurs rarely, after a software upgrade or a failure at the target site. The rule is writing an index file for fast access to the target database.
- Copying then Publishing, or Copying/Publishing - The Copying phase is when the rule copies files from the source volume to a *staging area* on the shadow volume. All files must copy successfully before the rule enters its Publishing phase. During the Publishing phase, the rule moves the files from the staging area to their final locations. These two phases occur when publish group is in effect.

  If publish is set to individual, this phase is Copying/Publishing; each file is copied to the staging area and then immediately moved to its final location.
- Pruning appears only for the first run of the shadow-copy rule, if at all. In this phase, the shadow volume scans its directory tree and prunes away all files and directories that are not on the source volume, too. For cases where the source and shadow volumes are known to be identical from the outset, you can disable this phase with the no prune-target command.
- Cleaning also appears very rarely. This means removing empty directories and files from a staging area behind the target volume. This is only necessary in a limited number of cases.
- Closing Database is the final step before a completed run.
- Completed - The operation is finished. Success or failure is indicated under Target Information, below.

Some phases may go by too quickly for you to ever see them on this screen. If the phase is Unknown, there may be a system anomaly. Disabled appears if the shadow-copy rule is disabled.

Guidelines: Target Status

The Target Information field shows the status of the most-recent (or current) shadow copy to the target volume:

- Successful indicates that the shadow-copy run finished and completely succeeded.
- Error(s) detected means that some problem occurred. Details appear in the shadow-copy report for the run. The show reports command shows all reports on the switch, including shadow-copy reports.
- Unable to establish connection indicates that the remote switch is unreachable through the Resilient Overlay Network (RON). Use show ron to view the status of your RON connections.
- Error(s) detected, Connection failed shows that the RON connection with the remote switch was lost after the shadow-copy run started. As above, use show ron to view the status of your RON connections.
- Target volume is offline indicates a problem at the remote ARX. Use show namespace at the remote ARX to find the status of the shadow volume.
- Target is not a shadow volume means that the shadow-copy rule has the wrong target (gbl-ns-vol-shdwcp), or that the target volume itself is not configured as a shadow volume. To convert a managed volume to a shadow volume, use the shadow command.
- Target volume is no modify shows that the shadow volume has the no modify setting. This prevents the shadow-copy rule from writing to the shadow volume. At the target switch, you can use the modify command in the shadow volume to allow write access.
- Target volume has insufficient space indicates that the file servers behind the shadow volume do not have enough space to hold a replica of the source volume. Consider using a larger share behind the shadow volume, or adding additional shares.
- Target host does not exist,
- Target namespace does not exist, and
- Target volume does not exist all indicate a bad target in the shadow-copy rule. Use the target (gbl-ns-vol-shdwcp) command to change the shadow-volume target.
- Target volume is unavailable indicates a problem at the remote ARX. Use show namespace at the remote ARX to find the status of the shadow volume.
- Target host is not reachable indicates that the remote switch is unreachable through the network. Use the ping command and/or expect traceroute to check the connection to the remote ARX.
- Target version not supported means that the software release on the target ARX is incompatible with the release on the source ARX. From either CLI, use show version to see the currently-running release of software. You can install new software through the CLI; refer to the ARX CLI Maintenance Guide for software-upgrade instructions.

Sample

```
bstnA# show shadow wwmed
```

shows the status of all shadow copies from the wwmed namespace. See Figure 29.2 on page 29-28 for sample output.

Related Commands

shadow

Figure 29.2 *Sample Output: show shadow wwmed*

```
bstnA# show shadow wwmed

Shadow Copy Status
==================

    Namespace                                 :   wwmed
    Source Volume                             :   /acct
    Shadow Rule                               :   SVrule
    Report File                               :   SVetc_201002270226.rpt
    Fileset                                   :   worthSaving (Files and Directories)
    Delta Threshold                           :   100 MB
    Publishing Mode                           :   Individual
    CIFS SID Translation                      :   Ignore Errors
    Shared Access Allowed                     :   Yes
    Inline Notifications                      :   Yes
    Bandwidth Limit                           :   5.0 Mb/s (625.0 kB/s)
    Retry Attempts                            :   5
    Retry Delay                               :   300 seconds
    Database Location                         :   metadata-share

    ================================================================================
      Processing Started                      :   Feb 27 02:26
      Processing Completed                    :   Feb 27 02:36
      Elapsed Time                            :   00:10:01
      Operating Mode                          :   Full tree walk
      Tree Walk Reason                        :   Initial run
      Current Phase                           :   Completed

      Target Information
      ------------------
        prtlndA: nemed:/acctShdw/             :   Successful

      Copy Phase Information
      ----------------------
        Phase Started                         :   Feb 27 02:26
        Phase Completed                       :   Feb 27 02:36
        Elapsed Time                          :   00:09:59
        Average Transmission Rate             :   4.7 Mb/s (599.4 kB/s)
        Total Files/Directories               :   4,848
        Files/Directories Scanned             :   4,848
        Files/Directories Skipped             :   0
        Files/Directories Processed           :   4,848
            Identical Files/Directories       :   0
            New Files/Directories             :   4,848
            Updated Files/Directories         :   0
        Full Update Bytes Sent                :   355,762,776 (339 MB)
        Delta Update Bytes Sent               :   0 (0 B)
            Effective Data Bytes              :   0 (0 B)

      Publishing Phase Information
      ----------------------------
        Phase Started                         :   Feb 27 02:36
        Phase Completed                       :   Feb 27 02:36
        Elapsed Time                          :   00:00:00
        Database Records Scanned              :   4,848
        Files/Directories Published           :   5,318
            New Files/Directories             :   4,848
            Renamed Files/Directories         :   0
            Updated Files/Directories         :   470
            Removed Files/Directories         :   0
```

sid-translation (gbl-ns-vol-shdwcp)

Purpose Behind either the source or the shadow volume, one of the back-end CIFS filers may be configured for *Local Groups*. These filers use local Security IDs (SIDs) for their group names instead of those issued by a Domain Controller (DC). You can configure a shadow-copy rule to translate all local SIDs to their corresponding names at the source filer, then translate the names back to SIDs at the target filer. This preserves local-group privileges, so that members of a local group at the source filer can access copies of their files and directories at the shadow volume. Use the sid-translation command to enable SID translation for the current rule.

Use no sid-translation to disable SID translation for this shadow-copy rule.

Mode gbl-ns-vol-shdwcp

Security Role(s) storage-engineer or crypto-officer

Syntax

```
sid-translation [fail-on-errors]
no sid-translation
```

fail-on-errors (optional) causes the rule to abort any file or directory copy where the SID translation fails. Translation failures can occur if a local group on the source filer does not exist on the target filer. For SID translation to reliably succeed, all local groups on the source volume's filers must be duplicated on the shadow volume's filers. If SID translation fails and you omitted this option, the rule copies the SID to the shadow volume without translating it; the Access-Control Entry (ACE) is therefore unusable while on the shadow volume, though it *will* be usable if someone later copies the file back to the source volume.

Default(s) `no sid-translation`

Guidelines

This command does not apply to a rule in an NFS-only namespace.

While SID translation is enabled, the rule translates each SID to a group name (at the source volume), then translates the group name back to a SID (at the target volume). This ensures that the final SID is one that is properly configured at the shadow volume.

SID translations are only needed if

- any filers behind the source volume have local groups defined, and
- CIFS clients in those local groups require read access to the files at the shadow volume.

Before you enable SID translations, all filers behind both volumes should support all of each other's local-group names. These are the filers behind the source volume *and* the shadow volume. If a file tagged with local-group name, "staff", is copied to remote filer without any definition for a "staff" group, the shadow-copy rule copies the SID for "staff" verbatim; this SID will be unusable at the destination filer. To avoid these translation failures, duplicate all local group names and user names on all of the filers behind both volumes.

The gbl-ns-vol-shr ignore-sid-errors command is designed for filers that return an error for an unknown SID but accept the file or directory anyway. If this is set for a share in the shadow volume, the shadow-copy rule ignores SID errors from that filer, even if you used fail-on-errors in the sid-translation command.

If the source and shadow volume are in two different Windows domains, your site may require additional preparation. Contact F5 Support if you want to configure SID translation for an inter-domain shadow-copy rule.

Use the gbl-ns-vol-shr sid-translation command to perform SID translations for file migrations within the volume.

Sample

```
bstnA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])# sid-translation
```

enables SID translation for the rule named "SVrule."

Related Commands

shadow-copy-rule
ignore-sid-errors
sid-translation

target (gbl-ns-vol-shdwcp)

Purpose Use the target command to choose a shadow-volume target for the current shadow-copy rule. A shadow-copy rule copies a fileset from the current volume to a shadow volume.

Mode gbl-ns-vol-shdwcp

Security Role(s) storage-engineer or crypto-officer

Syntax **`target [hostname`** ***`host`*** **`[namespace`** ***`name`*****`]] volume`** ***`shadow`*** **`[path`** ***`sub-path`*****`]`**

hostname ***host*** (optional, 1-32 characters) identifies the host for the shadow volume. The shadow volume must reside on the current switch or another ARX in the same RON. Use the show ron command to view all hosts in the RON.

namespace ***name*** (optional, 1-30 characters) identifies the shadow volume's namespace.

shadow (1-256 characters) identifies the shadow volume itself.

path ***sub-path*** (optional, 1-1024 characters) selects a subdirectory in the shadow volume.

Default(s) ***host*** - the local switch

name - the current namespace

Guidelines This command sets a storage target for the current shadow-copy rule.

A target volume must be configured as a shadow volume; use the gbl-ns-vol shadow command to change a standard volume into a read-only shadow volume. To use another switch as the host for the shadow volume, log into the CLI on that switch and add the volume from there.

To consolidate multiple shadow copies into a single volume, use the optional **path** argument. Multiple volumes can use the same shadow volume as a target, where each source volume is copied to a unique path. For example, the /etc volume can have a target volume and path of /shadow/etc while the /var volume has a target and path of /shadow/var. Both source volumes are thereby copied to the /shadow volume. This simplifies backups for the two volumes.

If you use paths, do not overlap them. For example, you cannot use a target of /shadow/var for one source volume in addition to /shadow/var/log for another.

A target volume must support all of the protocols for its source volume(s). For example, an NFSv3 and CIFS volume can be a shadow target for another volume with both of those protocols, or for a CIFS-only volume, or for an NFSv3 volume. It cannot be the target for an NFSv2 volume.

If the volume contains any imported NetApp share with multiple qtrees, it cannot be used as the source or target of a shadow-copy rule. This is not an issue if each NetApp share is a single qtree, or if none of the NetApp shares contain any qtrees.

Samples

```
bstnA(gbl-ns-vol-shdwcp[wwmed~/acct~SVrule])# target hostname prtlndA
namespace nemed volume /acctShadow
```

uses /acctShadow as the target for “SVrule.” This shadow volume resides on a remote switch, “prtlndA,” in the same RON.

```
bstnA(gbl-ns-vol-shdwcp[archives~/usr~buUsers])# target namespace
allBackups volume /bu path /usr
```

uses “allBackups~/bu/usr” as the target for the “buUsers” rule.

```
bstnA(gbl-ns-vol-shdwcp[archives~/log~buLogs])# target namespace
allBackups volume /bu path /log
```

uses “allBackups~/bu/log” as the target for the “buLogs” rule.

Related Commands

shadow-copy-rule
shadow

30

Snapshots

clear statistics snapshot

Purpose Volumes keep statistics on all of their snapshot operations. A *snapshot* is a full copy of the volume at a particular moment in time. A *snapshot rule* determines the names for a set of snapshots, the number of snapshots to retain, and (optionally) a schedule for taking snapshots. Use this command to clear the cumulative statistics for a particular snapshot rule.

Mode priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`clear statistics snapshot`** ***`namespace volume snapshot-rule`***

namespace (1-30 characters) identifies a namespace.

volume (1-1024 characters) is a volume that supports snapshots.

snapshot-rule (1-1024 characters) is the rule whose statistics should be cleared.

Default(s) None

Guidelines Thc CLI prompts for confirmation before clearing any snapshot statistics. Enter **yes** to proceed.

The show policy ... details command shows cumulative statistics for snapshot operations. Use this command to clear the statistical counters for one snapshot rule.

Sample

```
bstnA# clear statistics snapshot medarcv /lab_equipment dailySnap

Confirming this command causes all historic statistical data for the
snapshot rule to be reset.
Proceed? [yes/no] yes
```

clears all snapshot statistics for the “dailySnap” rule.

Related Commands show policy

enable (gbl-ns-vol-...snap)

Purpose

A *snapshot rule* is the basic configuration object for taking snapshots, or point-in-time copies, of an ARX volume. It must be enabled to create any ARX snapshots, either manually or by a schedule. Use the enable command to enable the current snapshot rule.

Use no enable to disable the current snapshot rule.

Mode

gbl-ns-vol-snap

gbl-ns-vol-replica-snap

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
enable
no enable
```

Default(s)

```
no enable
```

Guidelines

You must enable a rule for the policy engine to use it, or to invoke it manually with snapshot create.

Samples

```
bstnA(gbl-ns-vol-snap[access~/G~nightly])# enable
```

enables the current snapshot rule.

```
bstnA(gbl-ns-vol-snap[access~/G~hourly])# no enable
```

disables the “hourly” snapshot rule in the same volume.

Related Commands

namespace -> volume -> snapshot rule
namespace -> volume -> snapshot replica-snap-rule
snapshot create

exclude

Purpose

You can use the exclude command to exclude a particular volume share from the current snapshot or replica-snap rule.

Use no exclude to return to including a share in this volume's coordinated snapshots.

Mode

gbl-ns-vol-snap

gbl-ns-vol-replica-snap

Security Role(s)

storage-engineer or crypto-officer

Syntax

`exclude` ***`share-name`***
`no exclude` ***`share-name`***

share-name (1-64) identifies the share to exclude. This is the share name from the ARX-volume configuration, not the filer configuration.

Default(s)

`no exclude`

Guidelines

The no manage snapshots command excludes an entire filer from any ARX-snapshot operations, typically because the filer does not support ARX snapshots. That command is sufficient for most installations. Use this command, exclude, only under the advisement of F5 Support; it creates sparse snapshots that may be an issue later if clients need files restored. It is intended for filers where a particular back-end volume or file system is near capacity, but others on the same filer are not.

By default, a snapshot rule creates snapshots on all of the back-end shares that support them. You configure a filer to support snapshots with some external-filer commands, such as filer-type and manage snapshots. For back-end volumes that are currently nearing their full capacity, you can use this command to exclude them from this rule's snapshots. The negative form of the command, no exclude, can later re-instate the share for inclusion in ARX snapshots.

Files in an excluded share do not appear in the volume's coordinated snapshot(s). For example, if the excluded share contains "\bigDir\bigFile.wmv," clients cannot recover "bigFile.wmv" from the ARX view of the snapshot. If this is a concern at your site, you can use one or more place-rules to control which files reside on your excluded shares. We recommend that you contact F5 Support for specific guidance.

Guidelines: Snapshot Rules and Replica-Snap rules

A snapshot rule coordinates the snapshots on your standard back-end shares, and a snapshot replica-snap-rule coordinates snapshots on special replica-snap shares. A replica-snap share is a replica of one of the managed-volume's standard shares on a cheaper filer; the cheaper filer creates and stores an ever-growing collection of snapshots from that replicated share. From the perspective of an ARX client, standard snapshots are interleaved with replica snaps, and the two snapshot types are indistinguishable.

A snapshot rule always excludes all replica-snap shares, and a replica-snap rule always excludes all standard snapshot shares. The exclude command cannot change these rules. For example, you can exclude a replica-snap share from a replica-snap rule, but you cannot exclude the same share from a standard snapshot rule; it is already excluded by definition.

Sample

```
bstnA(gbl-ns-vol-snap[access~/G~daily])# exclude ronin4
```

excludes the “ronin4” share from the snapshot rule named “daily.”

Related Commands

namespace -> volume -> snapshot rule
namespace -> volume -> snapshot replica-snap-rule

report (gbl-ns-vol-...snap)

Purpose

Use this command to enable progress reports for the current snapshot or replica-snap rule.

Use no report to prevent progress reports.

Mode

gbl-ns-vol-snap

gbl-ns-vol-replica-snap

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
report file-prefix
no report
```

file-prefix (1-1024 characters) sets a prefix for all snapshot reports. We recommend starting all of your snapshot-report prefixes with a common string, such as "snap_"; this commonality helps you to prepare for snapshot reconstitution (described below). Each report has a unique name in the following format:

*file-prefix*_0_create_*YearMonthDayHourMinuteSecondsMilliseconds*.rpt

For example, snap_nightly_0_create_20070930130022142.rpt could be the name for one report with the prefix, "snap_nightly."

Default(s)

`no report`

Guidelines

Every time a snapshot or replica-snap rule fires and creates a coordinated snapshot, the rule can generate a report to show the details of the snapshot-create operation. This includes details about each snapshot at each of the volume's back-end filers. The reporting feature is disabled by default; we recommend that you enable reporting for all snapshot and replica-snap rules, to diagnose any problems that occur with snapshots and to make snapshot reconstitution possible.

Use show reports for a list of reports, or show reports *file-name* to show the contents of one report.

See Figure 30.1 on page 30-9 for a sample snapshot-create report.

Guidelines: Preparing for Snapshot Reconstitution

There are situations that require a transfer of coordinated snapshots from one ARX to another, such as the transfer from the ARX at a primary site to another at a disaster-recovery site (see cluster-name and activate configs for details about site-to-site failovers). A challenge for this task is mapping multiple back-end snapshots to each ARX snapshot. The process of *snapshot reconstitution* meets this challenge by parsing snapshot reports and producing a CLI script of snapshot manage commands. Each snapshot manage command in the CLI script pulls one back-end snapshot into its corresponding ARX snapshot.

To prepare for snapshot reconstitution, save a copy of every snapshot report and replica-snap report. Use the at command together with copy ftp, copy scp, copy tftp, or copy smtp to regularly copy snapshot reports off of the ARX. Use the common string for all of your snapshot reports, together with an "*" or other wildcard. The report repository should always hold the latest snapshot reports, so that they have the latest back-end snapshot names. For example, this command copies all reports starting with "snap" to an external IP each morning:

```
bstnA(cfg)# at 01:19:18 every 1 day do "copy reports snap*
ftp://ftpuser:ftpuser@172.16.100.183//var/arxSnapRpts/ format xml"

  The scheduled execution time for CLI command is: 3/30/09 1:19 AM.
```

You also need a copy of the "snap-recon.pl" script from the "software" directory on the ARX:

```
bstnA# copy software snap-recon.pl
ftp://ftpuser:ftpuser@172.16.100.183//var/arxSnapRpts/
```

Choose a filer with a large amount of free space, or monitor the filer's space frequently and remove outdated snapshot reports. Every snapshot report has a timestamp in its name, so it is unique.

The filer should be able to run Perl scripts, and requires the XML::Simple module. You can download this Perl module from CPAN (http://search.cpan.org) if your system does not already have it.

Guidelines: Reconstituting ARX Snapshots

You can perform snapshot reconstitution if you have the snap-recon.pl script and the latest set of snapshot reports on a host that supports Perl. Start by running the snap-recon.pl script on that host. This produces a CLI script with a sequence of snapshot manage commands. By default, the output script is named, "snapRecon.cli." This script has several options; execute the command without any options to get a complete list. You must use the "**--report-dir *directory***" clause to specify the *directory* that holds the reports. For example, this command sequence lists the files on "client2:/var/arxSnapRpts," runs snap-recon.pl on the reports in the current directory (.), and then shows the new file in the directory:

```
juser@client2:/var/arxSnapRpts$ ls
snap_daily_0_create_20090330010418883.xml  snap_hourly_0_create_20090330010652616.xml
snap_daily_0_create_20090330010454303.xml  snap_hourly_0_create_20090330010723024.xml
snap_daily_0_create_20090330010524663.xml  snap-recon.pl
snap_daily_2_remove_20090330010612220.xml
juser@client2:/var/arxSnapRpts$ ./snap-recon.pl --report-dir .
juser@client2:/var/arxSnapRpts$ ls
snap_daily_0_create_20090330010418883.xml  snap_hourly_0_create_20090330010652616.xml
snap_daily_0_create_20090330010454303.xml  snap_hourly_0_create_20090330010723024.xml
snap_daily_0_create_20090330010524663.xml  snapRecon.cli
snap_daily_2_remove_20090330010612220.xml  snap-recon.pl
juser@client2:/var/arxSnapRpts$
```

Guidelines: Reconstituting ARX Snapshots (Cont.)

You have the option to edit the CLI script before you run it on the ARX. For example, the ARX volumes at a disaster-recovery site may have different names than those at the primary site. The snapRecon.cli script has several options for substituting these names; you can use these to rebuild the CLI script, or you can manually edit the CLI script.

Once the CLI script is ready, you can download it to the ARX and run it. Use copy ftp, copy {nfslcifs}, or copy scp for the download, and use run to run it.

A large, complex CLI script may contain errors. If you discover any back-end snapshots that are mismatched with their ARX counterparts, you can use the snapshot clear command to remove the ARX snapshot (not the rule) from the configuration. Then you can edit and re-run the script, or you can use snapshot manage to manually incorporate the back-end snapshots into the correct ARX snapshots.

Samples

```
bstnA(gbl-ns-vol-snap[access~/G~nightly])# report snap_G
```

enables reports for the snapshot rule, "nightly."

```
bstnA(gbl-ns-vol-snap[access~/G~fileHist])# report fh_G error-only
```

enables reports for the snapshot rule, "fileHist." This rule, presumably designed to record file-history information, only generates reports for snapshots with errors.

Related Commands

namespace -> volume -> snapshot rule
namespace -> volume -> snapshot replica-snap-rule
show reports

Figure 30.1 *Sample Report: snap_daily_0_create_....rpt*

```
bstnA# show reports snap_daily_0_create_20120229004340055.rpt

  Snapshot Rule Summary
  ---------------------
   Namespace Name:               medarcv
   Volume Name:                  /lab_equipment
   Snapshot Rule Name:           dailySnap


  Snapshot Properties
  -------------------
   Snapshots Enabled:            Yes
   Guarantee Consistency:        Disabled
   Retain Count:                 7
   Schedule:                     daily4am
   CIFS Directory Name:          ~snapshot
   Directory Display:            All Exports
   Hidden File Attribute:        Not Set
   Restricted Access Configured: Yes
   VSS Mode:                     None
   Contents:
     Metadata:                   No
     Volume Configuration:       No
     User Snapshots:             Yes
   Archive:
     Total Archive Operations:   0
     Total Successful Operations: 0
     Total Failed Operations:    0
     Total Saved Metadata:       0 B
     Total Saved Volume Config:  0 B
```

```
   Average Copy Rate:           0 b/s

Snapshot Summary - dailySnap_0
------------------
 Snapshot Name:                 dailySnap_0
 Snapshot Operation:            Create
 Result:                        Success
 Time Requested:                02/29/2012 00:43:40 -0500
 Time Created:                  02/29/2012 00:43:40 -0500
 Last Time Verified:
 Request:                       Create
 Snapshot State:                Sparse
 Snapshot Origin:               Manual
 Report Name:                   snap_daily_0_create_20120229004340055.rpt


Included Shares
----------------
     Share Name:                equip (user data)
     Filer:
       Name:                    nas10
       CIFS Share:              equipment
       Volume:                  vol2
       Filer Snapshot:          acopia_1_201202290543_d9bdece8-9866-11d8-91e3-f48e42637d58_vol2

     Share Name:                leased (user data)
     Filer:
       Name:                    nas10
       CIFS Share:              for_lease
       Volume:                  vol1
       Filer Snapshot:          acopia_1_201202290543_d9bdece8-9866-11d8-91e3-f48e42637d58_vol1

     Share Name:                backlots (user data)
     Filer:
       Name:                    fs2
       CIFS Share:              backlot_records
       Volume:                  E:\
       Filer Snapshot:          {62a12821-d854-4279-87fb-0a74903cec91}
       Time Created:            02/29/2012 00:43:41 -0500


Excluded Shares
----------------
     Share Name:                scanners
     Filer:
       Name:                    fs5
       CIFS Share:              xraysScanners
     Reason:                    Snapshots were not supported on this type of back-end filer.
```

retain

Purpose

Each snapshot rule retains some maximum number of *snapshots*, or point-in-time copies, of its volume's files. Use this command to set the number of retained snapshots for the current rule.

Use no retain to revert to the default retention count.

Modes

gbl-ns-vol-snap

gbl-ns-vol-replica-snap

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
retain snap-count
no retain
```

snap-count (1-1024) is the maximum number of snapshots for this rule to retain.

Default(s)

`3`

Guidelines

Whenever the snapshot rule runs, it creates a new snapshot "0" (zero). The former snapshot "0" becomes snapshot "1," the former snapshot "1" becomes snapshot "2," and so on. The rule deletes the oldest snapshot if keeping it would exceed the retain count. For example, if the rule retains 5 snapshots, it only keeps snapshots 0-4. It removes the older snapshots from client presentations as well as the back-end filers. This prevents wasted snapshot space on the filers.

If you increase the retention count in a rule that already has a full count of snapshots, the rule does not delete any snapshots until it exceeds the count again. For example, if the retain count increases from three to four, the next time the rule runs it moves snapshot "3" to "4" and does not delete any snapshots. On the following run of the rule, it deletes snapshot "4" to avoid exceeding the retain count.

If you decrease the retain count in a rule that has already reached it, the volume removes all excess snapshots the next time the rule runs. The CLI prompts for confirmation before accepting this directive.

Guidelines: Snapshot Rules and Replica-Snap rules

A snapshot rule coordinates the snapshots on your standard back-end shares, and a snapshot replica-snap-rule coordinates snapshots on special replica-snap shares. A replica-snap share is a replica of one of the managed-volume's standard shares on a cheaper filer; the cheaper filer creates and stores an ever-growing collection of snapshots from that replicated share. From the perspective of an ARX client, standard snapshots are interleaved with replica snaps, and the two snapshot types are indistinguishable.

The replica-snap filer(s) are designed to hold larger numbers of snapshots than standard shares, so we recommend a much higher retain count for a snapshot replica-snap-rule than a snapshot rule.

Sample

```
bstnA(gbl-ns-vol-snap[access~/G~nightly])# retain 7
```

retains seven snapshots for the rule named "nightly."

Related Commands

namespace -> volume -> snapshot rule
namespace -> volume -> snapshot replica-snap-rule

schedule (gbl-ns-vol-...snap)

Purpose

Use this schedule command to assign a schedule to the current snapshot (or replica-snap) rule.

Use no schedule to remove the rule's schedule.

Mode

gbl-ns-vol-snap

gbl-ns-vol-replica-snap

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
schedule name
no schedule
```

name (1-64 characters) identifies the schedule. Use show schedule for a list of configured schedules.

Default(s)

None.

Guidelines

A snapshot rule can take *snapshots* (point-in-time copies) of an ARX volume on a regular schedule. This command determines which schedule to use, if any.

Whether or not a schedule is assigned to the snapshot (or replica-snap) rule, you can manually invoke snapshots with the snapshot create command.

To create a schedule, use the gbl-mode schedule command.

Guidelines: Snapshot Grouping

If more than one volume uses the same schedule, the ARX aggregates all of the filer snapshots at once. This is called *snapshot grouping*. By grouping the filer snapshots, the ARX avoids any duplicate snapshots on a given back-end volume.

Snapshot grouping may create a problem for volumes backed by EMC Celerra servers, which serialize their snapshots. Consider a site with 10 volumes backed by 10 different back-end volumes on a single EMC server. If all 10 ARX volumes use the same snapshot schedule, each ARX snapshot must wait for all 10 EMC checkpoints to finish, one after the other.

Each ARX volume must use the same schedule (by name) for snapshot grouping to occur. If you use schedules with different names but the same dates and times, the ARX volumes independently perform their snapshot operations. This could result in duplicate snapshots on the same back-end volume, but it avoids the EMC-serialization issue.

Guidelines: Snapshot Rules and Replica-Snap rules

A snapshot rule coordinates the snapshots on your standard back-end shares, and a snapshot replica-snap-rule coordinates snapshots on special replica-snap shares. A replica-snap share is a replica of one of the managed-volume's standard shares on a cheaper filer; the cheaper filer creates and stores an ever-growing collection of snapshots from that replicated share. From the perspective of an ARX client, standard snapshots are interleaved with replica snaps, and the two snapshot types are indistinguishable.

We recommend using a different schedule (by name) for a replica-snap rule than you use for a standard snapshot rule. The snapshot grouping (above) should be different between snapshot shares and replica-snap shares.

Samples

```
bstnA(gbl-ns-vol-snap[access~/G~nightly])# schedule daily4am
```

uses a schedule named “daily4am” for the “nightly” snapshot rule.

```
bstnA(gbl-ns-vol-snap[medarcv~/lab_equipment~test_snap])# no schedule
```

removes any schedule from the “test_snap” rule.

Related Commands

namespace -> volume -> snapshot rule
namespace -> volume -> snapshot replica-snap-rule
snapshot create
schedule
show schedule

show snapshots

Purpose Use the show snapshots command to see the current status of one or more snapshot rules. This shows snapshots that have been created by a snapshot rule, not snapshots created at the back-end filer.

Mode (any)

Security Role(s) operator or backup-operator (any)

Syntax

```
show snapshots
show snapshots namespace
show snapshots namespace vol-path
show snapshots namespace vol-path snapshot-rule
```

namespace (1-30 characters) identifies the namespace with the snapshot rule(s).

vol-path (1-1024 characters) focuses on one volume in the namespace.

snapshot-rule (1-1024 characters) focuses on one snapshot rule. This changes the output to a detailed view of the rule.

Guidelines: Summary View The simplest output shows a summary of all selected *snapshots*, or point-in-time copies of an ARX volume. For each volume, this contains a table with one line per snapshot. Each line contains the following fields:

Rule identifies the rule.

Type is the particular type of snapshot rule: Snapshot (snapshot rule), Replica (snapshot replica-snap-rule), or Notification (notification rule).

Name is the name of the particular ARX snapshot. This is typically the rule name with an integer ID that indicates the order of the snapshots ("0" is the newest, "1" is the next-newest, and so on).

Created indicates the time when the final filer finished its snapshot. If a snapshot is currently underway, or if the most-recent snapshot incurred an error, the current status appears in this field:

- Requested means that someone issued a snapshot create command, or that the rule's schedule just invoked a new snapshot.
- In Progress indicates that a snapshot is currently underway at the back-end filers.
- Incomplete means that the fencing timeout expired before the snapshot finished on all of its included back-end filers. This implies that snapshot consistency is enabled, causing the volume's VIP(s) to put up a *fence* against client access until the snapshots finish.
- Fence failed also implies that snapshot consistency is enabled, and means that the VIP failed to raise its fence against client access. This indicates an internal-software error; contact F5 Support if you see this.
- Policy pause failed indicates another internal-software error, and should not appear during normal operation. As above, contact F5 Support if you see this.

Guidelines: Summary View (Cont.)

– Failed indicates that the snapshot operation failed at one or more back-end filers or servers. Check the snapshot report (available in the detailed output) for more details.

Source is "Schedule" or "Manual."

– Schedule indicates that the snapshot was invoked by the snapshot rule's schedule (gbl-ns-vol-...snap), and

– Manual means that an administrator manually invoked a snapshot with the snapshot create command.

Guidelines: "Snapshot Rule Summary" in Detailed View

The detailed view of a snapshot rule appears when you specify the rule in the command. This shows two tables: "Snapshot Rule Summary" and "Snapshot Properties."

The "Snapshot Rule Summary" table summarizes the rule configuration:

Namespace Name,

Volume Name, and

Snapshot Rule Name are defined by the snapshot rule command.

Guidelines: "Snapshot Properties"

The "Snapshot Properties" Table provides details about the snapshot configuration:

Snapshots Enabled is "Yes" or "No," depending on the setting for enable (gbl-ns-vol-...snap).

Guarantee Consistency is "Enabled" or "Disabled," depending on whether or not the volume uses VIP fencing for its snapshots. The VIP fence, if enabled, blocks all client access to the volume while the filers take their coordinated snapshots. Use the snapshot consistency command to allow or disallow this fence.

Retain Count is the number of snapshots to retain for this rule. If the volume has this many snapshots when it takes a new one, it moves all of the snapshots to the next slot down and then deletes the oldest snapshot. For example, it creates a new "nightly_0" snapshot, moves the old "nightly_0" to "nightly_1," and so on, until it reaches the retain count; then it deletes any remaining snapshots. You can control this with the retain command.

Schedule is the name of the schedule for the snapshot rule, if any. Use the schedule (gbl-ns-vol-...snap) command to assign a schedule to the snapshot rule.

CIFS Directory Name only appears if the volume supports CIFS. This is the pseudo directory that well-informed CIFS clients (administrators) can use to access their snapshots. You can use the snapshot directory cifs-name command to change this name.

NFS Directory Name only appears if the volume supports NFS. This is the pseudo directory that well-informed NFS clients (administrators) can use to access their snapshots. You can use the snapshot directory nfs-name command to change this name.

Directory Display is "All Exports" (clients see the ~snapshot/.snapshot directory in any front-end share), "Volume Root Only" (clients only see the directory only in a front-end share of the volume's root directory) or "None." You can use the snapshot directory display command to change this.

Guidelines: "Snapshot Properties" (Cont.)

Hidden File Attribute only appears if the volume supports CIFS. This is "Set" if the special ~snapshot directory has its "hidden" DOS attribute raised, or "Not Set" otherwise. Use an optional argument in the snapshot directory display command to control this setting. This does not hide the ~snapshot directory from NFS clients.

Restricted Access Configured also only appears if the volume supports CIFS. This is "Yes" or "No" depending on whether or not someone used the snapshot privileged-access command. If this is "Yes," a small set of privileged CIFS clients can access the volume's snapshots. These clients are members of a Windows-Management-Authorization (*WMA*) group with permission to monitor snapshots. These privileged clients are typically administrators. This does not limit access by NFS clients in any way.

VSS Mode only appears if the volume supports CIFS. This field indicates the client-machine versions for which the volume supports the Volume Shadow-copy Service (*VSS*). VSS is an intuitive interface that Windows clients can use to access their snapshots. This is "Windows XP" (the volume supports VSS for Windows XP client machines, as well as newer machines), "Pre-Windows XP" (the volume also supports VSS for Windows-2000 clients), or "None." Use the snapshot vss-mode command to change this setting. Direct volumes do not support VSS, and NFS clients are unaffected by it.

Contents is relevant to a snapshot rule that sends its snapshots to a file-history archive. This shows the contents of this snapshot. Each of these fields shows one possible type of content with a "Yes" or "No:"

- Metadata is the volume's metadata,
- Volume Configuration is the volume-specific data from show global-config, and
- User Snapshots are the standard back-end snapshots. For a non-archive snapshot, this is the only content in the snapshot.

Archive is also relevant to a snapshot rule that sends its snapshots to a file-history archive. These field show high-level statistics for the rule's archiving operations.

Guidelines: "Snapshot Summary - *snapshot-name*"

After the rule has run at least once, additional tables appear for each snapshot. The first table, "Snapshot Summary - *snapshot-name*" summarizes the ARX's coordinated snapshot. This table contains the following fields:

Snapshot Name is the name of the coordinated snapshot.

Time Requested is the last date and time that someone issued an ARX command (such as snapshot create or snapshot verify) for this coordinated snapshot.

Time Created shows the date and time that the last filer behind the volume completed its snapshot or checkpoint operation. "In progress" appears if a snapshot is currently underway.

Last Time Verified is the last time someone ran snapshot verify against this ARX snapshot.

Guidelines: "Snapshot Summary - *snapshot-name*" (Cont.)

Request shows the currently-active request for this ARX snapshot. This is one of the following:

- Idle indicates that no command is currently running on this snapshot.
- Create means the rule is currently creating a coordinated snapshot. This indicates either a scheduled snapshot or a snapshot invoked by the snapshot create command.
- Manage shows that the rule is incorporating an existing snapshot to an ARX snapshot. This means that someone invoked the snapshot manage command.
- Delete indicates that filer snapshots are being removed from behind this rule. This request comes from the snapshot remove command.
- Remove indicates that someone invoked no snapshot rule to delete the current snapshot rule.
- Verify shows that someone invoked snapshot verify to verify the integrity of all filer snapshots behind this ARX snapshot.
- Cleanout indicates that a snapshot operation failed and the snapshot software is removing filer snapshots that are inconsistent with one another. This should be rare. If you see this, you can use the show cores command to check for core-memory files from the snapshot software. The clean-out operation generates a snapshot report that you can see with the show reports command.

Snapshot State shows the results of the most-recent snapshot operation. This is either "Complete," "Sparse," or "Incomplete." "Complete" indicates a successful snapshot on all of the volume's shares. "Sparse" means that at least one of the volume's shares is excluded from the volume snapshot, either administratively (through no manage snapshots or exclude) or because the share failed to create its snapshot. An "Incomplete" state only applies to a snapshot verify operation: this indicates that one of the included shares is now missing its back-end snapshot or checkpoint.

Snapshot Origin is "Schedule" or "Manual." This has the same possible values as the Source field in the summary view:

- Schedule indicates that the snapshot was invoked by the snapshot rule's schedule (gbl-ns-vol-...snap), and
- Manual means that an administrator manually invoked a snapshot with the snapshot create command.

Report Name identifies the report for the snapshot's most-recent action. You can use show reports *report-name* to view this report.

Guidelines: "Included Shares"

A separate table appears under the "Included Shares" heading for each of the shares in this snapshot. These are the component snapshots for the ARX snapshot described above. These snapshots are only the ones created by the ARX snapshot rule or imported into the rule with the snapshot manage command; they do not include any snapshots created independently on the filer.

Each table contains the following fields to describe the snapshot on a particular filer:

Share Name is the name of the share in the ARX volume, defined by the share command.

Filer is a header for the remaining fields.

Name is the external-filer name, defined with the external-filer command.

NFS Share is the name of the NFS export at the filer.

CIFS Share is the name of the share at the filer.

Volume is the name of the filer volume behind the above filer share.

Volume Snapshot is the name of the filer snapshot.

Guidelines: "Excluded Shares" and "Offline Shares"

If any of the volume's shares were administratively excluded from the snapshot, an "Excluded Shares" section appears to describe them. You can use no manage snapshots to exclude all shares on a filer, or exclude to exclude a particular share from the snapshot rule. These tables contain the same fields that are in the "Included Shares" tables, except the fields to describe the back-end snapshot.

If the snapshot operation failed at any of the volume's back-end share, an "Offline Shares" section appears to describe them. These share tables contain the same fields as those in the "Excluded Shares" section.

Samples

`bstnA# show snapshots`

shows a summary of all snapshots on the ARX. See Figure 30.2, below, for sample output.

`bstnA# show snapshots medarcv /lab_equipment hourlySnap`

shows details for the "hourlySnap" rule. See Figure 30.3 on page 30-19 for sample output.

Related Commands

snapshot rule

Figure 30.2 *Sample Output: show snapshots*

```
bstnA# show snapshots

Namespace: medco
Volume:    /vol

Rule                Type           Name                    Created                  Source
------------------- -------------- ----------------------- ------------------------ --------

Namespace: wwmed
Volume:    /acct

Rule                Type           Name                    Created                  Source
------------------- -------------- ----------------------- ------------------------ --------

Namespace: medarcv
Volume:    /rcrds
```

```
Rule                 Type           Name                      Created                   Source
-------------------- -------------- ------------------------- ------------------------- --------

Namespace: medarcv
Volume:    /lab_equipment

Rule                 Type           Name                      Created                   Source
-------------------- -------------- ------------------------- ------------------------- --------
hourlySnap           Snapshot       hourlySnap_0              03/07/2012 00:49:39 -0500 Manual
hourlySnap           Snapshot       hourlySnap_1              03/07/2012 00:49:08 -0500 Manual
dailySnap            Snapshot       dailySnap_0               03/07/2012 00:47:40 -0500 Manual
dailySnap            Snapshot       dailySnap_1               03/07/2012 00:47:10 -0500 Manual
mirrorSnap           Replica        mirrorSnap_0              03/07/2012 00:53:10 -0500 Manual
mirrorSnap           Replica        mirrorSnap_1              03/07/2012 00:52:22 -0500 Manual

Namespace: medarcv
Volume:    /test_results

Rule                 Type           Name                      Created                   Source
-------------------- -------------- ------------------------- ------------------------- --------

Namespace: insur
Volume:    /claims

Rule                 Type           Name                      Created                   Source
-------------------- -------------- ------------------------- ------------------------- --------
mpHourlySnap         Snapshot       mpHourlySnap_0            03/07/2012 01:07:16 -0500 Manual
mpHourlySnap         Snapshot       mpHourlySnap_1            03/07/2012 01:05:02 -0500 Manual
```

Figure 30.3 *Sample Output: show snapshots medarcv /lab_equipment hourlySnap*

```
bstnA# show snapshots medarcv /lab_equipment hourlySnap


  Snapshot Rule Summary
  ---------------------
   Namespace Name:               medarcv
   Volume Name:                  /lab_equipment
   Snapshot Rule Name:           hourlySnap


  Snapshot Properties
  -------------------
   Snapshots Enabled:            Yes
   Guarantee Consistency:        Disabled
   Retain Count:                 3
   Schedule:                     hourly
   CIFS Directory Name:          ~snapshot
   Directory Display:            All Exports
   Hidden File Attribute:        Not Set
   Restricted Access Configured: Yes
   VSS Mode:                     None
   Contents:
     Metadata:                   No
     Volume Configuration:       No
     User Snapshots:             Yes
   Archive:
     Total Archive Operations:   0
     Total Successful Operations: 0
     Total Failed Operations:    0
```

```
   Total Saved Metadata:        0 B
   Total Saved Volume Config:   0 B
   Average Copy Rate:           0 b/s

Snapshot Summary - hourlySnap_0
-----------------
 Snapshot Name:                 hourlySnap_0
 Time Requested:                03/07/2012 00:49:38 -0500
 Time Created:                  03/07/2012 00:49:39 -0500
 Last Time Verified:
 Request:                       Idle
 Snapshot State:                Complete
 Snapshot Origin:               Manual
 Report Name:                   snap_hourly_0_create_20120307004938985.rpt


Included Shares
---------------
     Share Name:                equip (user data)
     Filer:
       Name:                    nas10
       CIFS Share:              equipment
       Volume:                  vol2
       Filer Snapshot:          acopia_5_201203070549_d9bdece8-9866-11d8-91e3-f48e42637d58_vol2

     Share Name:                leased (user data)
     Filer:
       Name:                    nas10
       CIFS Share:              for_lease
       Volume:                  vol1
       Filer Snapshot:          acopia_5_201203070549_d9bdece8-9866-11d8-91e3-f48e42637d58_vol1

     Share Name:                backlots (user data)
     Filer:
       Name:                    fs2
       CIFS Share:              backlot_records
       Volume:                  E:\
       Filer Snapshot:          {3488228f-bc30-494d-a004-52ba0edcbb64}
       Time Created:            03/07/2012 00:49:40 -0500


Excluded Shares
---------------
     Share Name:                scanners
     Filer:
       Name:                    fs5
       CIFS Share:              xraysScanners
     Reason:                    Snapshots were not supported on this type of back-end filer.

Snapshot Summary - hourlySnap_1
-----------------
 Snapshot Name:                 hourlySnap_1
 Time Requested:                03/07/2012 00:49:08 -0500
 Time Created:                  03/07/2012 00:49:08 -0500
 Last Time Verified:
 Request:                       Idle
 Snapshot State:                Complete
 Snapshot Origin:               Manual
 Report Name:                   snap_hourly_0_create_20120307004908645.rpt


 Included Shares
```

```
----------------
     Share Name:               equip (user data)
     Filer:
       Name:                   nas10
       CIFS Share:             equipment
       Volume:                 vol2
       Filer Snapshot:         acopia_4_201203070549_d9bdece8-9866-11d8-91e3-f48e42637d58_vol2

     Share Name:               leased (user data)
     Filer:
       Name:                   nas10
       CIFS Share:             for_lease
       Volume:                 vol1
       Filer Snapshot:         acopia_4_201203070549_d9bdece8-9866-11d8-91e3-f48e42637d58_vol1

     Share Name:               backlots (user data)
     Filer:
       Name:                   fs2
       CIFS Share:             backlot_records
       Volume:                 E:\
       Filer Snapshot:         {bfcccd7f-1211-44f0-880a-06d849d00ba5}
       Time Created:           03/07/2012 00:49:09 -0500


Excluded Shares
----------------
     Share Name:               scanners
     Filer:
       Name:                   fs5
       CIFS Share:             xraysScanners
     Reason:                   Snapshots were not supported on this type of back-end filer.
```

snapshot clear

Purpose An ARX volume creates *snapshots*, or point-in-time copies, by coordinating snapshots and/or checkpoints at its back-end filers. Under rare circumstances, you may perform a *snapshot reconstitution* operation to connect (or re-connect) back-end snapshots to ARX snapshots. If the snapshot reconstitution causes a mismatch between the back-end snapshots and their ARX counterparts, you can use the snapshot clear command to remove the ARX snapshot(s).

Modes priv-exec

Security Role(s) storage-engineer, backup-operator, or crypto-officer

Syntax `snapshot clear` ***`namespace`*** `[`***`vol-path`*** `[`***`snap-rule`*** `[snapshot` ***`snap-instance`*** `[share` ***`share-name`***`]]]]`

namespace (1-30 characters) identifies a namespace with one or more snapshot rules.

vol-path (optional, 1-1024 characters) is one volume with snapshots. If you omit this, the command removes the configurations for all ARX snapshots in the namespace.

snap-rule (optional, 1-1024 characters) is a snapshot rule. This focuses in the snapshots in the chosen rule.

snap-instance (optional, 1-255 characters) identifies a particular ARX snapshot. Snapshots are typically named *snap-rule_n*, where *n* is 0 (zero) for the newest snapshot, 1 for the next-newest, 2 for the snapshot before 1, and so on.

share-name (optional, 1-64 characters) focuses on a single ARX share in the above volume. This removes the rule's references to snapshots on this share.

Default(s) None.

Guidelines The CLI prompts for confirmation before clearing any snapshots from the ARX configuration; enter **yes** to proceed.

This command is designed for situations where filer snapshots are incorporated into the wrong ARX snapshots. This can occur after a snapshot reconstitution with errors. (For details on snapshot reconstitution, see the Guidelines for the snapshot rule command.) You can use this command to clear the mismatched ARX snapshots, then you can edit the snapRecon.cli script and re-run it. Alternatively, you can use snapshot manage to re-incorporate one back-end snapshot at a time.

This command is not recommended for many other situations. The snapshot remove command removes the back-end snapshots behind one or more ARX snapshots, and the no snapshot rule command removes the snapshot rules.

Sample

```
bstnA# snapshot clear medarcv /lab_equipment hourlySnap snapshot
hourlySnap_2 share backlots

Confirmation of this command results in the clearing of snapshot
'hourlySnap_2' for share 'backlots' associated with snapshot rule
'hourlySnap' in namespace 'medarcv' volume '/lab_equipment'.  This
command does not remove the snapshot from the filer in the volume.

Proceed? [no] yes
```

clears the record for one snapshot on one back-end share. This applies to the "hourlySnap_2" snapshot in the ARX configuration; the command clears its reference to a snapshot on the "backlots" volume. This prevents the volume's clients from accessing any files in that back-end snapshot.

Related Commands

snapshot rule
snapshot manage
snapshot remove

snapshot consistency

Purpose

A *snapshot* is a complete, point-in-time copy of an ARX volume. The ARX volume coordinates snapshots at each of its back-end filers; you can use the snapshot consistency command to put up a *VIP fence* until the filers finish. Clients cannot access the volume's VIP(s) while the fence is up, so the snapshots are guaranteed to be consistent.

Most installations do not require strict snapshot consistency. Use no snapshot consistency to stop using the VIP fence for future snapshots.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
snapshot consistency
no snapshot consistency
```

Default(s)

```
no snapshot consistency
```

Guidelines

The VIP fence prevents any changes to the volume, which could result in an ARX snapshot that contains redundant or inconsistent files. For example, two filers may take their snapshots at slightly-different times, thus creating slight inconsistencies between the snapshots on each filer. The VIP fence prevents all inconsistencies between filer snapshots; the snapshot consistency command causes the volume to use the VIP fence.

Important

The consistency fence blocks all access to the volume's VIP or VIPs. It also stops access to the other volumes behind the VIP(s), if there are any, until the fence comes back down.

The VIP fence stops when the final back-end filer finishes its snapshot, or after a timeout. The ARX snapshot times out after 1 minute, plus 80 seconds for each back-end snapshot. If this volume uses the same snapshot schedule (gbl-ns-vol-...snap) as another volume, their snapshots are grouped and this timeout increases by 80 seconds for each additional back-end snapshot. If the timeout expires before all filers have finished their snapshots, the ARX rolls back all filer snapshots.

To create a snapshot rule (which establishes the name of a snapshot set, the number of snapshots to retain in the set, and an optional schedule), use the snapshot rule command.

Samples

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# snapshot consistency
```

enables VIP fencing in the "medarcv~/lab_equipment" volume.

```
bstnA(gbl-ns-vol[access~/G])# no snapshot consistency
```

disables VIP fencing in the "access~/G" volume. Future snapshots in this volume may have inconsistencies from one back-end snapshot to another.

Related Commands

snapshot rule

snapshot create

Purpose A *snapshot* is a full copy of a virtual volume at one point in time. A *snapshot rule* defines a name for the snapshots, the maximum number of snapshots to retain under that name, and an optional schedule. Use the snapshot create command to manually invoke a snapshot rule, creating a coordinated snapshot at the current point in time.

Modes priv-exec

Security Role(s) storage-engineer, backup-operator, or crypto-officer

Syntax `snapshot create` ***`namespace vol-path snapshot-rule`*** `[`***`snapshot-instance`***`]`

namespace (1-30 characters) identifies the namespace.

vol-path (1-1024 characters) is the name of the volume.

snapshot-rule (1-1024 characters) is the snapshot rule to invoke.

snapshot-instance (optional; 1-68 characters) specifies a name for the snapshot. If you omit this, the name defaults to "***snapshot-rule***_0" (for example, "nightly_0" for a snapshot rule named "nightly"). You cannot choose the name of an existing snapshot.

Default(s) None.

Guidelines The CLI prompts for confirmation after you enter this command. Enter **yes** to proceed.

This is not generally recommended for a scheduled snapshot rule, except to replace a scheduled snapshot that failed. If the rule has reached its maximum retain count, this deletes the oldest snapshot after creating the new one. A manual snapshot also disrupts the continuity of the snapshot schedule: for example, a Tuesday-morning snapshot may not be useful in a rule that takes weekly snapshots every Sunday night. Use show global-config namespace *namespace volume* to determine whether or not the rule has a schedule.

This command makes a snapshot-create report to show the configuration and progress of the snapshot operation. The name of the report appears on the command line after you confirm the command. See Figure 30.1 on page 30-9 for a sample snapshot-create report.

This command creates filer snapshots asynchronously, allowing you to continue entering CLI commands while the operation proceeds. You can use the tail reports *report-name* follow command to follow the progress of the snapshot creation. You can also use wait-for snapshot create to wait for all snapshots to finish on the back-end filers; this is especially useful in CLI scripts.

Some filers may remove their snapshots to conserve space. To verify that all filer snapshots are still available behind a snapshot rule, use the snapshot verify command. To remove filer snapshots from behind a snapshot rule yourself, use snapshot remove.

Sample

```
prtlndA# snapshot create nemed /vol44 COB
              Snapshot rule 'COB' in namespace 'nemed' volume '/vol44'
has a schedule associated with it.  Running this rule may cause
existing snapshots to be removed if the retain count is exceeded by
this request.
Proceed? [yes/no]  yes

  Starting snapshot operation in volume /vol44, report name:
snap_cob_0_create_200707180325.rpt
```

invokes the “COB” rule and creates a snapshot of the “nemed~/vol44” volume.

Related Commands

snapshot rule
schedule (gbl-ns-vol-...snap)
retain
wait-for snapshot create
snapshot verify
snapshot remove

snapshot directory cifs-name

Purpose

An ARX volume can periodically take *snapshots*, or point-in-time copies, of all of its contents. CIFS administrators can access the volume's snapshots through hidden pseudo directories, called "~snapshot" by default. This directory exists under every directory in the volume. Use the snapshot directory cifs-name command to rename this directory for CIFS clients.

Use no snapshot directory cifs-name to return to the default name for the container directory.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
snapshot directory cifs-name container-directory
no snapshot directory cifs-name
```

container-directory (1-32 characters) is the CIFS name you choose for the pseudo directory that contains snapshots. This directory exists under every directory in the volume. Do not use any characters that are illegal in CIFS names: the illegal characters are any control character, /, \, :, *, >, <, ", |, or ?. Also, avoid any name that a client might use for a standard file or directory.

Default(s)

~snapshot

Guidelines

The snapshot directory exists under every directory in the volume, but does not appear by default. For example, if drive "M:" maps to an ARX volume with snapshots and has two subdirectories, lab1 and lab2, all of the following snapshot directories exist: M:\~snapshot, M:\lab1\~snapshot, and M:\lab2\~snapshot. The ~snapshot directory does not appear when you type **dir** in M:\lab1, or in the Windows Explorer view of that directory, but a well-informed client can access it by name (for example, by entering **cd M:\lab1\~snapshot** in a DOS shell). You can use this command to change the name of the ~snapshot directory in all directories of the volume.

This is the name of the directory as seen by CIFS clients, not NFS clients. To control the directory name seen by NFS clients, use the snapshot directory nfs-name command.

To create a snapshot rule (which establishes the name of a snapshot set, the number of snapshots to retain in the set, and an optional schedule), use the snapshot rule command.

The snapshot privileged-access command makes snapshots accessible only to a privileged group of CIFS clients that you can create on the ARX. This applies to the snapshot directory as well as access through Properties -> Previous Versions in Windows Explorer. If this is set, clients without privileges can not enter or read any snapshot directories in the volume. If this is disabled, all clients can access the snapshot directories as long as they are aware of the directory name.

You can use the snapshot directory display command to control the volume exports that display this directory. You can choose to display the special directory only in exports of the volume-root directory, in exports of any directory in the volume, or in no exports of the volume.

Samples

```
bstnA(gbl-ns-vol[access~/G])# snapshot directory cifs-name ~ckpt
```

creates a sub directory named "~ckpt" under every directory in the "access~/G" volume. CIFS clients can access any directory's snapshots from the directory's "~ckpt" subdirectory.

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# no snapshot directory
cifs-name
```

reverts to the default snapshot-directory name (as seen through CIFS) in the "medarcv~/lab_equipment" volume.

Related Commands

snapshot rule
snapshot directory nfs-name
snapshot privileged-access
snapshot directory display

snapshot directory display

Purpose

A well-informed client can access a volume's *snapshots*, or point-in-time copies, through an undisplayed pseudo directory. The client must know the name of the pseudo directory (typically "~snapshot" for CIFS and ".snapshot" for NFS) to access it; the name does not appear in directory listings. You can use the snapshot directory display command to display the pseudo directory in all of the volume's front-end exports. You can also use this command to limit the display to exports of the volume's root directory, so that exports of lower-level directories do not display "~snapshot."

Use no snapshot directory display to stop displaying the snapshot directory in any export of the volume. With this setting, clients need to know the name of the directory to **cd** to it.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
snapshot directory display {all-exports | volume-root-only}
[hidden]

no snapshot directory display
```

all-exports | volume-root-only is a required choice. This selects the type of front-end export that displays the ~snapshot/.snapshot directory:

all-exports means to display the special directory in the root of any front-end export. For example, the /users volume would display the snapshot directory in an export of /users/jsmith or in an export of the volume root (/users).

volume-root-only displays the ~snapshot/.snapshot directory only in front-end exports of the volume root (/users in the previous example), not in exports of a lower-level directory (such as /users/tjefferson). This is useful for a site where most clients use exports below the volume root, and only administrators use an export at the root.

hidden raises the DOS "hidden" attribute for the ~snapshot directory. This limits the display to CIFS clients that are configured to override the "hidden" attribute. It has no effect on NFS clients.

Default(s)

```
no snapshot directory display
```

Guidelines

For CIFS clients, this command is often used in conjunction with the no snapshot vss-mode command, which stops snapshot access through the VSS interface. *VSS* (or Volume Shadowing Service) allows Windows clients to click on a file or directory in Windows Explorer, pull up the Properties pop-up, and access snapshots for the file or directory through the Previous Versions tab. This intuitive interface is designed for the vast majority of CIFS clients, whereas the ~snapshot directory is designed for well-informed CIFS administrators.

Direct volumes and NFS-only volumes do not support VSS, so the ~snapshot/.snapshot directory is the only means of snapshot access for direct-volume clients and NFS clients.

Guidelines (Cont.)

Administrative clients can access the pseudo directory whether or not you display it with this command; for example, a Windows administrator can use **cd M:\users\~snapshot** in a DOS shell to enter the ~snapshot directory under \users, or a Unix administrator can use **cd /home/.snapshot**. You can use the snapshot directory cifs-name command to change the directory name seen by CIFS clients, and you can use snapshot directory nfs-name to change the NFS name. If the volume uses snapshot privileged-access, only privileged CIFS clients can access the volume's snapshots, either through VSS or through the ~snapshot directory. (The snapshot privileged-access command has no effect on NFS clients.)

This change is visible to CIFS clients the next time they issue the **dir** command, or on the next refresh of their graphical view. For NFS clients, the directory appears on the next **ls** command. That is, the snapshot directory seems either appear or vanish from the volume's front-end exports.

To create a snapshot rule (which establishes the name of a snapshot set, the number of snapshots to retain in the set, and an optional schedule), use the snapshot rule command.

Samples

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# snapshot directory display
all-exports
```

shows the snapshot-container directory in all exports of the "medarcv~/lab_equipment" volume.

```
bstnA(gbl-ns-vol[access~/G])# no snapshot directory display
```

stops the display of the snapshot-container directory in exports of "access~/G."

Related Commands

snapshot vss-mode
snapshot privileged-access
snapshot directory cifs-name
snapshot directory nfs-name
snapshot rule
export (gbl-cifs)

snapshot directory nfs-name

Purpose

An ARX volume can periodically take *snapshots*, or point-in-time copies, of all of its contents. NFS clients can access the volume's snapshots through hidden pseudo directories, called ".snapshot" by default. This directory exists under every directory in the volume. Use the snapshot directory nfs-name command to rename this directory for NFS clients.

Use no snapshot directory nfs-name to return to the default name for the container directory.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
snapshot directory nfs-name container-directory
no snapshot directory nfs-name
```

container-directory (1-32 characters) is the NFS name you choose for the pseudo directory that contains snapshots. This directory exists under every directory in the volume. Avoid any name that a client might use for a standard file or directory.

Default(s)

.snapshot

Guidelines

The snapshot directory exists under every directory in the volume, but does not appear by default. For example, if the directory "/mnt/arx" is mounted to an ARX volume with snapshots and has two subdirectories, lab1 and lab2, all of the following snapshot directories exist: /mnt/.snapshot, /mnt/lab1/.snapshot, and /mnt/lab2./snapshot. The .snapshot directory does not appear when you type **ls** in /mnt/lab1, but a well-informed client can access it by name (for example, by entering **cd /mnt/lab1/.snapshot** in a Unix shell). You can use this command to change the name of the .snapshot directory in all directories of the volume.

This is the name of the directory as seen by NFS clients, not CIFS clients. To control the directory name seen by CIFS clients, use the snapshot directory cifs-name command.

To create a snapshot rule (which establishes the name of a snapshot set, the number of snapshots to retain in the set, and an optional schedule), use the snapshot rule command.

You can use the snapshot directory display command to control the volume exports that display this directory. You can choose to display the special directory only in exports of the volume-root directory, in exports of any directory in the volume, or in no exports of the volume.

Sample

```
bstnA(gbl-ns-vol[access~/G])# snapshot directory nfs-name .ckpt
```

creates a directory name of ".ckpt" under every directory in the "access~/G" volume. NFS clients can access their snapshots from any instance of the ".ckpt" directory.

Related Commands

snapshot rule
snapshot directory cifs-name
snapshot privileged-access
snapshot directory display

snapshot manage

Purpose Use the snapshot manage command to incorporate a filer snapshot into an ARX snapshot.

Modes priv-exec

Security Role(s) backup-operator, storage-engineer, or crypto-officer

Syntax **`snapshot manage`** ***`namespace vol-path share rule filer-snap`*** **`created-on`** ***`date-time`*** **`[report-prefix`** ***`prefix`*** **`] [verbose]`**

namespace (1-30 characters) identifies the ARX namespace.

vol-path (1-1024 characters) is the name of the ARX volume.

share (1-64 characters) is the name of the ARX share.

rule (1-1024 characters) is the snapshot rule to receive the filer snapshot.

filer-snap (1-255 characters) identifies the snapshot to incorporate from the back-end filer.

date-time is the date and time that the ARX snapshot was created, in *mm*/*dd*/*yyyy*:*HH*:*MM*:*SS* format. This identifies the specific ARX snapshot to include the ***filer-snap***. You can find this in the output of show snapshots. If you enter a time when no snapshot was previously taken, this creates a new ARX snapshot.

prefix (optional, 1-64 characters) is the prefix for a report file. The CLI logs its output to a report file named as follows:

prefix-namespace-volume-rule-yyyymmddHHMM.rpt

where *prefix*, *namespace*, *volume*, and *rule* are all chosen in this command, and *yyyymmddHHMM* is the current date and time.

verbose (optional) forces a report even for a successful operation; by default, the command only produces a report if there is an error.

Default(s) ***prefix*** - "snapshot-manage"

Guidelines This command incorporates the snapshot synchronously, before returning to the next CLI prompt. If the snapshot fails, or if you entered the **verbose** keyword, a report name appears. You can use show reports *report-name* to view the report's contents.

You can use the snapshot clear command to clear one or more snapshots from the ARX configuration. The snapshot clear command disconnects back-end snapshots from the ARX configuration, but does not remove any snapshot rules or back-end snapshots.

Use the snapshot rule command to create a snapshot rule for the ARX volume. The snapshot rule is required for incorporating an existing snapshot; it defines the schedule for creating future snapshots (if any) and the number of snapshots to retain.

Sample `prtlndA# snapshot manage nemed /vol7 vol2exp daily daily.0 created-on 2008/06/05:01:04:54`

incorporates an existing filer snapshot, "daily.0," into the "nemed~/vol7" volume.

Related Commands

snapshot rule
show snapshots
snapshot clear

snapshot privileged-access

Purpose

An ARX volume typically allows access to its *snapshots*, or point-in-time copies, to all CIFS clients. Some installations choose to limit the visibility of snapshots, so that only designated CIFS administrators can access them. You can manage this by identifying the administrators in a Windows-Management-Authorization (*WMA*) group, applying that group to the current namespace, and using the snapshot privileged-access command to enforce the WMA group for the volume's snapshots.

Use no snapshot privileged-access to open snapshot access to all CIFS clients, not just administrators in WMA groups.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
snapshot privileged-access
no snapshot privileged-access
```

Default(s)

```
no snapshot privileged-access
```

Guidelines

This command has no effect on NFS clients.

You can use this command to select a privileged set of CIFS clients with access to snapshots. CIFS clients outside the namespace's WMA group(s) cannot access the volume's snapshots by any method.

The no form of this command opens up snapshot access to all CIFS clients.

Use the windows-mgmt-auth command to create a WMA group, permit snapshot monitor (permit (gbl-mgmt-auth)) to authorize the group for snapshot access, and windows-mgmt-auth (gbl-ns) to assign the WMA group to the current namespace. If no such WMA groups are assigned to the namespace, this command disables snapshot access for all clients.

To create a snapshot rule (which establishes the name of a snapshot set, the number of snapshots to retain in the set, and an optional schedule), use the snapshot rule command.

Guidelines: Methods for Accessing Snapshots

CIFS clients have two methods for accessing snapshots. The most intuitive method is through the Windows Explorer interface: select a desired file or directory, pull up Properties, and click the Previous Versions tab. This is called the Microsoft Volume-Shadowing Service (*VSS*) interface. The second method is designed primarily for CIFS administrators: access snapshots through a pseudo directory named ~snapshot, which is not displayed by default. This is the only interface supported by direct volumes, and it is the only way NFS clients can access their snapshots. (For NFS clients, the directory name is ".snapshot".)

Guidelines: Alternatives for Controlling Snapshot Access

If snapshot privileged-access is enabled, CIFS clients within the right WMA group(s) can access snapshots. If not, all CIFS clients can access snapshots. In either case, clients with snapshot access are subject to restrictions from these commands:

- You can use no snapshot vss-mode to stop all snapshot access through VSS. This only applies to CIFS clients.
- You can snapshot directory display to control the visibility of the ~snapshot/.snapshot directory. By default, the volume does not display the directory; well-informed CIFS or NFS clients must enter its exact path to access its snapshots.
- You can use the snapshot directory cifs-name command to change the name of the snapshot-container directory, as seen by CIFS clients. You can also use snapshot directory nfs-name to change the name seen by NFS clients.

If this command, snapshot privileged-access, is enabled, no CIFS client *outside* the namespace's WMA groups can access snapshots at all. As mentioned above, this command has no effect on NFS clients.

Samples

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# snapshot privileged-access
```

makes snapshots in "medarcv~/lab_equipment" accessible to only a small group of CIFS clients: the properly-authorized members of the namespace's WMA group(s).

```
bstnA(gbl-ns-vol[access~/G])# no snapshot privileged-access
```

ignores WMA groups in the "access~/G" volume.

Related Commands

windows-mgmt-auth -> permit (gbl-mgmt-auth)
windows-mgmt-auth (gbl-ns)
snapshot rule
snapshot directory display
snapshot directory cifs-name
snapshot directory nfs-name
export (gbl-cifs)

snapshot remove

Purpose An ARX volume creates *snapshots*, or point-in-time copies, by coordinating snapshots and/or checkpoints at its back-end filers. A snapshot rule establishes the name of the ARX snapshots, the number of snapshots to retain under that name, and an optional schedule for taking snapshots. You can use the snapshot remove command to remove the filer snapshots behind a snapshot rule, without removing the rule itself.

Modes priv-exec

Security Role(s) storage-engineer, backup-operator, or crypto-officer

Syntax `snapshot remove` ***`namespace vol-path snapshot-rule`*** `[`***`snapshot-instance`***`]`

namespace (1-30 characters) identifies the namespace.

vol-path (1-1024 characters) is the name of the volume.

snapshot-rule (1-1024 characters) is the snapshot rule.

snapshot-instance (optional, 1-68) identifies a particular snapshot created by the snapshot rule.

Default(s) None.

Guidelines The CLI prompts for confirmation before removing any snapshots from any filers. Enter **yes** to continue.

This command produces a separate removal report for each ARX snapshot. Each report shows the configuration of the snapshot rule, a summary status for the each ARX snapshot, and removal details about the filer snapshots for each ARX snapshot. See Figure 30.5 on page 30-54 for a sample removal report.

To remove the rule configuration itself, use the no snapshot rule command. If you want to remove the snapshot rule without leaving any of its filer snapshots behind, use this command first.

This command removes filer snapshots asynchronously, allowing you to continue entering CLI commands while the operation proceeds. You can use the tail reports *report-name* follow command to follow the progress of each snapshot removal. You can also use wait-for snapshot remove to wait for all snapshot removals to finish; this is especially useful in CLI scripts.

You can use the snapshot create command to manually create a new snapshot behind an existing snapshot rule.

Sample

```
bstnA# snapshot remove medarcv /lab_equipment hourlySnap

Confirmation of this command results in the removal all snapshots
associated with snapshot rule 'hourlySnap' in namespace 'medarcv'
volume '/lab_equipment'.  The snapshot rule is not deleted.

 Proceed? [yes/no] yes

  Starting snapshot operation in volume /lab_equipment, report name:
snap_hourly_0_remove_20120229040753550.rpt


  Starting snapshot operation in volume /lab_equipment, report name:
snap_hourly_2_remove_20120229040753550.rpt


  Starting snapshot operation in volume /lab_equipment, report name:
snap_hourly_1_remove_20120229040753550.rpt

bstnA#
```

removes all filer snapshots behind the "hourlySnap" rule. For a sample report, see Figure 30.4 below.

Related Commands

snapshot rule
wait-for snapshot remove
snapshot create

Figure 30.4 *Sample Report: snap_hourly_1_remove_....rpt*

```
bstnA# show reports snap_hourly_1_remove_20120229040753550.rpt

  Snapshot Rule Summary
  ---------------------
   Namespace Name:               medarcv
   Volume Name:                  /lab_equipment
   Snapshot Rule Name:           hourlySnap


  Snapshot Properties
  -------------------
   Snapshots Enabled:            Yes
   Guarantee Consistency:        Disabled
   Retain Count:                 3
   Schedule:                     hourly
   CIFS Directory Name:          ~snapshot
   Directory Display:            All Exports
   Hidden File Attribute:        Not Set
   Restricted Access Configured: Yes
   VSS Mode:                     None
   Contents:
     Metadata:                   No
     Volume Configuration:       No
     User Snapshots:             Yes
   Archive:
     Total Archive Operations:   0
     Total Successful Operations: 0
     Total Failed Operations:    0
     Total Saved Metadata:       0 B
     Total Saved Volume Config:  0 B
     Average Copy Rate:          0 b/s
```

```
Snapshot Summary - hourlySnap_1
-----------------
 Snapshot Name:                hourlySnap_1
 Snapshot Operation:           Delete
 Result:                       Success
 Time Requested:               02/29/2012 04:01:55 -0500
 Time Created:                 02/29/2012 02:39:00 -0500
 Last Time Verified:           02/29/2012 04:01:57 -0500
 Request:                      Delete
 Snapshot State:
 Snapshot Origin:              Schedule
 Report Name:                  snap_hourly_1_remove_20120229040753550.rpt


Included Shares
---------------
     Share Name:               equip (user data)
     Filer:
       Name:                   nas10
       CIFS Share:             equipment
       Volume:                 vol2
       Filer Snapshot:         acopia_14_201202290739_d9bdece8-9866-11d8-91e3-f48e42637d58_vol2

     Share Name:               leased (user data)
     Filer:
       Name:                   nas10
       CIFS Share:             for_lease
       Volume:                 vol1
       Filer Snapshot:         acopia_14_201202290739_d9bdece8-9866-11d8-91e3-f48e42637d58_vol1

     Share Name:               backlots (user data)
     Filer:
       Name:                   fs2
       CIFS Share:             backlot_records
       Volume:                 E:\
       Filer Snapshot:         {0188047d-fd0d-4b95-b769-0fbceae4ad6b}
       Time Created:           02/29/2012 02:39:12 -0500


Excluded Shares
---------------
     Share Name:               scanners
     Filer:
       Name:                   fs5
       CIFS Share:             xraysScanners
     Reason:                   Snapshots were not supported on this type of back-end filer.

     Share Name:               equipSnap
     Filer:
       Name:                   nas11
       CIFS Share:             equipBkup
     Reason:                   Replica-snap share excluded from managed share snapshot.

     Share Name:               leasedSnap
     Filer:
       Name:                   nas11
       CIFS Share:             leasedBkup
     Reason:                   Replica-snap share excluded from managed share snapshot.
```

snapshot replica-snap-rule

Purpose

A *snapshot* is a copy of all files and directories in an ARX volume at a particular point in time. A standard snapshot rule determines the schedule for taking snapshots and the number snapshots to retain; these snapshots reside on the same back-end shares that store the volume's files. A *replica-snap rule* takes snapshots only on special replica-snap shares behind the volume. A filer-replication program (such as SnapMirror or RoboCopy) replicates a primary share to the replica-snap share regularly and independently. The server with the replica-snap share is presumed to have enough disk space for a large number of snapshots. You can use a replica-snap rule to take regular snapshots on the replica-snap share(s). This eases the storage burden on the filer behind the primary share(s). Use this command, snapshot replica-snap-rule, to start configuring a replica-snap rule.

Use the no form of the command to remove the replica-snap rule without removing any of the snapshots behind it.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
snapshot replica-snap-rule name
no snapshot replica-snap-rule name
```

name (1-1024 characters) is a name you choose for the replica-snap rule.

Default(s)

None

Guidelines: Filer Preparation

Before you begin, you must prepare the filers behind the volume. The volume must have one or more replica-snap shares, where each replica-snap share is backed by NetApp filers, EMC Celerra servers, EMC Data Domain systems, and/or Windows servers that support snapshots. In the case of Windows servers, WinRM must also be installed so that the ARX can invoke snapshots through its management API.

Each replica-snap share holds an updated duplicate of the files and directories in another share. The source share and the replica-snap share are both in the same managed volume. The managed volume presents the source share's files and directories to its clients, along with any snapshots in the replica-snap share. It may also present the source share's snapshots if you configure a standard snapshot rule for the volume.

The ARX volume creates a coordinated snapshot by issuing CLI commands to the back-end filer(s) that host its replica-snap shares. The ARX therefore needs information and credentials for accessing each filer's CLI. From gbl-filer mode, use the filer-type command to identify the filer vendor, use proxy-user (gbl-filer) to identify a proxy user with proper management-login credentials, and use manage snapshots to declare that the filer supports snapshots. You can use the ip address ... management command to designate the management-IP address at that station (by default, the ARX logs into the CLI through an external filer's primary-IP address, set with the simplest syntax for the ip address command).

Guidelines: Some Filer-Replication Applications Overwrite ARX Snapshots

Some filer-replication applications are volume-level, and copy the entire contents of the back-end share to the replica-snap share, including the filer snapshots. This would overwrite any snapshots that the ARX takes at the replica-snap share, and defeat the purpose of this rule. Do not use volume-level replication to copy the source share's contents to the replica-snap share.

Guidelines

An enabled replica-snap rule is the basis for managing replica snapshots on the ARX. You can apply a schedule to the rule so that it takes regular snapshots, or you can invoke the rule manually with the snapshot create command. This type of rule only creates snapshots on replica-snap shares, and ignores standard shares; use the snapshot rule command to create snapshots on standard shares, too.

When you create a new replica-snap rule, the CLI prompts for confirmation. Enter **yes** to create the rule. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.)

This command places you in gbl-ns-vol-replica-snap mode, where you enable the rule and where you have some options that you can apply to it. By default, a replica-snap rule retains three snapshots; whenever it successfully creates a new snapshot, it deletes the oldest snapshot so that there are never more than three. You can use the optional retain command to change the number of retained snapshots; a high retention count is typical for a replica-snap rule, since it only takes disk space on the replica-snap shares. You can set a regular schedule for the replica snapshots with the schedule (gbl-ns-vol-...snap) command; we recommend a different schedule from the one used for the volume's standard snapshots. To enable report-generation for each snapshot, use the report (gbl-ns-vol-...snap) command. You must use the enable (gbl-ns-vol-...snap) command enable this rule to take any snapshots at all, even manually.

A replica-snap rule has no effect in a direct volume, which cannot contain any replica-snap shares.

You can use snapshot manage to incorporate existing filer snapshots from the replica-snap shares into the replica-snap rule.

Each ARX snapshot has one or more component snapshots on its back-end filers. You can use the snapshot verify command to verify that all of the component snapshots still exist behind a replica-snap rule.

To see the current state of snapshots, including replica snapshots, use the show snapshots command.

Guidelines: Snapshot Presentation to Clients

By default, CIFS clients can access their snapshots with Windows Explorer. They select a file or directory, pull up its Properties, and find a list of snapshots for the file or directory in the Previous Versions tab. CIFS clients can use this interface to find and restore previous versions of their files and directories. Microsoft calls this the Volume Shadowing Service, or *VSS*, for Shared Folders. Managed volumes support VSS for Shared Folders, but direct volumes do not and NFS clients do not use VSS.

◆ Note

*If CIFS clients were connected to the volume before you create your first snapshot-related rule (*snapshot rule*,* notification rule*, or this rule), the clients must shut down and restart all instances of Windows Explorer before they can see the* Previous Versions *tab. Windows Explorer only checks for snapshot support when it first connects to the share.*

The snapshots from a replica-snap rule appear interleaved with any snapshots from a standard snapshot rule.

Some filer-replication applications may drop file attributes, such as Windows Security Descriptors. If this occurs, the replica snapshots do not include these attributes. The managed-volume software cannot detect or correct this issue.

For installations where only administrators are allowed to access snapshots, you have commands for controlling VSS access:

- You can use snapshot privileged-access to select a limited number of CIFS clients (the site administrators) who can access snapshots.
- You can also (or alternatively) stop VSS access altogether with no snapshot vss-mode. Administrators can continue to access snapshots through a special pseudo directory, ~snapshot, that is not displayed by default.

Both of the above options affect CIFS clients only; they have no effect on NFS clients.

You can use snapshot directory cifs-name to change the name of the ~snapshot directory. This is the name seen by CIFS clients only; you can use snapshot directory nfs-name to provide a different name for NFS clients. You can also control the display of the directory based on the export: use the snapshot directory display volume-root-only command to display this directory only in exports from the root of the volume, not in exports of the volume's subdirectories. These commands affect NFS clients as well as CIFS clients.

Guidelines: Data Domain Servers and Presentation Through NFS

EMC Data Domain servers provide a unique presentation to NFS clients; they are only accessible from volume root's .snapshot directory. The .snapshot directory under each lower directory may or may not include the Data Domain snapshots.

This limitation does not apply to CIFS clients, who can access the Data-Domain snapshots in every ~snapshot directory.

Guidelines: Filer-Snapshot Names

Each EMC or NetApp snapshot has a unique name to distinguish it as a backing snapshot for a particular ARX. The snapshot names use the following format

`acopia_`*`id_time-stamp_uuid_filer-volume`*

where

- *id* is an integer that the snapshot software uses internally.
- *time-stamp* is in *yyyymmddHHMM* format. This is the time that the snapshot was created.
- *uuid* is the Universally-Unique ID that identifies the ARX. In a redundant pair, this is the UUID for the peer that was originally configured as senior/active, no matter which peer is currently active. You can find an ARX's UUID in the output of show running-config.
- *filer-volume* is the name of the filer's volume.

Windows servers use a UUID in curly braces ({}) to identify each snapshot. This does not change for ARX-generated snapshots.

Guidelines: Configuration Changes in a Scheduled Rule

If you make a configuration change in a replica-snap rule that is running with a schedule, the configuration change is ineffective until the next time the rule runs.

Guidelines: Removing the Replica-Snap Rule

The no form of the command removes the replica-snap rule without removing any snapshots from the back-end filers. To remove the snapshots from the filers, use the snapshot remove command *before* you remove the rule. This is an efficient method for cleaning all of the supporting snapshots behind the rule. If supporting snapshots remain when you invoke no snapshot replica-snap-rule, the CLI lists all remaining snapshots when it prompts for confirmation.

Guidelines: Snapshot Reconstitution

An accidental replica-snap-rule removal would separate the back-end snapshots from the ARX configuration, requiring a *reconstitution* of the coordinated snapshot. There are other situations where you may require snapshot reconstitution, too, such as a site-to-site failover: if a filer mechanism duplicates all of the filer snapshots from Site A over to Site B, and each site is managed by its own ARX pair, the snapshots at Site B need to be reconstituted in Site B's ARX pair (see cluster-name and activate configs for details about site-to-site failovers). The snapshot-reconstitution process requires some preparation when you start adding snapshot and replica-snap rules to the configuration. These guidelines show the high-level process for preparing your snapshots, along with the process for reconstituting snapshots in the event of an issue.

Guidelines: Preparing for Snapshot Reconstitution

After you add your first snapshot or replica-snap rule, choose an external repository to hold all of the latest snapshot reports. Choose a filer with a large amount of free space, or monitor the filer's space frequently and remove outdated snapshot reports. Every snapshot report has a timestamp in its name, so it is unique.

The filer should be able to run Perl scripts, and requires the XML::Simple module. You can download this Perl module from CPAN (http://search.cpan.org) if your system does not already have it.

Use the at command together with copy ftp, copy scp, copy {nfs|cifs}, or copy smtp to regularly copy snapshot reports from the ARX to your chosen filer. Use the common string for all of your snapshot reports, together with an "*" or other wildcard. The report repository should always hold the latest snapshot reports, so that they have the latest back-end snapshot names. For example, this command copies all reports starting with "snap" to an external IP each morning:

```
bstnA(cfg)# at 01:19:18 every 1 day do "copy reports snap*
ftp://ftpuser:ftpuser@172.16.100.183//var/arxSnapRpts/ format xml"

  The scheduled execution time for CLI command is: 3/30/09 1:19 AM.
```

You also need a copy of the "snap-recon.pl" script from the "software" directory on the ARX:

```
bstnA# copy software snap-recon.pl
ftp://ftpuser:ftpuser@172.16.100.183//var/arxSnapRpts/
```

Guidelines: Reconstituting ARX Snapshots

You can perform snapshot reconstitution if you have the snap-recon.pl script and the latest set of snapshot reports on a host that supports Perl. Start by running the snap-recon.pl script on that host. This produces a CLI script with a sequence of snapshot manage commands. By default, the output script is named, "snapRecon.cli." This script has several options; execute the command without any options to get a complete list. You must use the "**--report-dir *directory***" option to specify the *directory* that holds the reports. For example, this command sequence lists the files on "client2:/var/arxSnapRpts," runs snap-recon.pl on the reports in the current directory (.), and then shows the new file in the directory:

```
juser@client2:/var/arxSnapRpts$ ls
snap_daily_0_create_20090330010418883.xml  snap_hourly_0_create_20090330010652616.xml
snap_daily_0_create_20090330010454303.xml  snap_hourly_0_create_20090330010723024.xml
snap_daily_0_create_20090330010524663.xml  snap-recon.pl
snap_daily_2_remove_20090330010612220.xml
juser@client2:/var/arxSnapRpts$ ./snap-recon.pl --report-dir .
juser@client2:/var/arxSnapRpts$ ls
snap_daily_0_create_20090330010418883.xml  snap_hourly_0_create_20090330010652616.xml
snap_daily_0_create_20090330010454303.xml  snap_hourly_0_create_20090330010723024.xml
snap_daily_0_create_20090330010524663.xml  snapRecon.cli
snap_daily_2_remove_20090330010612220.xml  snap-recon.pl
juser@client2:/var/arxSnapRpts$
```

Guidelines: Reconstituting ARX Snapshots (Cont.)

You have the option to edit the CLI script before you run it on the ARX. For example, the ARX volumes at a disaster-recovery site may have different names than those at the primary site. The snapRecon.cli script has several options for substituting these names; you can use these to rebuild the CLI script, or you can manually edit the CLI script.

Once the CLI script is ready, you can download it to the ARX and run it. Use copy ftp, copy {nfs|cifs}, or copy scp for the download, and use run to run it.

A large, complex CLI script may contain errors. If you discover any back-end snapshots that are mismatched with their ARX counterparts, you can use the snapshot clear command to remove the ARX snapshot (not the rule) from the configuration. Then you can edit and re-run the script, or you can use snapshot manage to manually incorporate the back-end snapshots into the correct ARX snapshots.

Samples

```
bstnA(gbl-ns-vol[access~/G])# snapshot rule nightly
This will create a new policy object.

Create object 'nightly'? [yes/no] yes
bstnA(gbl-ns-vol-snap[access~/G~nightly])# ...
```

instantiates a new snapshot rule, “nightly,” in the current namespace volume.

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# no snapshot rule hourlySnap

The following snapshots are being managed by the rule 'hourlySnap':

hourlySnap_0(acopia_27_200707170500_501f705c-8735-11d8-8936-a58a9e4556
df_datavol2)

hourlySnap_2(acopia_25_200707170300_501f705c-8735-11d8-8936-a58a9e4556
df_datavol2)

hourlySnap_1(acopia_26_200707170400_501f705c-8735-11d8-8936-a58a9e4556
df_datavol2)

If this command is confirmed, the rule is deleted and all associated
data related to the aforementioned snapshots will be removed from the
switch.  The snapshots on the filers will not be removed.

Proceed? [yes/no] yes
bstnA(gbl-ns-vol-snap[medarcv~/lab_equipment])# ...
```

removes a snapshot rule with some existing snapshots remaining on one of its back-end filers.

Related Commands

proxy-user
filer-type
proxy-user (gbl-filer)
ip address `... management`
manage snapshots

retain
exclude
schedule (gbl-ns-vol-...snap)
report (gbl-ns-vol-...snap)
enable (gbl-ns-vol-...snap)
show snapshots

snapshot vss-mode
snapshot directory display
snapshot directory cifs-name
snapshot directory nfs-name

snapshot manage
snapshot create
snapshot verify
snapshot remove

snapshot clear

snapshot rule

Purpose

A *snapshot* is a copy of all files and directories in an ARX volume at a particular point in time. A snapshot rule determines the schedule for regular snapshots (if any) and the number snapshots to retain. Use this command to start configuring a snapshot rule.

Use the no form of the command to remove the snapshot rule without removing any of the snapshots behind it.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
snapshot rule name
no snapshot rule name
```

name (1-1024 characters) is a name you choose for the snapshot rule.

Default(s)

None

Guidelines: Filer Preparation

Before you begin, you must prepare the filers behind the volume. The volume must be backed by NetApp filers, EMC Celerra servers, EMC Data Domain systems, and/or Windows servers that support snapshots. In the case of Windows servers, WinRM must also be installed so that the ARX can invoke snapshots through its management API.

The ARX volume creates a coordinated snapshot by issuing CLI commands to each of its back-end filers. The ARX therefore needs information and credentials for accessing each filer's CLI. From gbl-filer mode, use the filer-type command to identify the filer vendor, use proxy-user (gbl-filer) to identify a proxy user with proper management-login credentials, and use manage snapshots to declare that the filer supports snapshots. You can use the ip address ... management command to designate the management-IP address at that station (by default, the ARX logs into the CLI through an external filer's primary-IP address, set with the simplest syntax for the ip address command).

Guidelines

An enabled snapshot rule is the basis for managing snapshots on the ARX. You can apply a schedule to the rule so that it takes regular snapshots, or you can invoke the rule manually with the snapshot create command.

When you create a new snapshot rule, the CLI prompts for confirmation. Enter **yes** to create the rule. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.)

This command places you in gbl-ns-vol-snap mode, where you enable the rule and where you have some options that you can apply to it. By default, a snapshot rule retains three snapshots; whenever it successfully creates a new snapshot, it deletes the oldest snapshot so that there are never more than three. You can use the optional retain command to change the number of retained snapshots. You can set a regular schedule for the snapshots with the schedule (gbl-ns-vol-...snap) command. To enable report-generation for each snapshot, use the report (gbl-ns-vol-...snap) command. You must use the enable (gbl-ns-vol-...snap) command enable this rule to take any snapshots at all, even manually. To exclude one of the volume's shares from the coordinated snapshot (under the advisement of F5 Support), you can use the exclude command.

You can configure a snapshot rule in a managed volume or a direct volume. The snapshot rule does not require any metadata, so a direct volume can support it.

You can use snapshot manage to incorporate existing filer snapshots into the snapshot rule.

Each ARX snapshot has one or more component snapshots on its back-end filers. You can use the snapshot verify command to verify that all of the component snapshots still exist behind a snapshot rule.

To see the current state of snapshots, use the show snapshots command.

Guidelines: Coexistence with Replica-Snap Rules

If your back-end file servers replicate one or more of your tier-1 shares on a cheaper server, where you create and store a growing collection of snapshots, you can declare that server's shares as special replica-snap shares. A snapshot replica-snap-rule makes coordinated snapshots of all replica-snap shares in the current volume, and a snapshot rule (described here) makes coordinated snapshots of all the other shares in the current volume. ARX clients see the snapshots from a replica-snap rule interleaved with the snapshots from a standard snapshot rule.

Guidelines: VIP Fencing

During a snapshot-create operation, clients can access the volume and possibly make changes, rendering the filer snapshots inconsistent from one another. For most sites, this inconsistency is rarely an issue. For sites where consistency is important, you can use the snapshot consistency command to put up a *VIP fence* for snapshots. This fence prevents client access to any VIP that supports this volume. This may affect multiple volumes. The fence stays up until the last filer completes its snapshot or checkpoint, or until a timeout expires.

Guidelines: Snapshot Presentation to Clients

By default, CIFS clients can access their snapshots with Windows Explorer. They select a file or directory, pull up its Properties, and find a list of snapshots for the file or directory in the Previous Versions tab. CIFS clients can use this interface to find and restore previous versions of their files and directories. Microsoft calls this the Volume Shadowing Service, or *VSS*, for Shared Folders. Managed volumes support VSS for Shared Folders, but direct volumes and NFS-only volumes do not.

◆ Note

*If CIFS clients were connected to the volume before you create your first snapshot-related rule (*snapshot replica-snap-rule*, notification rule, or this rule), the clients must shut down and restart all instances of Windows Explorer before they can see the* Previous Versions *tab. Windows Explorer only checks for snapshot support when it first connects to the share.*

For installations where administrators are the only CIFS clients allowed to access snapshots, you have commands for controlling VSS access:

- You can use snapshot privileged-access to select a limited number of CIFS clients (the site administrators) who can access snapshots.
- You can also (or alternatively) stop VSS access altogether with no snapshot vss-mode. Administrators can continue to access snapshots through a special pseudo directory, ~snapshot, that is not displayed by default.

Both of the above options affect CIFS clients only; they have no effect on NFS clients.

You can use snapshot directory cifs-name to change the name of the ~snapshot directory that CIFS clients see. To change the directory name seen by NFS clients, use snapshot directory nfs-name. You can also control the display of the directory based on the export: use the snapshot directory display volume-root-only command to display this directory only in ARX exports from the root of the volume, not in exports of the volume's subdirectories. These commands apply to NFS clients as well as CIFS clients.

Clients only see the snapshots that were invoked by the ARX rule or added to the rule with snapshot manage. Snapshots made independently on the back-end filer are not shown.

Guidelines: Data Domain Servers and Presentation Through NFS

EMC Data Domain servers provide a unique presentation to NFS clients; they are only accessible from volume root's .snapshot directory. The .snapshot directory under each lower directory does not include the Data Domain snapshots.

This limitation does not apply to CIFS clients, who can access the Data-Domain snapshots in every ~snapshot directory.

You can avoid this NFS limitation if every directory on the Data Domain server(s) is striped to another filer type. To ensure universal striping, migrate all directories in the volume to a share from a different vendor and promote them to "master:"

- create a filename-fileset that recursively matches everything in the volume (using path exact / and recurse),
- create a place-rule, and
- use the from (gbl-ns-vol-plc) command to have that fileset match directories and promote them.

In a tiered volume, you can use this place rule to migrate all master directories to a Tier 1 share. This is a best practice for tiered volumes.

Guidelines: Filer-Snapshot Names

Each EMC or NetApp snapshot has a unique name to distinguish it as a backing snapshot for a particular ARX. The snapshot names use the following format

`acopia_`*`id_time-stamp_uuid_filer-volume`*

where

- *id* is an integer that the snapshot software uses internally.
- *time-stamp* is in *yyyymmddHHMM* format. This is the time that the snapshot was created.
- *uuid* is the Universally-Unique ID that identifies the ARX. In a redundant pair, this is the UUID for the peer that was originally configured as senior/active, no matter which peer is currently active. You can find an ARX's UUID in the output of show running-config.
- *filer-volume* is the name of the filer's volume.

Windows servers use a UUID in curly braces ({}) to identify each snapshot. This does not change for ARX-generated snapshots.

Guidelines: Configuration Changes in a Scheduled Rule

If you make a configuration change in a snapshot rule that is running with a schedule, the configuration change is ineffective until the next time the rule runs.

Guidelines: Removing the Snapshot Rule

The no form of the command removes the rule without removing any snapshots from the back-end filers. To remove the snapshots from the filers, use the snapshot remove command *before* you remove the rule. This is an efficient method for cleaning all of the supporting snapshots behind the rule. If supporting snapshots remain when you invoke no snapshot rule, the CLI lists all remaining snapshots when it prompts for confirmation.

Guidelines: Snapshot Reconstitution

An accidental snapshot-rule removal would separate the back-end snapshots from the ARX configuration, requiring a *reconstitution* of the coordinated snapshot. There are other situations where you may require snapshot reconstitution, too, such as a site-to-site failover: if a filer mechanism duplicates all of the filer snapshots from Site A over to Site B, and each site is managed by its own ARX pair, the snapshots at Site B need to be reconstituted in Site B's ARX pair (see cluster-name and activate configs for details about site-to-site failovers). The snapshot-reconstitution process requires some preparation when you start adding snapshot rules to the configuration. These guidelines show the high-level process for preparing your snapshots, along with the process for reconstituting snapshots in the event of an issue.

Guidelines: Preparing for Snapshot Reconstitution

After you add your first snapshot rule, choose an external repository to hold all of the latest snapshot reports. Choose a filer with a large amount of free space, or monitor the filer's space frequently and remove outdated snapshot reports. Every snapshot report has a timestamp in its name, so it is unique.

The filer should be able to run Perl scripts, and requires the XML::Simple module. You can download this Perl module from CPAN (http://search.cpan.org) if your system does not already have it.

Use the at command together with copy ftp, copy scp, copy {nfs|cifs}, or copy smtp to regularly copy snapshot reports from the ARX to your chosen filer. Use the common string for all of your snapshot reports, together with an "*" or other wildcard. The report repository should always hold the latest snapshot reports, so that they have the latest back-end snapshot names. For example, this command copies all reports starting with "snap" to an external IP each morning:

```
bstnA(cfg)# at 01:19:18 every 1 day do "copy reports snap*
ftp://ftpuser:ftpuser@172.16.100.183//var/arxSnapRpts/ format xml"

  The scheduled execution time for CLI command is: 3/30/09 1:19 AM.
```

You also need a copy of the "snap-recon.pl" script from the "software" directory on the ARX:

```
bstnA# copy software snap-recon.pl
ftp://ftpuser:ftpuser@172.16.100.183//var/arxSnapRpts/
```

Guidelines: Reconstituting ARX Snapshots

You can perform snapshot reconstitution if you have the snap-recon.pl script and the latest set of snapshot reports on a host that supports Perl. Start by running the snap-recon.pl script on that host. This produces a CLI script with a sequence of snapshot manage commands. By default, the output script is named, "snapRecon.cli." This script has several options; execute the command without any options to get a complete list. You must use the "**--report-dir *directory***" option to specify the *directory* that holds the reports. For example, this command sequence lists the files on "client2:/var/arxSnapRpts," runs snap-recon.pl on the reports in the current directory (.), and then shows the new file in the directory:

```
juser@client2:/var/arxSnapRpts$ ls
snap_daily_0_create_20090330010418883.xml  snap_hourly_0_create_20090330010652616.xml
snap_daily_0_create_20090330010454303.xml  snap_hourly_0_create_20090330010723024.xml
snap_daily_0_create_20090330010524663.xml  snap-recon.pl
snap_daily_2_remove_20090330010612220.xml
juser@client2:/var/arxSnapRpts$ ./snap-recon.pl --report-dir .
juser@client2:/var/arxSnapRpts$ ls
snap_daily_0_create_20090330010418883.xml  snap_hourly_0_create_20090330010652616.xml
snap_daily_0_create_20090330010454303.xml  snap_hourly_0_create_20090330010723024.xml
snap_daily_0_create_20090330010524663.xml  snapRecon.cli
snap_daily_2_remove_20090330010612220.xml  snap-recon.pl
juser@client2:/var/arxSnapRpts$
```

Guidelines: Reconstituting ARX Snapshots (Cont.)

You have the option to edit the CLI script before you run it on the ARX. For example, the ARX volumes at a disaster-recovery site may have different names than those at the primary site. The snapRecon.cli script has several options for substituting these names; you can use these to rebuild the CLI script, or you can manually edit the CLI script.

Once the CLI script is ready, you can download it to the ARX and run it. Use copy ftp, copy {nfs|cifs}, or copy scp for the download, and use run to run it.

A large, complex CLI script may contain errors. If you discover any back-end snapshots that are mismatched with their ARX counterparts, you can use the snapshot clear command to remove the ARX snapshot (not the rule) from the configuration. Then you can edit and re-run the script, or you can use snapshot manage to manually incorporate the back-end snapshots into the correct ARX snapshots.

Samples

```
bstnA(gbl-ns-vol[access~/G])# snapshot rule nightly
This will create a new policy object.

Create object 'nightly'? [yes/no] yes
bstnA(gbl-ns-vol-snap[access~/G~nightly])# ...
```

instantiates a new snapshot rule, “nightly,” in the current namespace volume.

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# no snapshot rule hourlySnap

The following snapshots are being managed by the rule 'hourlySnap':

hourlySnap_0(acopia_27_200707170500_501f705c-8735-11d8-8936-a58a9e4556
df_datavol2)

hourlySnap_2(acopia_25_200707170300_501f705c-8735-11d8-8936-a58a9e4556
df_datavol2)

hourlySnap_1(acopia_26_200707170400_501f705c-8735-11d8-8936-a58a9e4556
df_datavol2)

If this command is confirmed, the rule is deleted and all associated
data related to the aforementioned snapshots will be removed from the
switch.  The snapshots on the filers will not be removed.

Proceed? [yes/no] yes
bstnA(gbl-ns-vol-snap[medarcv~/lab_equipment])# ...
```

removes a snapshot rule with some existing snapshots remaining on one of its back-end filers.

Related Commands

proxy-user
filer-type
proxy-user (gbl-filer)
ip address `... management`
manage snapshots

retain
exclude
schedule (gbl-ns-vol-...snap)
report (gbl-ns-vol-...snap)
enable (gbl-ns-vol-...snap)
show snapshots

snapshot consistency

snapshot vss-mode
snapshot directory display
snapshot directory cifs-name
snapshot directory nfs-name

snapshot manage
snapshot create
snapshot verify
snapshot remove

snapshot clear

snapshot verify

Purpose A *snapshot* is a full copy of an ARX volume at one point in time. The ARX volume coordinates snapshots at each of its back-end filers. A back-end filer may clean up some of its snapshots, including those that back an ARX snapshot, to conserve disk space. Use the snapshot verify command to confirm that all filer snapshots are in place behind a snapshot rule, or behind a particular ARX snapshot.

Modes priv-exec

Security Role(s) storage-engineer, backup-operator, or crypto-officer

Syntax `snapshot verify` ***`namespace vol-path snapshot-rule`*** `[`***`snapshot-instance`***`]`

namespace (1-30 characters) identifies the namespace.

vol-path (1-1024 characters) is the name of the volume.

snapshot-rule (1-1024 characters) is the snapshot rule to verify.

snapshot-instance (optional, 1-68) identifies a particular snapshot created by the snapshot rule.

Default(s) None.

Guidelines This command produces a separate verification report for each ARX snapshot. Each report shows the configuration of the snapshot rule, a summary status for each ARX snapshot, and details about the filer snapshot(s) behind each ARX snapshot. See Figure 30.5 on page 30-54 for a sample verification report.

This command verifies filer snapshots asynchronously, allowing you to continue entering CLI commands while the operation proceeds. You can use the tail reports *report-name* follow command to follow the progress of each snapshot verification. You can also use wait-for snapshot verify to wait for all snapshot verifications to finish; this is especially useful in CLI scripts.

To create a snapshot rule (which establishes the name of a snapshot set, the number of snapshots to retain in the set, and an optional schedule), use the snapshot rule command. For a snapshot rule without a schedule, or for a scheduled snapshot that failed, you can use the snapshot create command to manually create a snapshot. To remove filer snapshots from behind a snapshot rule, use snapshot remove.

Samples

```
bstnA# snapshot verify medarcv /lab_equipment hourlySnap

  Starting snapshot operation in volume /lab_equipment, report name:
snap_hourly_0_verify_20120229040155141.rpt


  Starting snapshot operation in volume /lab_equipment, report name:
snap_hourly_2_verify_20120229040155141.rpt


  Starting snapshot operation in volume /lab_equipment, report name:
snap_hourly_1_verify_20120229040155141.rpt
```

verifies the integrity of all of the "hourlySnap" rule's snapshots.

```
bstnA# snapshot verify medarcv /lab_equipment dailySnap dailySnap_0

  Starting snapshot operation in volume /lab_equipment, report name:
snap_daily_0_verify_20120229040155259.rpt
```

verifies the most-recent snapshot from the "dailySnap" rule. The report contents appear in Figure 30.5, below.

Related Commands

wait-for snapshot verify
snapshot rule
snapshot create
snapshot remove

Figure 30.5 *Sample Report: snap_daily_0_verify_....rpt*

```
bstnA# show reports snap_daily_0_verify_20120229040155259.rpt

  Snapshot Rule Summary
  ---------------------
   Namespace Name:               medarcv
   Volume Name:                  /lab_equipment
   Snapshot Rule Name:           dailySnap


  Snapshot Properties
  -------------------
   Snapshots Enabled:            Yes
   Guarantee Consistency:        Disabled
   Retain Count:                 7
   Schedule:                     daily4am
   CIFS Directory Name:          ~snapshot
   Directory Display:            All Exports
   Hidden File Attribute:        Not Set
   Restricted Access Configured: Yes
   VSS Mode:                     None
   Contents:
     Metadata:                   No
     Volume Configuration:       No
     User Snapshots:             Yes
   Archive:
     Total Archive Operations:   0
     Total Successful Operations: 0
     Total Failed Operations:    0
```

```
   Total Saved Metadata:        0 B
   Total Saved Volume Config:   0 B
   Average Copy Rate:           0 b/s

Snapshot Summary - dailySnap_0
------------------
 Snapshot Name:                 dailySnap_0
 Snapshot Operation:            Verify
 Result:                        Success
 Time Requested:                02/29/2012 04:01:55 -0500
 Time Created:                  02/29/2012 04:00:01 -0500
 Last Time Verified:            02/29/2012 04:01:58 -0500
 Request:                       Verify
 Snapshot State:                Sparse
 Snapshot Origin:               Schedule
 Report Name:                   snap_daily_0_verify_20120229040155259.rpt


Included Shares
---------------
     Share Name:                equip (user data)
     Filer:
       Name:                    nas10
       CIFS Share:              equipment
       Volume:                  vol2
       Filer Snapshot:          acopia_18_201202290900_d9bdece8-9866-11d8-91e3-f48e42637d58_vol2

     Share Name:                leased (user data)
     Filer:
       Name:                    nas10
       CIFS Share:              for_lease
       Volume:                  vol1
       Filer Snapshot:          acopia_18_201202290900_d9bdece8-9866-11d8-91e3-f48e42637d58_vol1

     Share Name:                backlots (user data)
     Filer:
       Name:                    fs2
       CIFS Share:              backlot_records
       Volume:                  E:\
       Filer Snapshot:          {cff0aeb2-7da2-47da-a868-144cb84c9822}
       Time Created:            02/29/2012 04:00:03 -0500


Excluded Shares
---------------
     Share Name:                scanners
     Filer:
       Name:                    fs5
       CIFS Share:              xraysScanners
     Reason:                    Snapshots were not supported on this type of back-end filer.

     Share Name:                equipSnap
     Filer:
       Name:                    nas11
       CIFS Share:              equipBkup
     Reason:                    Replica-snap share excluded from managed share snapshot.

     Share Name:                leasedSnap
     Filer:
       Name:                    nas11
       CIFS Share:              leasedBkup
     Reason:                    Replica-snap share excluded from managed share snapshot.
```

snapshot vss-mode

Purpose Windows clients can view volume *snapshots*, or point-in-time copies, by clicking on any file or directory in the volume, opening the Properties menu, and selecting the Previous Versions tab. This is called the Volume Shadowing Service (*VSS*). The ARX supports VSS for Windows XP and later clients by default, but not Windows 2000 clients. You can use the snapshot vss-mode pre-xp command to support VSS for Windows 2000 clients, too, at the cost of a slight performance penalty.

Some sites prefer to allow snapshot access to administrators only; for those sites, you can use snapshot vss mode none. You can offer other methods to access snapshots to administrators, as described in the *Guidelines* below.

Use no snapshot vss-mode to return to your default VSS support: VSS for Windows XP (and later) clients.

Modes gbl-ns-vol

Security Role(s) storage-engineer or crypto-officer

Syntax

```
snapshot vss-mode {xp | pre-xp | none}
no snapshot vss-mode
```

xp | pre-xp | none is a required choice. This selects the Windows-client version(s) for whom the volume supports VSS:

xp causes the volume to support VSS for Windows XP and later clients. This option excludes Windows 2000 clients.

pre-xp extends the volume's VSS support to Windows 2000 clients as well as some later versions of Windows. This option may make VSS unusable for Windows 7 or later clients. You should only use this option if you have Windows 2000 clients, and no Windows 7 or later clients.

none prevents all VSS support from this volume. When this is set, all CIFS clients must use other means to access their snapshots. This can be useful in an installation where only administrators are allowed to access snapshots.

Default(s)

```
snapshot vss-mode xp
```

Guidelines

This command only applies to managed volumes and CIFS clients. Direct volumes do not support VSS, and NFS clients do not use VSS.

ARX-managed volumes support VSS by default, allowing all Windows-XP and later clients to access their snapshots. Support for the older Windows 2000 VSS precludes support for Windows 7 VSS, so the default excludes Windows 2000 clients; you can use this command to support Windows 2000 clients instead of Windows 7 clients if your site requires it. For installations where only administrators have access to snapshots, you can disable VSS altogether with the none option.

Whether or not VSS mode is supported, CIFS clients can always access snapshots through "~snapshot," a pseudo directory in the volume's front-end shares. (NFS clients see a different name for this directory; typically ".snapshot.") One ~snapshot directory resides in every directory in the volume. This directory is not displayed by default, but well-informed clients can access the directory by name (for example, by typing **cd ~snapshot** even though **dir** does not display the ~snapshot directory). To display the directory name in directory listings (such as **dir** and **ls**), use snapshot directory display. You can use the snapshot directory cifs-name command to change the directory's name for CIFS clients, and you can use the snapshot directory nfs-name command to change it for NFS clients.

To create a snapshot rule (which establishes the name of a snapshot set, the number of snapshots to retain in the set, and an optional schedule), use the snapshot rule command.

Samples

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# snapshot vss-mode none
```

stops supporting snapshot access through VSS in the "medarcv~/lab_equipment" volume. Windows administrators can continue to access snapshots through the volume's pseudo directories, typically named "~snapshot."

```
bstnA(gbl-ns-vol[accss~/G])# snapshot vss-mode xp
```

supports VSS for all Windows clients of the "access~/G" volume except Windows 2000 clients.

Related Commands

snapshot rule
snapshot directory cifs-name
snapshot directory nfs-name
snapshot directory display
export (gbl-cifs)

wait-for snapshot create

Purpose Use the wait-for snapshot create command to wait until a manual snapshot is created on all of the filers behind the volume.

Mode (any)

Security Role(s) storage-engineer, crypto-officer, or backup-operator

Syntax **`wait-for snapshot create`** ***`namespace vol-path rule`***
`[`*`snapshot-instance`*`] [timeout` ***`timeout`***`]`

namespace (1-30 characters) is the name of the namespace.

vol-path (1-1024 characters) identifies the volume.

rule (1-1024 characters) is the name of the snapshot rule.

snapshot-instance (optional; 1-255 characters) identifies a snapshot behind the rule. If you omit this, it waits for the "0" snapshot (for example, "hourly_0" or "COB_0"). This is useful for snapshot-create operations that use a non-default snapshot name.

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - none, wait indefinitely

snapshot-instance - the default name for the most-recent snapshot: ***rule***_0.

Guidelines This command is useful in CLI scripts.

When manually invoking a snapshot rule with the snapshot create command, you can use the wait-for snapshot create command to wait for the operation to complete on all of the back-end filers. This can be useful for CLI scripts, which you can copy onto the switch (with copy ftp, copy scp, copy {nfs|cifs}, or copy tftp), and then run.

If you set a ***timeout*** and it expires before all filer snapshots are finished, the command exits with a warning. To interrupt the wait-for snapshot create command, press **<Ctrl-C>**.

Sample `prtlndA# wait-for snapshot create nemed /vol44 COB timeout 30`
waits up to 30 seconds for the filers behind the 'nemed~/vol44' volume to create their snapshots.

Related Commands snapshot create

wait-for snapshot remove

Purpose Use the wait-for snapshot remove command to wait until a snapshot-removal operation has completed on every back-end filer behind the volume.

Mode (any)

Security Role(s) storage-engineer, crypto-officer, or backup-operator

Syntax **`wait-for snapshot remove`** ***`namespace vol-path rule`*** **`[`*`snapshot-instance`*`] [timeout`** ***`timeout`*****`]`**

namespace (1-30 characters) is the name of the namespace.

vol-path (1-1024 characters) identifies the volume.

rule (1-1024 characters) is the name of the snapshot rule.

snapshot-instance (optional; 1-255 characters) identifies the ARX snapshot. If you omit this, it waits for the removal of all filer snapshots behind the "0" snapshot (for example, "hourly_0" or "COB_0").

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - none, wait indefinitely

snapshot-instance - the default name for the most-recent snapshot: ***rule***_0.

Guidelines This command is useful in CLI scripts.

When removing filer snapshots with the snapshot remove command, you can use the wait-for snapshot remove command to wait for the operation to complete. That is, this command waits for all of the volume's back-end filers to remove the snapshots behind a particular ARX snapshot. This can be useful for CLI scripts, which you can copy onto the switch (with copy ftp, copy scp, copy {nfs|cifs}, or copy tftp), and then run.

If you set a ***timeout*** and it expires before the filer snapshots are removed, the command exits with a warning. To interrupt the wait-for snapshot remove command, press **<Ctrl-C>**.

Sample

```
bstnA> wait-for snapshot remove medarcv /lab_equipment hourlySnap
timeout 60
```

waits up to 60 seconds for the snapshots behind the 'hourlySnap_0' snapshot to be removed from all filer shares behind the 'medarcv~/lab_equipment' volume. (This command defaults to 'hourlySnap_0' because no ARX-snapshot instance was specified in the command.)

Related Commands snapshot remove

wait-for snapshot verify

Purpose You can use the snapshot verify command to confirm that all of the filer snapshots behind an ARX-snapshot rule still exist. Use the wait-for snapshot verify command to wait until a snapshot-verification has completed.

Mode (any)

Security Role(s) storage-engineer, crypto-officer, or backup-operator

Syntax **`wait-for snapshot verify`** ***`namespace vol-path rule`*** **`[`*`snapshot-instance`*`] [timeout`** ***`timeout`*****`]`**

namespace (1-30 characters) is the name of the namespace.

vol-path (1-1024 characters) identifies the volume.

rule (1-1024 characters) is the name of the snapshot rule.

snapshot-instance (optional; 1-255 characters) identifies a snapshot behind the rule. If you omit this, it waits for the verification of the "0" snapshot (for example, "hourly_0" or "COB_0").

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - none, wait indefinitely

snapshot-instance - the default name for the most-recent snapshot: ***rule***_0.

Guidelines This command is useful in CLI scripts.

When verifying the integrity of an ARX snapshot with the snapshot verify command, you can use the wait-for snapshot verify command to wait for the operation to complete. This can be useful for CLI scripts, which you can copy onto the switch (with copy ftp, copy scp, copy {nfs|cifs}, or copy tftp), and then run.

If you set a ***timeout*** and it expires before the verification is complete, the command exits with a warning. To interrupt the wait-for snapshot verify command, press **<Ctrl-C>**.

Sample

```
bstnA# wait-for snapshot verify medarcv /lab_equipment hourlySnap
timeout 90
```

waits up to 90 seconds for the ARX to verify all filer snapshots behind the 'hourlySnap_0' snapshot still exist on their back-end filers. (This command defaults to 'hourlySnap_0' because no ARX-snapshot instance was specified in the command.)

Related Commands snapshot verify

31

File Tracking

archive

Purpose

A *file-history archive* holds frequent snapshots of a managed volume's configuration and (typically) its metadata. You can use the archive command to make the current snapshot rule send its configuration/metadata snapshots to a pre-configured file-history archive. After the snapshot rule starts sending this data to the archive, you can later query the archive for the volume's state at different times. These queries are useful for finding the back-end locations of the volume's files on different dates, as the volume migrates them from filer to filer.

Use no archive to stop sending any snapshots to a file-history archive. This disables any record keeping by the current snapshot rule.

Mode

gbl-ns-vol-snap

Security Role(s)

storage-engineer or crypto-officer

Syntax

`archive` ***`archive-name`***
`no archive`

archive-name (1-64 characters) identifies the file-history archive where this rule should copy its configuration/metadata snapshots.

Default(s)

`no archive`

Guidelines

You can use show file-history archive for a list of archives on this ARX, and you can use show file-history archive ... contents to see the volume configuration and metadata that the archive holds.

Only volume configuration and volume metadata are copied to a file-history archive; a snapshot rule does not copy any client files or directories. Use the contents command to control the information that is gathered by this snapshot rule.

Sample

```
bstnA(gbl-ns-vol-snap[access~/G~daily])# archive betaArchive
```

uses the file-history archive named "betaArchive" to store configuration and metadata snapshots for the "access~/G" volume.

Related Commands

namespace **->** volume **->** snapshot rule
contents
file-history archive
show file-history archive
show file-history archive ... contents

cancel snapshot archive

Purpose A *file-history archive* holds daily snapshots of a managed volume's configuration and (typically) its metadata. You create a snapshot rule in the volume that snapshots this information and then sends the snapshot to the volume's archive. A file-history archive resides on an external filer. In case of an issue with copying the snapshot to the archive filer, you can use the cancel snapshot archive command to cancel the copy operation.

Mode priv-exec

Security Role(s) crypto-officer, storage-engineer, or backup-operator

Syntax **`cancel snapshot archive namespace`** ***`ns`*** **`volume`** ***`path`*** **`rule`** ***`rule-name`***

ns (1-30 characters) is the namespace where a snapshot archive is in progress.

path (1-1024 characters) identifies the specific volume, and

rule-name (1-1024 characters) is the snapshot rule.

Default(s) None

Guidelines The CLI prompts for confirmation before canceling the archive operation; enter **yes** to proceed.

This cancels the archiving process, not the snapshot process. Use this command to resolve an issue between the ARX and the filer behind the archive's location.

Sample

```
bstnA# cancel snapshot archive namespace medarcv volume /rcrds rule
rcrdsArchive

Confirming this command will cause all archiving operations in
namespace 'medarcv' volume '/rcrds'
associated with rule 'rcrdsArchive' to be cancelled. Proceed? [yes/no]
yes
```

cancels the archive operation for the "rcrdsArchive" rule.

Related Commands namespace `->` volume `->` snapshot rule
file-history archive
location

clear file-history archive

Purpose A *file-history archive* is where one or more managed volumes store their configuration and metadata snapshots. The managed volumes may store up to seven-years worth of snapshots, possibly taken on a daily basis. This is useful for querying past states of a managed volume, so that you can find file locations as they migrate between the volume's back-end filers. As time goes on, you can use the clear file-history archive command to remove older snapshots that are no-longer needed.

Modes priv-exec

Security Role(s) crypto-officer, storage-engineer, or backup-operator

Syntax **`clear file-history archive`** ***`name`*** **`[metadata-only] [namespace`** ***`ns`*** **`volume`** ***`vol`*** **`]`**

name (1-64 characters) identifies a particular file-history archive. This is the particular archive from which you will clear some (or all) file-history records.

metadata-only (optional) indicates that you want to delete volume metadata only, not the volume configuration. Volume metadata consumes much more space than the configuration, and it supports the show file-history virtual-service command for seeking file locations. That command would no longer be supported, but you can still use the show virtual path-history command to query the front-end service names of various times. If you omit this option, both metadata and configuration are deleted and neither query is supported.

namespace *ns* (optional, 1-30 characters) focuses on archive records from the given namespace.

volume *vol* (required after the **namespace** option, 1-1024 characters) narrows the focus to the archive records for a single volume.

Syntax: One Date **`clear file-history archive`** ***`name`*** **`[metadata-only] date`** ***`date`*** **`...`**

name is the file-history archive, as described above.

metadata-only (optional) is also described above.

date *date* (*mm/dd/yyyy*, such as 06/05/2009, or "today") is a specific date for the file-history records to be cleared.

... represents the **namespace** and **volume** options described above.

Syntax: Contents During a Range of Dates

```
clear file-history archive name [metadata-only] start-date s-date ...
clear file-history archive name [metadata-only] end-date e-date ...
clear ... [metadata-only] start-date s-date end-date e-date ...
```

name and **metadata-only** are described above.

start-date *s-date* (*mm/dd/yyyy*, such as 03/02/2009) is the beginning of the date range. If you provide no **end-date** (as shown in the first line above), the end date is today.

end-date *e-date* (*mm/dd/yyyy*, such as 11/12/2009) is the end of the date range. If you provide no **start-date** (as shown in the second line above), the start is the date of the first snapshot in the file-history archive(s).

... at the end of each syntax line represents the **namespace** and **volume** options described above.

Syntax: Contents During a Time Leading Up to a Date

```
clear file-history archive name [metadata-only] count {days|weeks|quarters|months|years} before e-date ...
```

name and **metadata-only** are described above.

count (1-100) is the number of days, weeks, or whatever time unit you choose next.

days|weeks|quarters|months|years chooses the time unit. For example, “2 weeks,” “3 quarters,” or “7 months.”

before *e-date* (*mm/dd/yyyy*, such as 09/30/2009, or “today”) is the end of the time period.

... at the end of each syntax line represents the **namespace** and **volume** options described above.

Guidelines

The CLI prompts for confirmation before clearing any records from the archive filer; enter **yes** to proceed.

Before you use this command, you can use show file-history archive ... contents to list the file-history records on the archive.

Sample

```
bstnA# clear file-history archive fileRecordsMed date 01/07/2009

WARNING: Confirming this command removes all archived metadata and
volume configuration data dated
 between '01/07/2009' and '01/07/2009' inclusive from archive
'fileRecordsMed'.

Proceed? [yes/no] yes
```

clears all the file-history records from the fileRecordsMed archive that were recorded on January 7.

Related Commands

file-history archive
show file-history archive ... contents

contents

Purpose A *file-history archive* holds daily snapshots of a managed volume's configuration and (typically) its metadata. You can use the contents command to choose the contents of this rule's snapshots: client-visible files, volume-configuration data, and/or volume metadata.

Use no contents to eliminate one or more types of data from the current rule's snapshots.

Mode gbl-ns-vol-snap

Security Role(s) storage-engineer or crypto-officer

Syntax

```
contents {user-data | volume-config [metadata]}
no contents {user-data | volume-config}
```

user-data | volume-config [metadata] is a required choice.

user-data represents files and directories that are visible to clients. This causes the rule to create snapshots on the volume's storage filers.

volume-config represents a configuration file that describes this volume. This option only applies to a managed volume. This does not result in any snapshots unless you also choose the metadata option. The volume configuration can support queries with the show virtual path-history command, but not the show file-history virtual-service command.

– **metadata** (optional) includes the managed volume's metadata along with its configuration. This causes the rule to take a snapshot at the filer behind the volume's metadata share. The metadata occupies more disk space than the volume configuration, but it also supports file-specific queries with the show file-history virtual-service command.

Default(s)

```
contents user-data
```

Guidelines A snapshot rule typically includes volume configuration/metadata *or* user data, not both. If the rule is designed for backing up client data, the default is sufficient. If the rule sends its data to a file-history archive, the contents are usually limited to volume configuration and metadata.

A snapshot rule, if dedicated to volume configuration and metadata, deletes its metadata snapshot as soon as the snapshot is successfully copied to the archive filer. That is, the snapshot rule removes the snapshot from the filer behind the metadata share, and it removes all records of the snapshot from the ARX database. This conserves space on the metadata filer; the snapshot is not required after its contents are safely in the archive.

For a snapshot rule that contains all forms of data, user-data and volume-config, the rule adheres to its retain count for all of its snapshots. This includes the snapshots in the volume's filer shares as well as the volume's metadata snapshots. The snapshot remove command can remove all of the rule's filer snapshots, including those on the volume's metadata filer.

Sample

```
bstnA(gbl-ns-vol-snap[wwmed~/var~a])# contents volume-config metadata
```

causes the "a" snapshot rule to keep the volume's configuration and metadata in its snapshots.

```
bstnA(gbl-ns-vol-snap[wwmed~/var~a])# no contents user-data
```

stops the "a" snapshot rule from taking snapshots on its user-data filers.

Related Commands

namespace -> volume -> snapshot rule
archive
file-history archive

description (gbl-archive)

Purpose Use the optional description command to set a descriptive string for the current file-history archive. A *file-history archive* is where managed volumes can store file-location records over a long period of time.

Use the no form of the command to delete the description.

Mode gbl-archive

Security Role(s) storage-engineer or crypto-officer

Syntax

```
description text
no description
```

text (1-255 characters) is your description. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s) `no description`

Guidelines This description appears in the output of the show file-history archive command.

Sample

```
bstnA(gbl-archive[flArch])# description “archives for Med Cos”
```

specifies a description for the current file-history archive.

Related Commands file-history archive
show file-history archive

file-history archive

Purpose

Over time, a file may migrate through multiple shares in a managed volume. You can use the *file-tracking* facility to track these migrations. The file-history archive command creates (or edits) a storage depot for file-location records. A single file-history archive can hold these records for multiple managed volumes.

Use the no form of the command to remove the configuration for a file-history archive from the ARX database.

Modes

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

`file-history archive` ***`name`***
`no file-history archive` ***`name`***

name (1-64 characters) is a name you choose for the file-history archive.

Default(s)

None

Guidelines

The CLI requests confirmation before creating the new archive object; enter **yes** to confirm. This places you into gbl-archive mode, where you use the location command to establish the filer share or ARX volume where the archive resides. You can also enter an optional archive description (gbl-archive) in this mode.

You can create up to 24 file-history-archive objects on the ARX.

The file-history records in an archive are snapshots, or point-in-time copies, of a managed volume's metadata and configuration. To keep these records for a particular managed volume, you create a snapshot rule in the volume. The contents of the snapshot rule are the volume's configuration and metadata instead of (or in addition to) client files and directories; use the contents command to set these options. Then use the archive command to send copies of the volume data to this archive.

After one or more snapshot rules post their snapshots in a file-history archive, you can query the archive for file or directory locations at the time of each snapshot. Use the show file-history archive command to list all file-history archives, and use show file-history archive ... contents to list their contents (that is, their snapshots of volume configuration and metadata). To find which client-visible services existed at a time in the past, use the show virtual path-history command. The show file-history virtual-service command shows the locations of client files and/or directories on a given date or a range of dates. You can use these commands to find the location of a file on a date when it was backed up, and to find the location of the same file today.

Samples

```
bstnA(gbl)# file-history archive flArch
This will create a new archive.

Create archive 'flArch'? [yes/no] yes
bstnA(gbl-archive[flArch])# ...
```

instantiates a new file-history archive named "flArch."

```
bstnA(gbl)# no file-history archive betaArchive
bstnA(gbl)# ...
```

removes a file-history archive.

Related Commands

location
description (gbl-archive)
snapshot rule
contents
archive
show file-history archive
show virtual path-history
show file-history virtual-service

find

Purpose Use the find command to locate one file's back-end filer and path. This finds a file's location as of now, not from a previous time.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

`find host` *`hostname-or-ip`* `{nfs| cifs}` *`share-name`* `path` *`path`* `[verbose]`

hostname-or-ip identifies a virtual server by its external DNS hostname or Virtual-IP address. Use show virtual service for a list of all virtual servers and their Virtual IP addresses. (To look up DNS names at an external DNS server, you must first identify the server with the ip name-server command.)

nfs | cifs is a required choice, determining the type of share that follows.

share-name (1-4096 characters) is the name of the front-end share, as seen by clients. The show server-mapping command shows all front-end shares on the ARX, in the left column.

path (1-4096 characters) is the virtual-file path within the front-end share.

verbose (optional) adds the file's NFS filehandles to the output.

`find global-server` *`fqdn`* `{nfs| cifs}` *`share-name`* `path` *`path`* `[verbose]`

fqdn (1-255 characters) identifies a global server by its fully-qualified domain name (FQDN). Use show global service for a list of all global servers and their FQDNs.

The remaining options are described above.

`find namespace` *`namespace`* `path` *`path`* `[verbose]`

namespace (1-30 characters) identifies the namespace. Use show namespace for a list of all namespaces.

path (1-4096 characters) is the virtual-file path within the namespace. This is the file path as seen by the client, starting with the volume path.

verbose (optional) adds the file's NFS filehandles to the output.

`find wins` *`wins-name`* `{nfs| cifs}` *`share-name`* `path` *`path`*

wins-name (1-255 characters) identifies a virtual server by its optional wins-name or one of its optional wins-alias names.

The remaining options are described above.

Default(s) None.

Guidelines

This command finds one file's physical location. To find all file locations behind a namespace, use the nsck ... report metadata-only command.

For files in multi-protocol (NFS and CIFS) namespaces, this shows the file's physical location from an NFS point of view and a CIFS point of view.

The output contains several fields to map the ARX-volume location of the file to the file's physical location on the filer:

Namespace identifies the namespace that contains the file.

Logical Path shows the full path to the file from the client's perspective. This path starts from the root of the volume.

NFS Physical Location is an NFS path to the physical file. This appears in *ip-address*:/*path-to-file* format.

Managed Volume Path indicates that the ***path*** is on a direct volume, and the physical location is on an ARX managed-volume that is standing in as a filer. This field does not appear in the output unless it applies to this query.

CIFS Physical Location is a CIFS path to the physical file. This appears in //*ip-address*/*path-to-file* format. If this is an NFS-only file in a multi-protocol (CIFS and NFS) volume, a message here indicates that the file has an inconsistent name. You can use the nsck ... report inconsistencies command to find all of the volume's inconsistent names.

The **verbose** output shows two NFS filehandles for the file:

Virtual File Handle is the filehandle that the ARX presents to clients.

Physical File Handle is the filehandle that the server gave to the ARX.

The filehandles may be incomplete for the find namespace command. Use one of the other versions (such as find host) for complete NFS filehandles.

If there are any NFS symlinks (or symbolic links) in the ***path***, they appear next. Each symlink in the path appears on one line, with a text arrow (->) pointing to the symlink's target file or directory. For example, a ***path*** of "/myvol/dirA/Ydir/sub" could show this entry, indicating that "dirA" is a symlink:

```
/myvol/dirA -> /myvol/Xdir
```

Each symlink path is fully resolved up to (and not including) the name of the symlink itself. For example, the above symlink indicates the real path for the "dirA" symlink, "/myvol/Xdir." If "sub" was also a symlink, its entry starts with that real path, "/myvol/Xdir," instead of the "dirA" name used in the ***path***:

```
/myvol/Xdir/Ydir/sub -> myvol/test/dir3
```

This second symlink entry shows the real path up to the "sub" symlink, and points to the real path for "sub."

The object is a symlink appears at the end if the final file or directory in the path is a symlink. This appears in the "*symlink-path -> real-object-path*" format shown above, where the *symlink-path* is a real path up to (but not including) the final name in the path.

Guidelines: File Tracking

This identifies a file's location now; for a file's location in the past, you can use *file tracking*. This can be useful for finding the correct filer-backup tape for a lost or compromised file.

The file-tracking feature requires some configuration before you can make any file queries. Specifically, a managed volume requires a snapshot rule to regularly copy its configuration and metadata to a file-history archive. After some copies of metadata have been archived, you can use the show file-history virtual-service command to query file locations at different dates.

Samples

```
bstnA# find namespace wwmed path /acct/index.html

Namespace:             wwmed
Logical path:          /acct/index.html
NFS Physical location: 192.168.25.19:/exports/budget/index.html
```

finds a file in a namespace.

```
bstnA# find host 192.168.25.14 nfs /claims path /stats/piechart.ppt

Namespace:              insur
Logical path:           /claims/stats/piechart.ppt
NFS Physical location:  192.168.25.21:/vol/vol1/NTFS_QTREE/insurance/stats/piechart.ppt
CIFS Physical location: The file/directory on //192.168.25.21/insurance has an inconsistent
name.
```

finds a file behind a virtual server. The virtual server is exporting the file from a multi-protocol namespace, so it shows both NFS and CIFS locations for the file. This file has an inconsistent, NFS-only name.

```
bstnA# find global-server ac1.MEDARCH.ORG nfs /claims path /stats/piechart.ppt

Namespace:              insur
Logical path:           /claims/stats/piechart.ppt
NFS Physical location:  192.168.25.21:/vol/vol1/NTFS_QTREE/insurance/stats/piechart.ppt
CIFS Physical location: The file/directory on //192.168.25.21/insurance has an inconsistent
name.
```

finds the same file behind its global server.

```
bstnA# find global-server insur.medarch.org cifs CLAIMS path /index.html verbose

Namespace:              insur
Logical path:           /claims/index.html
NFS Physical location:  192.168.25.21:/vol/vol1/NTFS_QTREE/insurance/index.html
CIFS Physical location: The file/directory on //192.168.25.21/insurance has an inconsistent
name.

  NFSv3
    Virtual File Handle  (32 bytes):
      0x00000001140000000100000000000700000000000000000000000000f00000000
    Physical File Handle (32 bytes):
      0x92660a0074e7b00020000000000202c09254ed00a58c004840000000bb438000
```

provides a verbose display for another file behind the same global server.

```
bstnA# find global-server insur.medarch.org cifs CLAIMS path
/common/reagentLists/stats/acmeIns.txt

Namespace:              insur
Logical path:           /claims/common/reagentLists/stats/acmeIns.txt
NFS Physical location:  192.168.25.21:/vol/vol1/NTFS_QTREE/insurance/stats/acmeIns.txt
CIFS Physical location: //192.168.25.21/insurance/stats/acmeIns.txt

The following symlinks were dereferenced to find this path:
/claims/common/reagentLists -> /claims/specs
/claims/specs/stats -> /claims/stats
```

finds another file behind the same global server, where the logical path contains two symlinks.

Related Commands nsck ... report metadata-only

location

Purpose

Use the location command to identify the filer or ARX volume to be used for the current file-history archive. A *file-history archive* is where managed volumes can store file-location records over a long period of time. This is the specific storage location for those file-location records.

Use the no form of the command to disconnect from the location and disable any file-tracking queries against it.

Mode

gbl-archive

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
location filer name {nfs3 | nfs3tcp} share [path path]
location filer name cifs share proxy-user proxy [path path]
location namespace ns volume vol-path [path path]
no location
```

name (1-64 characters) is external-filer name for this archive's back-end filer. This is the filer's name from the ARX configuration, which may or may not match the actual host name of the filer.

nfs3 | nfs3tcp chooses the file-access protocol for this location. This is a required choice between NFSv3 over UDP (**nfs3**) or NFSv3 over TCP (**nfs3tcp**).

share (1-64 characters) is the name of the back-end export or share.

path (optional, 1-1024 characters) is a specific sub path in the back-end share (or ARX volume) where file-location records should be stored. We recommend a path that is hidden from most clients.

proxy (for a CIFS filer, 1-32 characters) is the proxy-user configuration that the ARX uses as its identity when writing to or reading from the back-end-CIFS share. This requires read and write permissions, but does not necessarily need to belong to the Backup Operators group on the filer.

ns (1-30 characters) identifies a namespace on this ARX.

vol-path (1-1024 characters) identifies a volume in the namespace above. If you use this syntax to place the file-history archive in a managed volume, you can use migration policies to balance the file-history records among multiple back-end filers.

Default(s)

```
no location
```

Guidelines

This command determines the physical location for the current file-history archive. Choose a filer or ARX volume with enough storage space for up to 7 years of file-history records. Each file-history record is a snapshot of the volume's configuration (and, typically, its metadata). Use the snapshot rule command to create a snapshot rule for each desired volume, then use the contents command to determine whether or not metadata is included in the volume's snapshots.

Multiple managed volumes can share the same file-history archive. The schedule (gbl-ns-vol-...snap) for each volume's snapshot rule determines the frequency of these records, and the amount of space required over 7 years. For a current reading on the disk space consumed in the archive, you can use the show file-history archive ... contents command.

Guidelines: Archive Directory Structure

The snapshot rules create the following directory structure in the location filer's ***share*** and ***path***:

```
/acopia_file_history
    /ARX-name-ARX-GUID
        /yyyy
            /mm
                /dd
                    /namespace-name
                        /volume-name
                    /namespace-name
                        /volume-name
                /dd
                ...
            /mm
            ...
        /yyyy
        ...
```

Guidelines: Client Exposure to the Archive

The only clients who should access this "acopia_file_history" structure are management personnel who understand its purpose. Uninformed clients may accidentally delete important archive records. Choose a location that is inaccessible to most clients. If you use a managed volume as your archive location, use a hidden directory for this data, or dedicate the entire volume to file-history archives and only export/share it to management clients (using export (gbl-cifs) or export (gbl-nfs)).

Guidelines: Moving the Location

For cases where the filer location is running low on disk space, you can use the no location command to stop all archive writes to the filer. You can then enter a new location command to identify the replacement filer or volume. To preserve all of your current file-history records at the new location, move the "acopia_file_history" directory from the former location to the new one.

Samples

```
bstnA(gbl-archive[fileRecordsMed])# location filer fs4 cifs
arx_file_archv proxy-user repCred
```

sets the location for the "fileRecordsMed" archive to a CIFS share on a back-end filer.

```
bstnA(gbl-archive[flArch])# location namespace wwmed volume /acct path
fileArch
```

sets the location for the "flArch" archive to a path in an ARX volume, "wwmed~/acct."

Related Commands

file-history archive
show file-history archive ... contents
contents

show file-history archive

Purpose

Over time, a file may migrate through multiple shares in a managed volume. You can set up a managed volume to periodically record its file locations on a file-history archive. Later, you can use commands to query the archive for the location of a file on a given date or dates. The show file-history archive command lists the archives on this ARX, their configurations, and the managed volumes that use them.

Modes

(any)

Security Role(s)

crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show file-history archive
show file-history archive name [namespace ns volume vol]
```

name (optional, 1-64 characters) identifies a particular file-history archive. This invokes a detailed view of the archive, including its configuration and the volumes that use it. If you omit this option, the output is a list of all file-history archives on the ARX.

namespace *ns* (optional, 1-30 characters) focuses on archive records from the given namespace.

volume *vol* (required after the **namespace** option, 1-1024 characters) narrows the focus to the archive records for a single volume.

Guidelines

You can use this command to list the current file-history archives on the ARX, or to show the configuration of a particular archive. Before you can list any archives, you must use the file-history archive command to create an archive, and the location command to choose the archive's home filer.

The related show file-history archive ... contents command shows the archive's contents. The archive gets its contents from one or more snapshot rules in one or more managed volumes. By showing the contents, you can learn how much disk space the archive is using on its back-end filer or ARX volume.

Guidelines: Summary Output

The summary version of the command lists all file-history archives in a table. The table contains one row per archive with the following fields:

Archive Name is the name of the archive, defined by the file-history archive command.

Description offers a short description for the archive. You can set or change this with the description (gbl-archive) command.

Guidelines: Detailed Output

If you identify a particular archive in the command, the output contains configuration details for that archive:

Name is the name of the archive, defined by the file-history archive command.

Location is the archive's home filer or home volume. You can use the location command to change this.

Description offers a short description for the archive. You can set or change this with the description (gbl-archive) command.

Path is the path under the above location where the archive resides. The location command controls this.

Protocol is "CIFS" or "NFS." This is the file-access protocol that the ARX uses to copy files to the archive. For an archive on an external filer, the location command controls this. For an archive on a managed volume, this is controlled by the volume's namespace configuration.

Proxy User is the name of the ARX's proxy-user credentials for accessing the archive filer or managed volume. After the name of the proxy-user configuration, the windows-domain (gbl-proxy-user) and user (gbl-proxy-user) appears in parentheses.

Status is "Online" for an archive that is currently usable.

FreeSpace shows the current free space on the archive.

In use by is the title for a table of snapshot rules that add to this archive. The table contains one row per snapshot rule:

Namespace is the namespace that contains the snapshot rule.

Volume identifies the managed volume for this rule.

Snapshot Rule is the name of the snapshot rule that has written to this archive.

Samples

`bstnA(gbl)# show file-history archive`

lists all the file-history archives on the ARX. See Figure 31.1 for sample output.

`bstnA(gbl)# show file-history archive fileRecordsMed`

shows details about a particular file-history archive. See Figure 31.2 for sample output.

Related Commands

file-history archive
show file-history archive ... contents

Figure 31.1 *show file-history archive*

```
bstnA(gbl)# show file-history archive

Archive Name          Description
--------------------  ---------------------------------------
fileRecordsMed        archive share for ARX file histories
```

Figure 31.2 *show file-history archive fileRecordsMed*

```
bstnA(gbl)# show file-history archive fileRecordsMed

Name:           fileRecordsMed
Location:       \\192.168.25.29\arx_file_archv
Description:    archive share for ARX file histories
Path:           /
Protocol:       CIFS
```

```
Proxy User:     acoProxy2
Status:         Online
FreeSpace:      1.1 GB

In use by:

Namespace             Volume               Snapshot Rule
--------------------  -------------------  --------------------
medarcv               /rcrds               rcrdsArchive
medarcv               /lab_equipment       labArchive
```

show file-history archive ... contents

Purpose

Over time, a file may migrate through multiple shares in a managed volume. You can set up a managed volume to periodically record its file locations on a file-history archive. Later, you can use commands to query the archive for the location of a file on a given date or dates. The show file-history archive ... contents command lists the file-history records stored in a particular file-history archive, and shows the amount of disk space that the records occupy.

Modes

(any)

Security Role(s)

crypto-officer, storage-engineer, or backup-operator

Syntax: Contents on One Date

```
show file-history archive name contents date date [namespace ns
volume vol] [report prefix]
```

name (1-64 characters) is the file-history archive.

contents is a required keyword.

date *date* (*mm/dd/yyyy*, such as 05/06/2009, or "today") is the date for your query.

namespace *ns* (optional, 1-30 characters) focuses on archive records from the given namespace.

> **volume *vol*** (required after the **namespace** option, 1-1024 characters) narrows the focus to the archive records for a single volume.

report *prefix* sends the output to a report file with the given ***prefix*** (1-255 characters). The report name is *prefix_archive-name_mmddyyyyHHMM*.rpt, where archive name is "all" or the name you provide in the next option, and the rest is the time the report was created. The exact report name appears at the CLI prompt after you type the command.

Syntax: Contents During a Range of Dates

```
show file-history archive name contents start-date s-date ...
show file-history archive name contents end-date e-date ...
show ... contents start-date s-date end-date e-date ...
```

start-date *s-date* (*mm/dd/yyyy*, such as 05/06/2009) is the beginning of the date range. If you provide no **end-date** (as shown in the first line above), the end date is today.

end-date *e-date* (*mm/dd/yyyy*, such as 01/12/2009) is the end of the date range. If you provide no **start-date** (as shown in the second line above), the start is the date of the first snapshot in the file-history archive(s).

... at the end of each syntax line represents the **namespace**, **volume**, and **report** options described above.

Syntax: Contents During a Time Leading Up to a Date

```
show file-history archive name contents count
{days|weeks|quarters|months|years} before e-date ...
```

count (1-100) is the number of days, weeks, or whatever time unit you choose next.

days|weeks|quarters|months|years chooses the time unit. For example, "3 weeks," "2 quarters," or "5 months."

before *e-date* (*mm/dd/yyyy*, such as 09/04/2009) is the end of the time period.

... at the end of each syntax line represents the **namespace**, **volume**, and **report** options described above.

Guidelines

You can use this command to list the snapshots on a file-history archive. Before you can list any archive contents, you must use the file-history archive command to create an archive, and the location command to choose the archive's home filer or ARX volume.

The archive gets its contents from one or more snapshot rules in one or more managed volumes. A volume's snapshot rule contributes to a file-history archive through the following options: the rule must choose the volume's configuration (and possibly its metadata) with the contents command, and must identify the file-history archive with the archive command.

As time goes on, the metadata snapshots on the archive filer occupy progressively-more disk space. You can use this command to assess the consumed space, and you can use the clear file-history archive command to remove old snapshots from the archive filer.

For a list of all file-history archives, use show file-history archive.

Guidelines: Output

The output of the archives contents contains three tables.

Query Parameters displays all of your command-line choices with the following fields:

- Archive,
- Start date,
- End date,
- Namespace, and
- Volume.

Switch shows the hostname and the GUID for the ARX that writes to this archive. It is possible for multiple ARXes to use the same archive filer. You can also see the GUID at the top of the output for show running-config.

The next table lists the individual archive records. Each record appears in one row, with the following information:

- Archive Date/Time is the time stamp when the snapshot rule copied the volume configuration (and possibly the volume metadata) into the archive.
- Namespace:Volume identifies the managed volume for this file-history record.
- Snapshot Rule is the name of the snapshot rule that created the record.
- Config shows the size of the volume-configuration file.
- Metadata shows the size of the volume metadata. This is blank if the snapshot rule's contents did not include metadata snapshots. Without this, you cannot make the show file-history virtual-service query against this file-history record.

Summary shows the total disk space consumed by these records: total metadata space, total configuration space, and a grand total. You can use this as a guideline to determine the disk space you could reclaim with the clear file-history archive command.

Sample

`bstnA(gbl)# show file-history archive fileRecordsMed contents end-date today`

shows all of the file-history records in the "fileRecordsMed" archive, from the first up to today's. See Figure 31.3 for sample output.

Related Commands

file-history archive
show file-history archive
clear file-history archive

Figure 31.3 *show file-history archive ... contents end-date ...*

```
bstnA(gbl)# show file-history archive fileRecordsMed contents end-date today

Query Parameters
----------------
   Archive: fileRecordsMed
Start date: Jan  1 1970
  End date: Aug 31 2010
 Namespace:
    Volume:


Switch: bstnA (d9bdece8-9866-11d8-91e3-f48e42637d58)

Archive Date/Time     Namespace:Volume          Snapshot Rule  Config  Metadata
```

```
-------------------  ------------------------  -------------  ------  --------
Aug 31 2010 01:29:00  medarcv:/lab_equipment     labArchive     2.1 kB  32 kB
Aug 31 2010 01:34:00  medarcv:/rcrds             rcrdsArchive   2.2 kB  64 kB


Summary:
  Total space consumed by Volume Metadata:       96 kB
  Total space consumed by Volume Configuration:  4.3 kB
  Grand total space consumed:                    100 kB
```

show file-history virtual-service

Purpose Managed volumes migrate files between back-end filers according to your policy settings. This can create a challenge for sites that use data-protection devices to back-up their data; if such a device backs up files from Filer A that later migrate to Filer B, you would create a metadata inconsistency in the managed volume if you restored the files back to Filer A. This command, along with some archiving commands, finds the filer location for any file on any date. You can use a snapshot rule to record a managed volume's configuration and metadata to a *file-history archive*, and later use the show file-history virtual-service command to query that archive. The output from this command shows the filer locations for a file or directory on any given date or date range.

Modes (any)

Security Role(s) crypto-officer, storage-engineer, backup-operator, or operator

Syntax

```
show file-history virtual-service fe-service fe-share
time-frame {[file file-name [path dir-path]]} [options]
```

fe-service (1-1024 characters) is the name of an ARX service used by clients. This is a front-end service created by either the cifs or nfs commands. You can use the show virtual path-history command to find the services that existed on an earlier date, and the paths from those services to your filer shares; service-to-filer mappings may change over time. This can also be the VIP (from the virtual server command), a wins-name, or a wins-alias that existed at the time.

fe-share (1-1024 characters) identifies a particular front-end share in the above service. This is created by either the export (gbl-cifs) or export (gbl-nfs) commands, which export storage from an ARX volume to ARX clients. As above, you may need to use the show virtual path-history command to find the shares that existed at the time you are querying. For NFS shares, this is case-sensitive (so that "/vol/aShare" is different from "/vol/ASHARE"), but this is not case-sensitive for CIFS shares.

time-frame is a date or range of dates. The output shows the file or directory locations as of these dates. The options are:

`date {today | date}`
for file locations on a specific date. Enter the ***date*** in ***mm/dd/yyyy*** format.

`start-date date`
for file locations after the given ***date***.

`end-date date`
for file locations before the given ***date***.

`start-date date-1 end-date date-2`
for file locations between two dates.

`count {days|weeks|quarters|months|years} before date`
for file locations for a time leading up to a date. The ***count*** (1-100) is the number of days, weeks, or other time-unit leading up to the ***date***.

Syntax (Cont.)

[file [*file-name*] [path *dir-path*] is a required choice: you can choose a file name, a directory path, or both.

file-name (1-255 characters) specifies the name of one or more files. You can use special *wildmat* characters (such as ***** for zero or more characters, **?** for any single character, **[a-z]** for any lowercase character, or **[^0-9]** for any character except a number) for a group of similar files.

dir-path specifies the name of one or more directories. Use forward slashes (/) for this path, even for CIFS shares. As with the file-name, you can use special *wildmat* characters to specify a group of paths. If you omit this, the path is the root of the volume.

options (all optional) are various search and output options that you can apply to any form of this command. If you use more than one of them, they must be in this order:

`recurse`

searches all directories in the volume or **path**. Without this option, the command only searches the root of the volume or path.

`case-sensitive`

makes the search case-sensitive, so that a search for "*.DOC" does not match the "myfile.doc" file anymore. By default, "*.DOC" matches "myfile.doc," "yourfile.Doc," "herFile.DOC," or any other case combination for the ".doc" extension.

`verbose`

adds more detail to the output. For a search with only a **path**, this cause the output to show all of the *files* inside all of the matching directories.

`report prefix`

sends the output to a report file with the given ***prefix*** (1-255 characters). The report name is *prefix_archive-name_mmddyyyyHHMM*.rpt, where archive name is "all" or the name you provide in the next option, and the rest is the time the report was created. The exact report name appears at the CLI prompt after you type the command.

`archive archive-name`

identifies a particular file-history archive to be searched.

Guidelines

You can use this command to find file/directory locations at different times, typically to identify a backup tape for a lost file or directory. Before you can find any archived files or directories, you must use the file-history archive command to create an archive, and the location command to choose the archive's home (either an external filer or an ARX volume). The archive gets its file-history records from one or more managed volumes. Each managed volume requires a snapshot rule to send its configuration and metadata to the archive filer. The snapshot rule contributes to a file-history archive through the following options: the rule must choose the volume's configuration *and* its metadata with the contents command, and must identify the file-history archive with the archive command.

This command finds file locations from the time when a file-history record was copied to a file-history archive. To find the location of a file now, use the find command or the nsck ... report metadata-only command.

To restore a file to a managed volume, copy it from its backup tape to a staging area on the back-end filer. The staging area is a directory that is *outside* any imported volume. You can mount both the staging area and the ARX service as a client, then copy the file(s) from the staging area to the volume. Alternatively, you can use the restore data command to automatically move the file(s) from the staging area to the volume.

Guidelines: Output

The output contains a table showing the parameters of your query, followed by some configuration information for the file-history archive and ARX service, and then a separate table for each archive record containing one of the chosen files.

Query Parameters contains several fields to summarize your command-line options:

- Virtual Server,
- Export,
- Start date,
- End date,
- File Name,
- Path,
- Archive,
- Case Sensitive,
- Recursive, and
- Verbose.

The next two fields identify the archive and its host ARX. These fields and the tables below only appear if the operation found at least one of the selected files or directories:

- Archive is the name of the file-history archive containing these records.
- Hosting Switch identifies this ARX by its hostname and by its GUID (also shown in the output of show running-config).

Guidelines: Output (Cont.)

The final series of tables contains the found files and/or directories. Each table shows a single matching file in one archive record:

Archive Date/Time is the specific time for the current file-location record.

Global Server is the global server that clients use to access the chosen files.

WINS Name and

WINS Aliases are aliases that clients may use to access the service. You can set these with the wins-name and wins-alias commands.

Dynamic DNS Names are service names that are registered at the local dynamic-DNS servers (typically, Windows DCs). You can use the dynamic-dns command to add a new dynamic-DNS name for this service.

VIP is the virtual-IP address used by clients to access this ARX service, set by the virtual server command.

Namespace is the ARX namespace that held the file at the above time.

Volume Path identifies the file's volume.

File Server(s) shows the single filer that held the file at the time, or the series of filers that may have held the directory. Each filer has three fields to describe the exact location of the file or directory:

- Shared Path is the path to the file or directory as seen from the network. This starts with the filer's CIFS-share name or NFS-export name.
- Physical Path identifies the exact, file-system path to the file or directory. This path is the one that could be used to find a backup on an external data-protection device.

  The ARX accesses the filer's CLI to get the correct physical-path information. The external-filer configuration on the ARX must have the proper filer-type and proxy-user (gbl-filer) for this access. Otherwise, this path is a "best guess:" for an NFS share, it uses the back-end-export path identified in the filer command; for a CIFS share, it sends a path query to the filer, requiring the namespace's proxy-user (gbl-ns) to be in the Administrators group.
- File Name is the name of the file or directory.
- Symlink Name appears instead of File Name if the file or directory is a symlink. This is the name of the symlink itself, not the symlink's target. You must retrieve the symlink from a backup tape to discover its target.

Samples

```
bstnA(gbl)# show file-history virtual-service ac1.medarch.org labs date
04/23/2009 file index.html
```

shows the location of a file named "index.html" in the "ac1.medarch.org~labs" share. The file's location is as of a particular date, April 23. See Figure 31.4 for sample output.

```
bstnA# show file-history virtual-service ac1.medarch.org ARCHIVES date
today path planA recurse
```

shows the location of all directories named "planA" in the "ac1.medarch.org~ARCHIVES" share. These directories are in the given locations as of today. See Figure 31.5 for sample output.

Related Commands show virtual path-history
file-history archive
find
restore data

Figure 31.4 *show file-history virtual-service ... date ... file ...*

```
bstnA(gbl)# show file-history virtual-service ac1.medarch.org labs date 08/30/2010 file index.html

Query Parameters
----------------
Virtual Server: ac1.MEDARCH.ORG
        Export: labs
    Start date: Aug 30 2010
      End date: Aug 30 2010
     File Name: index.html
          Path:
       Archive: All
Case Sensitive: No
       Recurse: No
       Verbose: No

Archive:           fileRecordsMed
Hosting Switch:    bstnA (d9bdece8-9866-11d8-91e3-f48e42637d58)

  ------------------------------------------------------------------------------
  Archive Date/Time: Aug 30 2010 01:16:00
  Global Server:     ac1.MEDARCH.ORG
  WINS Name:
  WINS Aliases:
  Dynamic DNS Names: ac1, fs1, fs2, fs5
  VIP:               192.168.25.15
  Namespace:         medarcv
  Volume Path:       /lab_equipment
  File Server(s):

      Shared Path: \\nas10\for_lease
    Physical Path: /vol/vol1/NTFS-QTREE/for_lease
        File Name: index.html

bstnA#
```

Figure 31.5 *show file-history virtual-service ... date ... path ... recurse*

```
bstnA# show file-history virtual-service ac1.medarch.org ARCHIVES date today path planA recurse

Query Parameters
----------------
Virtual Server: ac1.MEDARCH.ORG
        Export: ARCHIVES
    Start date: Aug 30 2010
      End date: Aug 30 2010
     File Name:
          Path: planA
       Archive: All
Case Sensitive: No
       Recurse: Yes
       Verbose: No

Archive:           fileRecordsMed
Hosting Switch:    bstnA (d9bdece8-9866-11d8-91e3-f48e42637d58)
```

```
------------------------------------------------------------------------------
Archive Date/Time: Aug 30 2010 01:21:00
Global Server:     ac1.MEDARCH.ORG
WINS Name:
WINS Aliases:
Dynamic DNS Names: ac1, fs1, fs2, fs5
VIP:               192.168.25.15
Namespace:         medarcv
Volume Path:       /rcrds/2004/planA
File Server(s):

    Shared Path: \\fs1\histories
  Physical Path: d:\exports\histories\2004\planA


------------------------------------------------------------------------------
Archive Date/Time: Aug 30 2010 01:21:00
Global Server:     ac1.MEDARCH.ORG
WINS Name:
WINS Aliases:
Dynamic DNS Names: ac1, fs1, fs2, fs5
VIP:               192.168.25.15
Namespace:         medarcv
Volume Path:       /rcrds/2005/planA
File Server(s):

    Shared Path: \\fs1\histories
  Physical Path: d:\exports\histories\2005\planA


------------------------------------------------------------------------------
Archive Date/Time: Aug 30 2010 01:21:00
Global Server:     ac1.MEDARCH.ORG
WINS Name:
WINS Aliases:
Dynamic DNS Names: ac1, fs1, fs2, fs5
VIP:               192.168.25.15
Namespace:         medarcv
Volume Path:       /rcrds/2001/planA
File Server(s):

    Shared Path: \\fs1\histories
  Physical Path: d:\exports\histories\2001\planA


------------------------------------------------------------------------------
Archive Date/Time: Aug 30 2010 01:21:00
Global Server:     ac1.MEDARCH.ORG
WINS Name:
WINS Aliases:
Dynamic DNS Names: ac1, fs1, fs2, fs5
VIP:               192.168.25.15
Namespace:         medarcv
Volume Path:       /rcrds/2000/planA
File Server(s):

    Shared Path: \\fs1\histories
  Physical Path: d:\exports\histories\2000\planA
```

show virtual path-history

Purpose

Service and volume configurations change over time, possibly making it difficult to make a proper file-tracking query about an earlier date. To make a query about file locations as of six-months ago, you must first know which ARX services existed at the time, and the paths from those client-visible services to the back-end filers and their file-systems. This is known as the *path history* of your ARX. From any mode, you can use the show virtual path-history command to list the front-end services, namespaces, and volumes from a given date or range of dates.

Modes

(any)

Security Role(s)

crypto-officer, storage-engineer, or backup-operator

Syntax: One Date

```
show virtual path-history date date [report prefix] [archive
arch-name] [verbose]
```

date (*mm/dd/yyyy*, such as 11/12/2009, or **today**) is the exact date for the query.

report *prefix* (optional, 1-255 characters) sends the output to a report file. Each report has a unique name in the following format:

prefix_archive_yyyymmddHHMM.rpt

where the *prefix* is part of this option, the *archive* is "all" or a specific ***arch-name*** (see below), and *yyyymmddHHMM* is the time the report was created. For example, "fRec_all_200903031200.rpt" could be the name for one report with the "fRec" prefix.

archive *arch-name* (optional, 1-64 characters) focuses on path history from a single file-history archive.

verbose (optional) expands the output with additional details.

Syntax: A Range of Dates

```
show virtual path-history start-date s-date ...
show virtual path-history end-date e-date ...
show virtual path-history start-date s-date end-date e-date ...
```

start-date *s-date* (*mm/dd/yyyy*, such as 05/06/2009) is the beginning of the date range. If you provide no **end-date** (as shown in the first line above), the end date is today.

end-date *e-date* (*mm/dd/yyyy*, such as 10/24/2009) is the end of the date range. If you provide no **start-date** (as shown in the second line above), the start is the date of the first snapshot in the file-history archive(s).

... represents the **report**, **archive**, and **verbose** options described above.

Syntax: Time Leading Up to a Date

```
show virtual path-history count
{days|weeks|quarters|months|years} before e-date ...
```

count (1-100) is the number of days, weeks, or whatever time unit you choose next.

days|weeks|quarters|months|years chooses the time unit. For example, "2 weeks," "5 days," or "6 months."

before *e-date* (*mm/dd/yyyy*, such as 03/02/2009) is the end of the time period.

... represents the **report**, **archive**, and **verbose** options described above.

Guidelines

You can use this command to find ARX-service and storage configurations at different times, typically to prepare for using the show file-history virtual-service command. Before you can find any volume-configuration information, you must use the file-history archive command to create an archive, and the location command to choose the archive's home (an ARX volume or an external filer). The archive gets its volume-configuration records from one or more managed volumes. Each managed volume requires a snapshot rule to send its volume configuration to the archive filer. The snapshot rule contributes to a file-history archive through the following options: the rule must choose the volume's configuration with the contents command, and must identify the file-history archive with the archive command.

The output contains a table with the chosen date range followed by one table per matching file-history record.

Query Parameters contains several fields to summarize your command-line options:

- Start date,
- End date,
- Archive, and
- Verbose.

As of *date* is the heading for every file-history record in your date range. If more than one volume sends its records to this archive, you may see more than one record with the same date.

The next series of fields describes the configuration for the archive and the service:

- Archive is the name of the file-history archive containing this record.
- Hosting Switch identifies this ARX by its hostname.
- Global Server is the global server that clients use to access the Volume, below.
- WINS Name and
- WINS Alias are aliases that clients may use to access the service. You can set these with the wins-name and wins-alias commands.
- Dynamic DNS Names are service names that are registered at the local dynamic-DNS servers (typically, Windows DCs). You can use the dynamic-dns command to add a new dynamic-DNS name for this service.
- VIP is the virtual-IP address used by clients to access this ARX service, set by the virtual server command.
- Namespace is the ARX namespace to which this record applies.
- Volume identifies the volume for this file-history record.

Guidelines (Cont.) A table of the volume's shares appears next. This shows the back-end shares that were included in the volume when the record was created. Each row shows one share, with the following fields:

Share is the ARX name of the share.

Shared Path is the filer's network-accessible path for the above share. The volume imported the filer share through this path, and used the share's storage for client files and directories when the record was created.

-> Physical Path is the actual file-system path on the filer. An external data-recovery application uses file-system paths to identify the files that it has backed up. You can use this to identify the backup files that apply to the volume as of this particular date. The verbose option adds the filer's IP address to the end of this output.

The ARX accesses the filer's CLI to get the correct physical-path information. The external-filer configuration on the ARX must have the proper filer-type and proxy-user (gbl-filer) for this access. Otherwise, this path is a "best guess:" for an NFS share, it uses the back-end-export path identified in the filer command; for a CIFS share, it sends a path query to the filer, requiring the namespace's proxy-user (gbl-ns) to be in the Administrators group.

The final table shows the exports that were used by the ARX's clients when the record was taken. Each row contains one client-visible export of the above volume, with the following fields:

Export is the exported name of the managed volume, or one of its sub paths.

Relative Virtual Path is the managed-volume path that clients saw as the root of the above Export.

Sample

```
bstnA(gbl)# show virtual path-history end-date 06/05/2009 report
pathsTilJune verbose
Generating report: pathsTilJune_all_200901260629.rpt
```

shows all path histories up to June 5, including verbose details, and sends them to a report.

Related Commands file-history archive
show file-history virtual-service

Figure 31.6 *show virtual path-history date ... verbose*

```
bstnA(gbl)# show virtual path-history date today verbose

Query Parameters
----------------
Start date: Aug 31 2010
  End date: Aug 31 2010
   Archive: All
   Verbose: Yes

As of Aug 31 2010 01:29:00

  Archive:           fileRecordsMed
  Hosting Switch:    bstnA (d9bdece8-9866-11d8-91e3-f48e42637d58)
  Global Server:     ac1.MEDARCH.ORG
  WINS Name:
  WINS Aliases:
  Dynamic DNS Names: ac1, fs1, fs2, fs5
```

```
  VIP:               192.168.25.15
  Namespace:         medarcv
  Volume:            /lab_equipment

  Share                      Shared Path -> Physical Path
  -------------------------  ------------------------------------------------
  equip                      \\nas10\equipment -> /vol/vol2/NTFS-QTREE/equipment (192.168.8.141)
  leased                     \\nas10\for_lease -> /vol/vol1/NTFS-QTREE/for_lease (192.168.8.141)
  backlots                   \\fs2\backlot_records -> e:\exports\backlot_records (192.168.25.27)
  scanners                   \\fs5\xraysScanners -> e:\exports\xraysScanners (192.168.25.71)

  Export                     Relative Virtual Path
  -------------------------  ------------------------------------------------
  labs                       \
  xraysScanners              \
  acopia#lab_equipment$      \


As of Aug 31 2010 01:34:00

  Archive:           fileRecordsMed
  Hosting Switch:    bstnA (d9bdece8-9866-11d8-91e3-f48e42637d58)
  Global Server:     ac1.MEDARCH.ORG
  WINS Name:
  WINS Aliases:
  Dynamic DNS Names: ac1, fs1, fs2, fs5
  VIP:               192.168.25.15
  Namespace:         medarcv
  Volume:            /rcrds

  Share                      Shared Path -> Physical Path
  -------------------------  ------------------------------------------------
  rx                         \\fs4\prescriptions -> d:\exports\prescriptions (192.168.25.29)
  charts                     \\fs1\histories -> d:\exports\histories (192.168.25.20)
  bulk                       \\fs2\bulkstorage -> e:\exports\bulkstorage (192.168.25.27)

  Export                     Relative Virtual Path
  -------------------------  ------------------------------------------------
  ARCHIVES                   \
  Y2005                      \2005
  bulkstorage                \
  acopia#rcrds$              \
  CELEBS                     \VIP_wing
  MP3S                       \2011\mp3downloads
  Y2004                      \2004
  Y2010                      \2010
```

32

Restore Data

cancel restore data

Purpose A *restore data* operation restores files from an external filer to a managed volume; the cancel restore data command stops a restore operation.

Modes priv-exec

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or backup-operator

Syntax **`cancel restore data`** ***`namespace`*** **`volume`** ***`volume`*** **`path`** ***`path`***

namespace (1-30 characters) is the namespace to which files are being restored. Use the show restore data command for a list of namespaces where restores are running.

volume (1-1024 characters) specifies the namespace volume where the restore is occurring.

path (1-1024 characters) specifies the volume directory. This path is relative to the ***volume*** root, above. Use forward slashes (/) for NFS or CIFS paths.

Default(s) None

Guidelines The CLI prompts for confirmation before canceling the restore operation; enter **yes** to continue.

Use the restore data command to start restoring files and directories to a volume. Restore operations happen in the background while you continue to use the CLI; you can use the wait-for restore data command to block the CLI until the restore is complete. To show all restore operations, use show restore data.

Sample

```
bstnA# cancel restore data wwmed volume /acct path /
Cancel specified restore operations? [yes/no] yes
bstnA#
```

cancels the restore operation in progress for the root directory in the "wwmed~/acct" volume.

Related Commands restore data
wait-for restore data
show restore data

clear restore data

Purpose The ARX holds its restore-operation records indefinitely, so that you can see the results of every restore operation through show restore data. Use the clear restore data command to remove restore operations from this running history.

Mode priv-exec

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or backup-operator

Syntax

```
clear restore data
clear restore data namespace
clear restore data namespace volume vol-path
clear restore data namespace volume vol-path path path
```

The ***namespace***, ***vol-path***, and ***path*** arguments mimic the ones used in the restore data command. These identify one or more restore operations; if you omit all of them, the command clears all restore-operation records from the switch.

namespace (1-30 characters) specifies a namespace where one or more restores have occurred.

vol-path (optional, 1-1024 characters) specifies a volume. Use forward slashes (/) for all volumes, including CIFS.

path (optional, 1-1024 characters) identifies a particular path that has been restored. As with volumes, use forward slashes in all paths.

Default(s) None

Guidelines The CLI prompts for confirmation before deleting any restore records. Enter **yes** to continue.

Use show restore data to show all recorded restore operations. This command clears them; after you run it, the show restore data command no longer displays the cleared operation(s).

The restore data command starts a restore operation (which runs in the background while you continue to use the CLI), cancel restore data cancels a running restore, and wait-for restore data blocks the CLI until a running restore completes.

Samples

```
bstnA# clear restore data wwmed
Clear restore data records? [yes/no] yes
bstnA#
```

clears all restore-operation records for the “wwmed” namespace.

```
bstnA# clear restore data
Clear restore data records? [yes/no] yes
bstnA#
```

clears all restore records from the switch. The show restore data command displays nothing after you do this.

Related Commands

restore data
show restore data
cancel restore data
wait-for restore data

restore data

Purpose To restore files and/or directories to a namespace volume, use the restore data command.

Modes priv-exec

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or backup-operator

Syntax

```
restore data namespace volume volume path dest-path
filer ext-filer {nfs nfs-export | cifs cifs-share}+ source-path
src-path [recurse] [remove-source]
```

namespace, **volume** ***volume***, and **path** ***dest-path*** specify the destination for the restored files:

namespace (1-30 characters) specifies the destination namespace for the restore operation. Use the show namespace command for a list of configured namespaces.

volume (1-1024 characters) is the volume to be restored. All of the namespace's volumes appear in the show namespace output.

dest-path (1-1024 characters) is the file or directory to be restored (for example, "var/log"). This path is relative to the ***volume*** root, above. If this is a directory, the restored files are placed into a new subdirectory under this path named "restore." Use forward slashes (/) for NFS or CIFS paths.

filer ***ext-filer***, **{nfs ... | cifs ...}+**, and **source-path** ***src-path*** define the source share for the restore operation:

ext-filer (1-64 characters) is the external filer where the backed-up files reside. This must be defined on the ARX as an external filer. Use the show external-filer command for a list of configured external filers.

{nfs ***nfs-export*** **| cifs** ***cifs-share*****}+** is a required choice. For multi-protocol (NFS and CIFS) filers, you must specify both.

- ***nfs-export*** (1-1024 characters) is an NFS-export path.
- ***cifs-share*** (also 1-1024 characters) is a CIFS-share name.

src-path (1-1024 characters) chooses the file or directory on the source share (for example, "/backups/vol3"). This path is relative to the ***nfs-export*** and/or ***cifs-share*** root, above. Use forward slashes (/) for NFS or CIFS paths.

You can enter the final two options in any order, or you can omit them:

recurse causes the restore operation to descend into subdirectories.

remove-source removes all of the selected files and directories from the source share after a successful restore. If there is any failure, the restore operation leaves the source share intact.

Default(s) **recurse** is off; the restore only restores from the top of the ***src-dir-path*** to the top of the ***dir-path*** without descending into any child directories.

remove-source is off; the source share is unaffected by the restore operation.

Guidelines: Backups

You can perform backup operations on back-end filers using the same procedure that was employed prior to the introduction of the ARX. You can use a data-protection device, called a *backup server* in this manual, to coordinate NDMP backups from your filers to tape.

Guidelines: Finding the Backup File

Before you begin a restore operation, you need to find the backed-up version of the lost or corrupted file. To accomplish this, you must first find the filer that held the file at the time of the desired backup. Assuming you know the service FQDN or VIP, the front-end-share name, and the file path of the lost file, you can use the find command to find the actual namespace, volume, and filer where the file currently resides.

For sites where migrations are frequent, or where backups are infrequent, you may need *file tracking* to find a file location on a particular backup date. The file-tracking process involves regular snapshots of the managed volume's configuration and metadata, which are stored on a file-history archive. Once this archive process is running, you can use the show virtual path-history command to query the archive for service names at the time of an earlier backup. This also shows the namespace, volumes, and filers that mapped to the service names at the time. Armed with this historical information, you can use show file-history virtual-service to query the archive for the file location at a given time.

Once you find the filer location for the lost file, you can go to your backup server to recover the file.

Guidelines: Sending the File to a Staging Area

Recover the file from the backup server and place it onto a *staging area* on a filer (typically, the same filer where the file was originally backed up). A staging area is a directory that is *not imported* into any managed volume.

The staging-area directory usually resides on a filer that is already defined as an external filer in the ARX configuration. If not, use the external-filer and ip address commands to add the new filer to the ARX database.

◆ Note

The external-filer configuration is required so that the ARX can support multi-homed backup servers for its restore operations.

Guidelines: Restoring Files from an External Staging Area

The restore data command restores the files and directories from the staging-area directory to their managed volume, which in turn sends the files/directories to the filer(s) specified by your storage policies.

Each restore operation has a unique job ID and generates a report of its progress. The reports follow this naming convention: "restore.*job-id*.*volume*.rpt." Use show reports to list all reports, including restore reports. To follow the progress of the restore operation, you can use tail reports report-name follow. A sample of this report appears in Figure 32.1 on page 32-8.

To view all restore operations (past and present), use show restore data. You can use the cancel restore data command to cancel a restore operation that is currently underway. To remove old records of restore operations, you can use the clear restore data command. Restore operations occur in the background so that you can continue to issue CLI commands; to pause the CLI until the restore operation to finishes, use wait-for restore data.

Guidelines: Invoking Restore from a Remote Host

From any remote host that runs the Secure SHell (SSH) program, you can run the restore data command and show the output of the report. Use the following syntax with **ssh** at the remote machine:

ssh *admin-user@mip "restore-data..."*

where

admin-user is a valid user (use show users to list all of them),

mip is a management-IP address for this ARX (use show interface mgmt and/or show interface vlan to list them), and

restore-data... is the restore data command, surrounded by quotation marks ("").

Samples

```
bstnA# restore data wwmed volume /acct path / filer das1 nfs
/exports/backups source-path /acct recurse
Scheduling restore operation on switch bstnA, report name:
restore.5._acct.rpt
bstnA#
```

restores files to the "wwmed~/acct" volume from an NFS-only share on a filer named "das1."

```
bstnA# restore data insur volume /claims path /stats filer nasE1 nfs
/root_vdm_4/backups cifs BACKUPS source-path /stats recurse
Scheduling restore operation on switch bstnA, report name:
restore.7._claims.rpt
bstnA#
```

restores files to the "insur~/claims" volume from a multi-protocol filer named "nasE1." This restores one directory tree in the volume, "/stats." A sample report appears below in Figure 32.1.

Related Commands

show server-mapping
export-mapping
external-filer
show restore data
cancel restore data
clear restore data
wait-for restore data
show reports

Figure 32.1 *Sample Report: restore.7._claims.rpt*

```
bstnA# show reports restore.7._claims.rpt
**** Restore Data Report: Started at Sat Feb 27 01:43:53 2010 ****
**** Software Version: 5.02.000.12541 (Feb 23 2010 20:12:44) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:
**** Namespace:         insur
**** Volume:            /claims
**** Path:              /stats/restore
**** Source Filer:      192.168.74.90
**** Source NFS Export: /root_vdm_4/backups
**** Source CIFS Share: BACKUPS
**** Source Path:       /stats
**** Options:           recurse

**** Legend:
 ****   OK = Object transferred without error.
```

```
 ****   SK = Skipped object due to naming problems.
 ****   FL = Encountered error during object transfer.

 ****   SY = Error re-syncing directory attributes after data restore.
 ****   DS = Error deleting source file or directory.
 ****   NR = Did not recurse due to problematic directory contents.
 ****   SA = Failed due to strict attribute consistency requirement.

 ****   FE = File entry.
 ****   DE = Directory entry.
 ****   LE = Hard link.
 ****   SE = Symbolic link.

**** Multi-Protocol Issues:
 ****   CC = Case-blind collision.
 ****   NE = Entry found with NFS that was not found with CIFS.
 ****   CE = Entry found with CIFS that was not found with NFS.
 ****   SL = Symbolic link found with NFS that was not found with CIFS.
 ****   IC = CIFS invalid characters found in NFS name.
 ****   NM = CIFS name has characters that are not mappable to the NFS encoding.
 ****   FN = A portion of the name contains a filer-generated pattern.
 ****   NC = Unable to copy CIFS data due to a filer-generated name.

Entry Type              Size   Object
-------------  -------------   --------------------------------------------------
[         CC]  [           ]   /stats/piechart.ppt
[         CC]  [           ]   /stats/PieChart.ppt
[         IC]  [           ]   /stats/on_the_job:2004.cnv
[         IC]  [           ]   /stats/on_the_job:2003.cnv
[         IC]  [           ]   /stats/in_home:2005
[OK    FE NC]  [  3,575,056]   /stats/PieChart.ppt
[OK    FE   ]  [      6,562]   /stats/acmeIns.txt
[OK    FE NC]  [    722,589]   /stats/piechart.ppt
[OK    FE NC]  [      1,679]   /stats/on_the_job:2003.cnv
[OK    FE NC]  [      9,509]   /stats/on_the_job:2004.cnv
[         IC]  [           ]   /stats/in_home:2005/age:11-21yrs.csv
[         IC]  [           ]   /stats/in_home:2005/age:>21yrs.csv
[         IC]  [           ]   /stats/in_home:2005/age:<10yrs.csv
[         IC]  [           ]   /stats/in_home:2005/age:>21yrs.csv
[         IC]  [           ]   /stats/in_home:2005/age:<10yrs.csv
[         IC]  [           ]   /stats/in_home:2005/age:11-21yrs.csv
[OK    DE NC]  [           ]   /stats/in_home:2005
[OK    FE NC]  [    335,872]   /stats/in_home:2005/age:11-21yrs.csv
[OK    FE NC]  [    267,264]   /stats/in_home:2005/age:>21yrs.csv
[OK    FE NC]  [     59,712]   /stats/in_home:2005/age:<10yrs.csv
[OK    FE   ]  [      2,432]   /stats/cleanBU.csh
[OK    FE   ]  [      4,545]   /stats/update.csh
[OK    FE   ]  [      1,423]   /stats/index.html
[OK    FE   ]  [      1,985]   /stats/makeCd.pl
[OK    FE   ]  [      4,110]   /stats/carrierCrossCheck.html
[OK    FE   ]  [          0]   /stats/21yrs.csv
[OK    DE   ]  [           ]   /stats/otj_latest
[OK    FE   ]  [      3,438]   /stats/otj_latest/feb.xls
[OK    FE   ]  [     71,462]   /stats/otj_latest/april.xls
[OK    FE   ]  [     17,279]   /stats/otj_latest/july.xls
[OK    FE   ]  [     17,629]   /stats/otj_latest/october.xls
[OK    FE   ]  [      3,321]   /stats/otj_latest/jan.xls

**** Total Found Items:                           21
**** Total Transferred Items:                     21
**** Total Failures:                               0
**** Total Bytes Restored:                 5,105,867
```

```
**** Total processed:              21
**** Elapsed time:           00:00:01
**** Restore Data Report: DONE at Sat Feb 27 01:43:54 2010 ****
```

show restore data

Purpose Use the show restore data command to show the results of all restore operations.

Modes (any)

Security Role(s) operator or backup operator (any)

Syntax

```
show restore data
show restore data namespace
show restore data namespace volume vol-path
show restore data namespace volume vol-path path path
```

namespace (optional, 1-30 characters) is the name of a namespace. Use show namespace for a list of all configured namespaces. If you omit this (and all other options), the output includes all recorded restore operations on all namespaces.

vol-path (optional, 1-1024 characters) focuses on a particular volume. Use forward slashes (/) for all volumes, including CIFS.

path (optional, 1-1024 characters) narrows the focus further, to a particular path within the volume. As with volumes, use forward slashes in all paths.

Default(s) None.

Guidelines This command shows all restore operations, past and present. They are organized by namespace, volume, and in-volume path that was restored:

Namespace is a namespace where at least one restore operation occurred.

Volume is one volume that was restored.

Path is the specific path that was restored. This contains details about the operation:

Options are "recurse" and/or "remove-source." This does not appear unless the user selected one of these options. They are options in the restore data command.

Source Filer is the name and IP address of the external filer used as a source. This is also specified in the restore data command.

- NFS Source identifies the NFS export used as a source share. This only appears if NFS was used in the restore operation.
- CIFS Source identifies the CIFS share used as a source. This only appears if CIFS was used in the restore operation.

Status is "Pending," "Restoring," "Pending Cancel," "Canceling," "Success," or "Failed." A typical restore operation goes from "Pending" to "Restoring" to "Success." Use the cancel restore data command to cancel the restore operation. For details on the restore operation's progress, you can look at the restore report, identified below.

Report Name is the name of a detailed report about the restore operation (and the ARX where the report is stored). You can use the show reports command to view the report.

Error only appears if the restore failed. This summarizes the problem.

Guidelines (Cont.)

Items Found are the files and/or directories that were found on the Source Filer share.

Items Transferred are the files and/or directories that were successfully restored to the target volume.

Failures is the number of files and/or directories that were *not* restored.

Total Bytes Restored is the number of bytes in the successfully-transferred files and directories.

Completed at shows the date and time when the restore operation finished.

This output grows with every run of the restore data command. Use clear restore data to clear all restore operations from this output (without deleting any of the detailed restore reports; use delete to remove those).

Samples

`bstnA# show restore data`

displays the state of all restore operations for all namespaces. Figure 32.2 shows sample output.

`bstnA# show restore data insur volume /claims`

displays the state of all restore operations for a multi-protocol (NFS and CIFS) volume, "insur~/claims." Figure 32.3 on page 32-13 shows sample output.

Related Commands

restore data
clear restore data
cancel restore data

Figure 32.2 *Sample Output: show restore data*

```
bstnA# show restore data

Namespace: wwmed
    Volume: /acct
        Path: /
            Options:              recurse
            Source Filer:         das1 [192.168.25.19]
            NFS Source:           /exports/backups/acct
            Status:               Success
            Report Name:          restore.6._acct.rpt (bstnA)
            Items Found:          415
            Items Transferred:    415
            Failures:             0
            Total Bytes Restored: 12,780,403
            Completed at:         Sat Feb 27 01:44:03 2010


Namespace: insur
    Volume: /claims
        Path: /stats
            Options:              recurse
            Source Filer:         nasE1 [192.168.25.51]
            NFS Source:           /root_vdm_4/backups/stats
            CIFS Source:          BACKUPS/stats
            Status:               Success
            Report Name:          restore.7._claims.rpt (bstnA)
            Items Found:          21
            Items Transferred:    21
            Failures:             0
```

```
            Total Bytes Restored: 5,105,867
            Completed at:         Sat Feb 27 01:43:54 2010
```

Figure 32.3 *Sample Output: show restore data insur volume /claims*

```
bstnA# show restore data insur volume /claims

Namespace: insur
    Volume: /claims
        Path: /stats
            Options:              recurse
            Source Filer:         nasE1 [192.168.25.51]
            NFS Source:           /root_vdm_4/backups/stats
            CIFS Source:          BACKUPS/stats
            Status:               Success
            Report Name:          restore.7._claims.rpt (bstnA)
            Items Found:          21
            Items Transferred:    21
            Failures:             0
            Total Bytes Restored: 5,105,867
            Completed at:         Sat Feb 27 01:43:54 2010
```

wait-for restore data

Purpose Use the wait-for restore data command to block the CLI until one or more restore operations completes.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or backup-operator

Syntax

```
wait-for restore data [timeout timeout]
wait-for restore data namespace [timeout timeout]
wait-for restore data namespace volume vol-path [timeout timeout]
wait-for restore data namespace volume vol-path path path [timeout timeout]
```

The ***namespace***, ***vol-path***, and ***path*** arguments mimic the ones used in the restore data command. These identify one or more restore operations; if you omit all of them, the command waits for all restore operations on the switch.

namespace (1-30 characters) specifies a namespace where one or more restore operations are occurring.

vol-path (optional, 1-1024 characters) specifies a volume. Use forward slashes (/) for all volumes, including CIFS.

path (optional, 1-1024 characters) identifies a particular path that is being restored. As with volumes, use forward slashes in all paths.

timeout (optional, 1-2096) is the timeout value, specified in seconds. If this timer expires before the restore operation completes, the command exits with a warning.

Default(s) ***timeout*** - 0 (zero, meaning "wait indefinitely")

Guidelines The wait-for restore data command waits for a specified restore operation to complete.

To interrupt the wait-for restore data command, press **<Ctrl-C>**. Use cancel restore data to cancel the restore operation altogether.

Samples

```
bstnA# wait-for restore data timeout 60
```

waits for all restore operations on the switch to complete. If they do not complete in 60 seconds, the command returns with a warning.

```
bstnA> wait-for restore data insur timeout 120
```

waits for all restore operations in the "insur" namespace to complete. If they do not complete in 120 seconds, the command returns with a warning.

Related Commands

restore data
show restore data
cancel restore data

33

Running and Global-Configs

clear global-config

Purpose

** To be used for evaluations and emergency recovery procedures only. Contact your F5 representative for more information.**

Use the clear global-config command to erase all global-config parameters.

The *global-config* parameters are shared among both ARX peers in a redundant pair: primarily, namespaces and global servers. The traditional running-config applies only to the local switch.

Mode

priv-exec

Security Role(s)

network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

`clear global-config`

Default(s)

None

Guidelines

This causes the switch to reboot. A prompt asks for confirmation before the CLI reboots the switch; enter **yes** to proceed.

If the switch is part of a redundant pair, you can only run this command on the senior switch. This reboots both switches in the pair.

Use this command to remove all namespaces, global servers, and other global-config parameters. This can be a useful alternative to delete startup-config, which deletes all running-config and global-config parameters. By preserving the running-config parameters, the network configuration remains intact; you can use the in-band (VLAN) management interfaces after clearing the global config.

However, this may leave back-end filers in a non-deterministic state. We strongly recommend using nsck ... destage on every active namespace before using this command.

Sample

```
bstnA# clear global-config
This removes all global configuration and reboots both chassis in the
redundant pair in order to ensure services are removed properly.

Are you sure? [yes/no] yes
bstnA#
Broadcast message (Fri Sep 24 13:53:15 2004):

The system is going down for reboot NOW!
...
```

deletes all namespaces, global servers, and other shared parameters.

Related Commands

delete startup-config
nsck ... destage

copy global-config

Purpose The *global-config* parameters are shared among both ARXes in a redundant pair. The traditional running-config applies only to the local switch. We recommend copying the global config to a remote server after any global-configuration change. Use the copy global-config command to copy the switch's global configuration to a local file, a remote server, an ARX volume, a remote ARX, or an email recipient.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax: FTP Upload

```
copy global-config ftp://[user[:password]@]server/file-name
```

ftp://[*user*[:***password*]@]*server*/*file-name*** (1-255 characters) is the syntax for uploading the startup-config:

***user*[:*password*]@** (optional) are the credentials for FTP access. If you omit them altogether, they default to the credentials set by the ip ftp-user command. If you enter the *user* but omit the *password*, the CLI prompts for the password before continuing.

server is the IP address or hostname for the FTP server.

file-name is the destination-file path. Lead with an extra slash ("/") if the path is absolute (for example, "ftp://10.1.1.5**//**var/cfgs/ac-gbl.cfg"). Use only one slash if the path is local to the home directory for ***user*** (for example, "ftp://10.1.1.5**/**ac-gbl.cfg").

Syntax: SCP Upload

```
copy global-config scp://user@server:dest-file
[accept-host-key]
```

scp://*user*@*server*:*dest-file* (1-255 characters) is the URL for the destination file:

***user*@** is the username to present to the other end of the SCP connection. This user must be valid at the remote host.

***server*:** is the IP address or hostname for SCP host. End with a colon (:).

dest-file is the destination-file path. Lead with a slash ("/") if the path is absolute (for example, "scp://root@10.1.1.5:**/**var/configs/gblConf"). Use no slash if the path is local to the home directory for ***user*** (for example, "scp://root@10.1.1.5**:**myGlobal.cfg").

accept-host-key (optional) indicates that if the other end of the connection has an unknown SSH host key (that is, if it is new, or if its key has changed since the last time the host was contacted), the ARX should accept the new host key and continue with the upload. Otherwise, the ARX stops the upload if the host presents an unknown key.

Syntax: Sending to an ARX Volume

`copy global-config {nfs|cifs} namespace vol file-name`

nfs | cifs is a required choice. This chooses the protocol for the file transfer.

namespace (1-30 characters) identifies the namespace to hold the global-config file.

vol (1-1024 characters) is the volume name.

file-name (1-255 characters) is the path to the config file, starting at the root of the volume.

Syntax: Sending to Another ARX

`copy global-config ron remote-arx {configs|scripts} file-name`

remote-arx (1-30 characters) is the hostname of the remote ARX to which you are sending the copy. This must be an ARX in the same Resilient-Overlay Network (RON); use show ron for a full list of hosts on the RON. The remote host checks your administrative credentials to confirm you have permission to copy the configuration file: the credentials are the ip ron-user, or the credentials you used to log into the CLI if no RON user is set.

configs | scripts is the destination directory at the remote host. This is a required choice.

file-name (1-255 characters) is a destination-file name, for example, gblcfgbeta01.cfg, to which you copy/save the switch's global-configuration information.

Syntax: TFTP Upload

`copy global-config tftp://server/file-name`

tftp://*server/file-name* (1-255 characters) is the URL for the global-config file:

server is the machine name (for example, "mymachine.myco.com").

file-name is the desired path name for the file. Lead with an extra slash (/) if the path is absolute: for example, "...myserver//var/configs/gblConf" specifies /var/configs/gblConf on "myserver." Use only one slash if the path is local to the server's "tftpboot" directory. This conforms with the specification for FTP URLs in RFC 1738.

Syntax: Send in an Email

`copy global-config smtp://[email-address/]file-name`

smtp://[*email-address/*]*file-name* (1-255 characters) is an email destination for the file:

smtp:// is required. This declares that the destination is an email address.

email-address (optional) is the recipient of the email in *username@host* format (for example, "jsmith@myco.com"). If you omit this, the CLI uses the default address set by the cfg-smtp to command.

file-name is the name of the global-config file. This is sent as an attachment to the outbound email message.

You must configure SMTP on the ARX before you use this syntax. Start with the smtp command in cfg mode.

Syntax: Local Copy

`copy global-config {configs | diag-info | scripts}` ***`file-name`***

configs | diag-info | scripts is the destination directory. This is a required choice.

file-name (1-255 characters) is a destination-file name, for example, gblcfgbeta01.cfg, to which you copy/save the switch's global-configuration information.

Default(s)

None

Guidelines

If you upload using SCP (a secure protocol), the CLI asks for a password. Provide a password for ***user***. The user name and password must be valid at the SCP server. The CLI also prompts for a password if you use FTP and enter a *user* without a *password*.

Use this command to save the current global config to a file. The global-config contains all parameters that you set in gbl mode, such as namespaces, filers, and global servers. These parameters are shared between both peers in a redundant pair. You should make a copy of the global config after any change to these parameters. To save the global-config on a regular schedule, you can use this command with the at command.

You can use this file later to recreate the global configuration, if necessary; copy it to the scripts or config directory and use the run scripts *config-file-name* command.

Note

A global-config run changes the global config for both members of a redundant pair.

Use show directories to view all the local-directory listings.

To save all the local-config parameters, such as layer-2 and layer-3, use the copy running-config command. To save both running- and global-config into a single file, use copy startup-config.

To view the global config, use show global-config.

Note that this command, copy global-config, immediately creates the global-config file from the database and then copies it, so its date stamp is either the time you invoked the command or the time that the file was created at the remote server. This is different from copy configs boot-config, which copies a file that was created when the system was installed: the boot-config file may have a time stamp with the system-install time.

Samples

```
bstnA# copy global-config scripts ron.cfg
```

copies the global config into a new file, "ron.cfg," in the scripts directory.

```
bstnA# copy global-config
ftp://ftpuser:ftpuser@10.1.1.19/acopia-gbl.cfg
```

copies the global config to an FTP server at 10.1.1.19.

```
bstnA# copy global-config scp://juser@10.1.1.35:acopia-gbl.cfg
Password: jpasswd
```

copies the running config to an SCP server at 10.1.1.35.

Related Commands

copy ftp
copy scp
copy {nfslcifs}
copy ron
copy tftp
copy smtp
run
ip ftp-user
ip ron-user
smtp
show directories
copy running-config
copy startup-config
show global-config
at

copy running-config

Purpose The *running-config* is the parameters that apply to the local switch only, as opposed to a redundant peer. We strongly recommend saving the running-config after any configuration change. Use the copy running-config command to copy the switch's current running configuration to a local file, upload it to a remote server, place it into an ARX volume, or send it in an email message.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax: FTP Upload `copy running-config ftp://[user[:password]@]server/file-name`

ftp://[*user*[:***password*]@]*server*/*file-name*** (1-255 characters) is the syntax for uploading the startup-config:

***user*[:*password*]@** (optional) are the credentials for FTP access. If you omit them altogether, they default to the credentials set by the ip ftp-user command. If you enter the *user* but omit the *password*, the CLI prompts for the password before continuing.

server is the IP address or hostname for the FTP server.

file-name is the destination-file path. Lead with an extra slash ("/") if the path is absolute (for example, "ftp://10.1.1.5**//**var/cfgs/ac-run.cfg"). Use only one slash if the path is local to the home directory for ***user*** (for example, "ftp://10.1.1.5**/**ac-run.cfg").

Syntax: SCP Upload `copy running-config scp://user@server:dest-file [accept-host-key]`

scp://*user*@*server*:*dest-file* (1-255 characters) is the URL for the destination file:

***user*@** is the username to present to the other end of the SCP connection. This user must be valid at the remote host.

***server*:** is the IP address or hostname for SCP host. End with a colon (:).

dest-file is the destination-file path. Lead with a slash ("/") if the path is absolute (for example, "scp://root@10.1.1.5:**/**var/configs/runConf"). Use no slash if the path is local to the home directory for ***user*** (for example, "scp://root@10.1.1.5**:**myrun.cfg").

accept-host-key (optional) indicates that if the other end of the connection has an unknown SSH host key (that is, if it is new, or if its key has changed since the last time the host was contacted), the ARX should accept the new host key and continue with the upload. Otherwise, the ARX stops the upload if the host presents an unknown key.

Syntax: Sending to an ARX Volume

`copy running-config {nfs|cifs} namespace vol file-name`

nfs | cifs is a required choice. This chooses the protocol for the file transfer.

namespace (1-30 characters) identifies the namespace to hold the running-config file.

vol (1-1024 characters) is the volume name.

file-name (1-255 characters) is the path to the config file, starting at the root of the volume.

Syntax: Sending to Another ARX

`copy running-config ron remote-arx {configs|scripts} file-name`

remote-arx (1-30 characters) is the hostname of the remote ARX to which you are sending the copy. This must be an ARX in the same Resilient-Overlay Network (RON); use show ron for a full list of hosts on the RON. The remote host checks your administrative credentials to confirm you have permission to copy the configuration file: the credentials are the ip ron-user, or the credentials you used to log into the CLI if no RON user is set.

configs | scripts is the destination directory at the remote host. This is a required choice.

file-name (1-255 characters) is a destination-file name, for example, rnngcfgPrtlnd.cfg, to which you copy/save the switch's running-configuration information.

Syntax: TFTP Upload

`copy running-config tftp://server/file-name`

tftp://*server/file-name* (1-255 characters) is the URL for the running-config file:

> ***server*** is the machine name (for example, "mymachine.myco.com").
>
> ***file-name*** is the desired path name for the file. Lead with an extra slash (/) if the path is absolute: for example, "...myserver//var/configs/runConf" specifies /var/configs/runConf on "myserver." Use only one slash if the path is local to the server's "tftpboot" directory. This conforms with the specification for FTP URLs in RFC 1738.

Syntax: Send in an Email

`copy running-config smtp://[e-mail-address/]file-name`

smtp://[*e-mail-address/*]*file-name* (1-255 characters) is an E-mail destination for the file:

> **smtp://** is required. This declares that the destination is an E-mail address.
>
> ***e-mail-address*** (optional) is the recipient of the E-mail in *username@host* format (for example, "jsmith@myco.com"). If you omit this, the CLI uses the default address set by the cfg-smtp to command.
>
> ***file-name*** is the name of the running-config file. This is sent as an attachment to the outbound E-mail message.

You must configure SMTP on the ARX before you use this syntax. Start with the smtp command in cfg mode.

Syntax: Local Copy `copy running-config {configs | diag-info | scripts} file-name`

configs | diag-info | scripts is the destination directory. This is a required choice.

file-name (1-255 characters) is a file name (for example, arxbeta02.cfg) to which you copy the switch's running-config information.

Default(s) None

Guidelines If you upload using SCP (a secure protocol), the CLI asks for a password. Provide a password for ***user***. The user name and password must be valid at the SCP server. The CLI also prompts for a password if you use FTP and enter a *user* without a *password*.

Use this command to save the current running configuration to a file in the event you need to recreate the local configuration later. The running-config contains all of the parameters that you create under cfg mode, such as layer-2 and layer-3 configuration. These parameters apply only to the current switch, as opposed to a redundant peer. You should save a copy of the running config after any configuration change. To save the running-config on a regular schedule, you can use this command with the at command.

To copy global parameters to a file, such as namespaces and global servers, use the copy global-config command. To save both running- and global-config into a single file, use copy startup-config.

The show running-config command shows the config at the CLI prompt.

To recreate the running config (either on the same switch or a replacement switch), copy it to the scripts or config directory and use the run scripts config-file-name command.

Note that this command, copy running-config, immediately creates the running-config file from the database and then copies it, so its date stamp is either the time you invoked the command or the time that the file was created at the remote server. This is different from copy configs boot-config, which copies a file that was created when the system was installed: the boot-config file may have a time stamp with the system-install time.

Samples

```
bstnA# copy running-config config aco-run.cfg
```

copies the running config into the new "aco-run.cfg" file in the config directory.

```
bstnA# copy running-config ftp://10.1.1.19/acopia-run.cfg
```

copies the running config to an FTP server at 10.1.1.19.

```
bstnA# copy running-config smtp://jsmith@myco.com/ARX-run.cfg
```

sends the running config in an E-mail to "jsmith@myco.com."

```
bstnA# copy running-config scp://juser@10.1.1.35:acopia-run.cfg
Password: jpasswd
```

copies the running config to an SCP server at 10.1.1.35.

Related Commands

copy ftp
copy scp
copy {nfs|cifs}
copy ron
copy tftp
copy smtp
run
ip ftp-user
ip ron-user
smtp
show configs
copy global-config
copy startup-config
show running-config
at

copy startup-config

Purpose Use the copy startup-config command to copy the switch's full configuration (both running-config and global-config) to a specified directory and file. You can also use this command to upload the file to a remote server or send it in an E-mail message.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax: FTP Upload

```
copy startup-config ftp://[user[:password]@]server/file-name
```

ftp://[*user*[:***password*]@]*server*/*file-name*** (1-255 characters) is the syntax for uploading the startup-config:

***user*[:*password*]@** (optional) are the credentials for FTP access. If you omit them altogether, they default to the credentials set by the ip ftp-user command. If you enter the *user* but omit the *password*, the CLI prompts for the password before continuing.

server is the IP address or hostname for the FTP server.

file-name is the destination-file path. Lead with an extra slash ("/") if the path is absolute (for example, "ftp://10.1.1.5**//**var/cfgs/ac-start.cfg"). Use only one slash if the path is local to the home directory for ***user*** (for example, "ftp://10.1.1.5**/**ac-start.cfg").

Syntax: SCP Upload

```
copy startup-config scp://user@server:dest-file
[accept-host-key]
```

scp://*user*@*server*:*dest-file* (1-255 characters) is the URL for the destination file:

***user*@** is the username to present to the other end of the SCP connection. This user must be valid at the remote host.

***server*:** is the IP address or hostname for SCP host. End with a colon (:).

dest-file is the destination-file path. Lead with a slash ("/") if the path is absolute (for example, "scp://root@10.1.1.5:**/**var/configs/startConf"). Use no slash if the path is local to the home directory for ***user*** (for example, "scp://root@10.1.1.5**:**mystart.cfg").

accept-host-key (optional) indicates that if the other end of the connection has an unknown SSH host key (that is, if it is new, or if its key has changed since the last time the host was contacted), the ARX should accept the new host key and continue with the upload. Otherwise, the ARX stops the upload if the host presents an unknown key.

Syntax: Sending to an ARX Volume

```
copy startup-config {nfs|cifs} namespace vol file-name
```

nfs | cifs is a required choice. This chooses the protocol for the file transfer.

namespace (1-30 characters) identifies the namespace to hold the startup-config file.

vol (1-1024 characters) is the volume name.

file-name (1-255 characters) is the path to the config file, starting at the root of the volume.

Syntax: Sending to Another ARX

`copy startup-config ron` ***`remote-arx`*** `{configs|scripts}` ***`file-name`***

remote-arx (1-30 characters) is the hostname of the remote ARX to which you are sending the copy. This must be an ARX in the same Resilient-Overlay Network (RON); use show ron for a full list of hosts on the RON. The remote host checks your administrative credentials to confirm you have permission to copy the configuration file: the credentials are the ip ron-user, or the credentials you used to log into the CLI if no RON user is set.

configs | scripts is the destination directory at the remote host. This is a required choice.

file-name (1-255 characters) is a destination-file name, for example, strtUpBstn.cfg, to which you copy/save the switch's startup-configuration information.

Syntax: TFTP Upload

`copy startup-config tftp://`***`server/file-name`***

tftp://*server/file*.tgz (1-255 characters) is the URL for the startup-config file:

server is the machine name (for example, "mymachine.myco.com").

file-name is the desired path name for the file. Lead with an extra slash (/) if the path is absolute: for example, "...myserver//var/configs/startConf" specifies /var/configs/startConf on "myserver." Use only one slash if the path is local to the server's "tftpboot" directory. This conforms with the specification for FTP URLs in RFC 1738.

Syntax: Send in an E-Mail

`copy startup-config smtp://[`***`e-mail-address/`***`]`***`file-name`***

smtp://[*e-mail-address/*]*file-name* (1-255 characters) is an E-mail destination for the file:

smtp:// is required. This declares that the destination is an E-mail address.

e-mail-address (optional) is the recipient of the E-mail in *username@host* format (for example, "jsmith@myco.com"). If you omit this, the CLI uses the default address set by the cfg-smtp to command.

file-name is the name of the startup-config file. This is sent as an attachment to the outbound E-mail message.

You must configure SMTP on the ARX before you use this syntax. Start with the smtp command in cfg mode.

Syntax: Local Copy

`copy startup-config {configs | diag-info | scripts}` ***`file-name`***

configs | diag-info | scripts selects a local destination directory. You can run the file as a script later if you save it to the configs or scripts directory.

file-name (1-255 characters) is a local-file name (for example, arxbeta02.cfg) to which you copy the switch's startup-config information.

Default(s) None

Guidelines

If you upload using SCP (a secure protocol), the CLI asks for a password. Provide a password for ***user***. The user name and password must be valid at the SCP server. The CLI also prompts for a password if you use FTP and enter a *user* without a *password*.

Use this command to save both the running-config and the global config to a file in the event you need to recreate the configuration. The running-config has all of the configuration parameters that apply to the current switch, and the global config has all the shared parameters for the redundant pair. To copy these configs individually, use copy running-config and/or copy global-config. After any configuration change, we recommend saving a copy of the changed configuration to an external server.

You can use this file later to recreate the full configuration, if necessary; copy it to the scripts directory and use the run scripts config-file-name command.

◆ Note

A startup-config run changes the global config for both members of a redundant pair.

To erase all configuration parameters and restart the configuration, you can delete the startup-config file (delete startup-config) and reboot (reload). This brings you back to the initial-boot script.

Note that this command, copy startup-config, immediately creates the startup-config file from the database and then copies it, so its date stamp is either the time you invoked the command or the time that the file was created at the remote server. This is different from copy configs boot-config, which copies a file that was created when the system was installed: the boot-config file may have a time stamp with the system-install time.

Samples

```
bstnA# copy startup-config config aco-startup.cfg
```

copies the startup config into the config directory.

```
bstnA# copy startup-config ftp://10.1.1.19/aco-start.cfg
```

copies the startup config to an FTP server at 10.1.1.19.

```
bstnA# copy startup-config scp://juser@10.1.1.19:aco-start.cfg
Password: jpasswd
```

copies the startup config to an SCP server at 10.1.1.19.

Related Commands

copy ftp
copy scp
copy {nfs|cifs}
copy ron
copy tftp
copy smtp
copy running-config
copy global-config
run
ip ftp-user
ip ron-user
smtp
delete startup-config
show configs
show running-config

delete startup-config

Purpose

The startup-config contains all parameters in both the running config and the global config. When the switch starts up, it runs all of these CLI commands. Without this file, the switch has no network configuration (other than the OOB management interface) or global configuration (namespace, global servers, and so on). The only administrative users that remain are those with crypto-officer roles.

Use this command to delete the startup-config file.

Mode

priv-exec

Security Role(s)

network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
delete startup-config
```

Default(s)

None

Guidelines

Before you delete the configuration, we recommend saving a copy with the copy startup-config command.

The CLI prompts for confirmation before deleting the startup-config file. Enter **yes** to continue.

After you delete the startup-config file, the chassis comes up from its next reboot with minimal configuration. You can use the reload command to invoke a reboot. Before the next reboot, you can use restore startup-config to recover a mistakenly-deleted startup-config.

Note that this deletes most of your network parameters, so you will only be able to log in through the OOB management port after reboot. To clear only namespaces, global servers, and other global parameters, you can use clear global-config instead.

To remove the remaining parameters and go back to the initial-boot script after the next reboot, you can also use delete configs boot-config. This eliminates the OOB management port, too, requiring you to log in through the Console port after reboot.

Sample

```
bstnA# delete startup-config
```

Related Commands

restore startup-config
copy ftp
copy scp
copy {nfs|cifs}
copy tftp
copy startup-config
show configs
show running-config
clear global-config
delete **configs boot-config**

restore startup-config

Purpose The startup-config contains all parameters in both the running config and the global config. When the switch starts up, it runs all of these CLI commands. Without this information, the switch has no network configuration (other than the OOB management interface) or global configuration (namespace, global servers, and so on). The only administrative users that remain are those with crypto-officer roles. If the startup config was previously deleted, but no reboot has occurred since that delete, you can use this command to restore it.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `restore startup-config`

Default(s) None

Guidelines The CLI prompts for confirmation before restoring the startup-config file. Enter **yes** to continue.

This recreates the current configuration and expresses it as a series of CLI commands. You can use this command after you previously removed the startup-config with the delete startup-config command. This recovers from a mistakenly-deleted configuration. To confirm that the command was successful, you can use the show configs command and confirm that a new "startup-config" entry is in the listing.

If you run delete startup-config and reload before running this command, the chassis boots up without any configuration parameters. The restore startup-config command is no-longer useful after the chassis reboots.

Sample

```
bstnA# restore startup-config

Restore file 'startup-config' in directory 'configs'? [yes/no] yes
```

Related Commands delete startup-config
show configs
copy startup-config

save boot-config

Purpose The boot-config file contains all of the parameters that are set by the initial-boot script (described in the *Hardware Installation* Guides). If this file is absent when the switch boots, the switch reruns this script. Use the save boot-config command to ensure that the switch keeps its boot-config for the next reboot.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `save boot-config`

Default(s) None

Guidelines After you delete the boot-config file (with delete configs boot-config) and reboot, the Console goes to the initial-boot script (described in the *Hardware Installation* Guides). The CLI starts after you answer all of the questions in the initial-boot script.

This command prevents re-running the initial-boot script. It creates a text file named "boot-config," which you can see with the show configs command. You can use the copy command (copy ftp, copy scp, copy {nfs|cifs}, or copy tftp) to upload it to an external host, or to download a new boot-config script.

Sample

```
bstnA# save boot-config
```

Related Commands delete **configs boot-config**
show configs
copy ftp
copy scp
copy {nfs|cifs}
copy tftp

show global-config

Purpose Use the show global-config command to view an ordered list of CLI commands required to recreate the current global-config state.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show global-config
show global-config [filer|global-server|schedule|security|
  config-replication]
show global-config archive [file-history-archive]
show global-config {nfs|cifs} [fqdn]
show global-config namespace [name [volume]]
```

filer | global-server | policy | security | config-replication are optional choices to focus on one section of the global-config report.

filer specifies external-filer details.

global-server specifies global-server configurations.

schedule specifies policy schedules.

security specifies security information such as management services authentication providers.

config-replication focuses the output on config-replication rules.

archive is an optional choice to focus on the "archive" section of the global-config report. This section has the configuration for every file-history archive on the ARX. Unless you specify a particular *file-history-archive* name, this shows the configurations for all of them.

file-history-archive (optional, 1-64 characters) identifies a particular archive.

nfs | cifs specifies that you want to focus on NFS or CIFS services only. Unless you include an ***fqdn***, this shows all front-end services of the given type.

fqdn (optional, 1-128 characters) identifies a particular NFS or CIFS service (for example, "www.medarcv.com"). For a list of all NFS and CIFS services, use the show virtual service command.

namespace specifies that you want to focus on namespaces only. If you omit the ***name***, all namespaces are shown.

name (optional, 1-30 characters) identifies a particular namespace (for example, "medarcv"). Use show namespace for a list of all namespaces.

volume (optional, 1-1024 characters) identifies a particular namespace volume (for example, "/usr").

Default(s) None

Guidelines

The output appears in table format by sections. The sections correspond to groups of commands for setting up major portions of the global configuration (such as namespaces, global servers, filers, policy, security, and so on).

The *global-config* parameters are shared among both peers in a redundant pair; they do not include parameters that apply to the current switch only. For parameters that apply to the local switch, such as layer 2 and 3 parameters, use show running-config.

After changes to the configuration, we recommend using copy global-config to save this text off to a file.

Later, to recreate a saved global-config, paste the text at the CLI prompt (from priv-exec mode) or copy the file from a remote server and use the run command to run it as a script.

Passwords and private SSH keys appear in encrypted form. Only the current switch (or a replacement for the current switch) can interpret and use this private information.

◆ Note

A global-config run changes the global config for both members of a redundant pair.

Samples

`bstnA> show global-config`

shows the full global-config. Figure 33.1 shows a sample report.

`bstnA> show global-config filer`

shows all filer-configuration settings. See Figure 33.2 on page 33-38 for sample output.

`bstnA> show global-config namespace wwmed`

shows the CLI commands to recreate one namespace, wwmed. See Figure 33.3 on page 33-41 for sample output.

`bstnA> show global-config cifs`

shows all configuration settings for all front-end-CIFS services.

Related Commands

copy ftp
copy scp
copy {nfs|cifs}
copy tftp
copy global-config
show running-config
show virtual service
show namespace

Figure 33.1 *Sample Output: show global-config*

```
bstnA> show global-config
; ARX-4000
; Version 6.04.000.15204 (Apr  9 2013 14:41:03) [nbuilds]
; Database version: 604000.4
; Generated global-config Wed Apr 10 11:28:21 2013
;
terminal character-set unicode-utf-8
```

```
;================================ global ================================
global
  kerberos health-check threshold 3500
  nfs tcp timeout 30
  smb2-allowed

;================================ user ================================
  user adm1 encrypted-password 156AVf1b0MLHivMMYCPe1FJQze+9bidzjC31+kuN13w=
    exit

  user adm12 encrypted-password A2tnjv1b0MLHivMMYCPe1FJQze+9bidzpEBrXxgtlCzzVDd1eUxgwg==
    exit

  user admin encrypted-password 2QG4Mv1b0MLHivMMYCPe1FJQze+9bidzKFv3aUkAsP8=
    exit

  user newadmin encrypted-password tkbg0P1b0MLHivMMYCPe1FJQze+9bidzyAGJvioGXpE0HsjzW6ikTg==
    exit

;================================ group ================================
  group Administrators
    role backup-operator
    role crypto-officer
    role network-engineer
    role operator
    role storage-engineer
    windows-domain MEDARCH.ORG
    exit

  group "Backup Operators"
    role backup-operator
    role operator
    windows-domain MEDARCH.ORG
    exit

  group "Domain Admins"
    role backup-operator
    role crypto-officer
    role network-engineer
    role operator
    role storage-engineer
    windows-domain MEDARCH.ORG
    exit

  group "Domain Users"
    role operator
    windows-domain MEDARCH.ORG
    exit

  group "Enterprise Admins"
    role backup-operator
    role crypto-officer
    role network-engineer
    role operator
    role storage-engineer
    windows-domain MEDARCH.ORG
    exit

  group admins
    role storage-engineer
    user adm1
    user adm12
```

```
    exit

  group crypto-officer
    user admin
    user newadmin
    exit

  group operator
    user admin
    user newadmin
    user adm1
    user adm12
    exit

;============================ radius-server ==============================
  radius-server 192.168.25.201
    exit

  radius-server 192.168.25.207
    auth-port 5555
    retries 4
    timeout 10
    exit

;============================ ntlm-auth-db ===============================
  ntlm-auth-db ntlmMap2
    user lab encrypted-password /VvQwseK8wxgI97UUlDN771uJ3NVUbD2LI926fq0mf2RVO5z
    exit

;========================== active-dir-forest ============================
  active-directory-forest MEDARCH.ORG
    forest-root MEDARCH.ORG 192.168.25.102 preferred
    forest-root MEDARCH.ORG 192.168.25.104 preferred
    forest-root MEDARCH.ORG 192.168.25.111 preferred
    forest-root MEDARCH.ORG 192.168.25.112 preferred
    forest-root MEDARCH.ORG 192.168.25.109 preferred
    forest-root MEDARCH.ORG 192.168.25.110 preferred
    tree-domain bostonmed.org 172.16.74.89 preferred
    tree-domain bostonmed.org 172.16.74.88 preferred
    tree-domain fdtestnet.net 172.16.168.21 preferred
    tree-domain fdtestnet.net 172.16.168.22 preferred
    child-domain bostoncifs.fdtestnet.net 10.19.230.94 preferred
    child-domain bostoncifs.fdtestnet.net 10.19.230.74 preferred
    child-domain bostoncifs.fdtestnet.net 10.19.230.88 preferred
    child-domain westcoast.medarch.org 192.168.202.9 preferred
    child-domain westcoast.medarch.org 192.168.202.10 preferred
    child-domain westcoast.medarch.org 192.168.202.11 preferred
    child-domain westcoast.medarch.org 192.168.202.16 preferred
    child-domain ma.ne.medarch.org 192.168.25.103 preferred
    child-domain ma.ne.medarch.org 192.168.25.105 preferred
    child-domain ne.medarch.org 172.16.124.73 preferred
    child-domain ne.medarch.org 172.16.124.19 preferred
    name-server MEDARCH.ORG 192.168.25.102
    name-server MEDARCH.ORG 192.168.25.104
    exit

  active-directory-forest WELLS.ME.ORG
    forest-root wells.me.org 172.16.108.136
    forest-root wells.me.org 172.16.108.139
    child-domain adk.wells.me.org 172.16.110.8
    child-domain adk.wells.me.org 172.16.110.5
    child-domain york.wells.me.org 172.16.120.22
```

```
    child-domain york.wells.me.org 172.16.120.5
    exit

  active-directory-forest VT.COM
    forest-root VT.COM 172.16.213.79
    forest-root VT.COM 172.16.213.8
    forest-root VT.COM 172.16.213.9
    tree-domain ATLANTIC.ME.ORG 172.16.210.14
    tree-domain ATLANTIC.ME.ORG 172.16.210.7
    child-domain MCNIELS.VT.COM 172.16.240.70
    child-domain MCNIELS.VT.COM 172.16.240.88
    child-domain BSH.ATLANTIC.ME.ORG 172.16.110.11
    child-domain BSH.ATLANTIC.ME.ORG 172.16.210.9
    exit

  active-directory forest-trust MEDARCH.ORG wells.me.org
;============================= proxy-user ==============================
  proxy-user acoProxy1
    description "jq's admin account"
    user jqprivate encrypted-password 2AmT2P1b0MLHivMMYCPe1FJQze+9bidzSaJG6Br+JpnkL1BsDw3YZA==
    windows-domain WWMEDNET.COM
    exit

  proxy-user acoProxy3
    user jqtester encrypted-password ralvnv1b0MLHivMMYCPe1FJQze+9bidzMp4xSVs2FNGf6IMrbZPSyQ==
    windows-domain FDTESTNET.COM pre-win2k-name BOSTONCIFS
    exit

  proxy-user cifs_admin
    user Administrator encrypted-password 2QG4Mv1b0MLHivMMYCPe1FJQze+9bidzKFv3aUkAsP8=
    windows-domain MEDARCH.ORG
    exit

  proxy-user emc_admin
    user nasadmin encrypted-password O6OVHv1b0MLHivMMYCPe1FJQze+9bidz+invUycOkMZb+lxHMLhCvA==
    exit

  proxy-user nas_admin
    user root encrypted-password 2QG4Mv1b0MLHivMMYCPe1FJQze+9bidzKFv3aUkAsP8=
    exit

  proxy-user nas_admin2
    user root encrypted-password 2QG4Mv1b0MLHivMMYCPe1FJQze+9bidzKFv3aUkAsP8=
    exit

  proxy-user ny_admin
    user jqpublic encrypted-password agwwC/1b0MLHivMMYCPe1FJQze+9bidzwGwfqayn5sU=
    windows-domain WELLS.ME.ORG
    exit

  proxy-user acoProxy2
    description "user with backup and admin creds on our servers"
    user jqpublic encrypted-password agwwC/1b0MLHivMMYCPe1FJQze+9bidzwGwfqayn5sU=
    windows-domain MEDARCH.ORG
    exit

;============================= nis-domain ==============================
  nis domain wwmed.com
    ip address 192.168.25.201
    ip address 192.168.25.204
    ip address 192.168.25.205
    exit
```

```
;============================ nfs-access-list ============================
  nfs-access-list eastcoast
    anonymous-gid 100
    anonymous-uid 100
    description "allowable subnets in MA, WELLS, & DC"
    nis domain wwmed.com
    permit 172.16.100.0 255.255.255.0 read-write root squash
    permit 172.16.204.0 255.255.255.0 read-only root allow
    permit 172.16.0.0 255.255.0.0 read-write root squash
    permit netgroup surgeons read-write root allow
    permit netgroup medtechs read-only root squash
    deny 192.168.77.0 255.255.255.0
    deny 192.168.202.0 255.255.255.0
    permit 192.168.98.0 255.255.255.0 read-write root allow
    permit 192.168.0.0 255.255.0.0 read-write root squash
    exit

  nfs-access-list westcoast
    permit 172.209.3.0 255.255.255.0 read-write root squash
    permit 172.214.1.0 255.255.255.0 read-write root squash
    exit

;============================ external-filer ============================
  external-filer das1
    description "financial data (LINUX filer, rack 14)"
    ip address 192.168.25.19
    exit

  external-filer das2
    description "Solaris filer 2 (rack 16)"
    ip address 192.168.25.22
    exit

  external-filer das3
    description "Solaris filer 3 (rack 16)"
    ip address 192.168.25.23
    exit

  external-filer das7
    description "Redhat-LINUX filer 1"
    ip address 192.168.25.24
    exit

  external-filer das8
    description "Redhat-LINUX filer 2"
    ip address 192.168.25.25
    exit

  external-filer fs1
    cifs-port 445
    description "misc patient records (Table 3)"
    ip address 192.168.25.20
    exit

  external-filer fs2
    cifs-port 445
    description "bulk storage server (Table 3)"
    filer-type windows port 80
    ip address 192.168.25.27
    proxy-user cifs_admin
    manage snapshots
```

```
  exit

external-filer fs3
  cifs-port 445
  description "Hematology lab server (Table 8)"
  ip address 192.168.25.28
  exit

external-filer fs4
  cifs-port 445
  description "prescription records (Table 3)"
  ip address 192.168.25.29
  exit

external-filer fs5
  cifs-port 445
  description "docs, invoices, for scanners (Table 7)"
  ip address 192.168.25.71
  exit

external-filer fs6
  cifs-port 445
  description "records of lab animals - lab rats (rack C)"
  ip address 192.168.25.30
  exit

external-filer fs7
  cifs-port 445
  description "lab animal test results - lab rats (rack C)"
  ip address 192.168.25.41
  exit

external-filer nas1
  cifs-port 445
  description "NAS filer 1 (rack 31)"
  filer-type network-appliance management-protocol ssh
  ip address 192.168.25.21
  ip address 192.168.25.61 secondary
  ip address 192.168.25.62 secondary
  nfs tcp connections 16
  proxy-user nas_admin2
  manage snapshots
  exit

external-filer nas10
  cifs-port 445
  description "NAS filer 10 (rack 38)"
  filer-type network-appliance management-protocol rsh
  ip address 192.168.25.49
  proxy-user nas_admin
  manage snapshots
  exit

external-filer nas11
  cifs-port 445
  description "filer 11 (rack 38)"
  filer-type network-appliance management-protocol rsh
  ip address 192.168.25.48
  proxy-user nas_admin
  manage snapshots
  exit
```

```
  external-filer nas2
    description "NAS filer 2 (rack 31)"
    ip address 192.168.25.44
    exit

  external-filer nas3
    description "NAS filer 3 (rack 32)"
    ip address 192.168.25.47
    exit

  external-filer nasE1
    cifs-port 445
    description "NAS filer E1"
    filer-type emc management-protocol ssh
    ignore-name lost+found
    ip address 192.168.25.51
    ip address 192.168.25.52 management
    proxy-user emc_admin
    manage snapshots
    exit

  external-filer smb1
    cifs connection-limit 500
    description "Samba filer"
    ip address 192.168.25.48
    exit

;=============================== schedule ================================
  schedule backupWindow
    description "regular backup times"
    duration 04:00:00
    every 1 days
    start 11/12/2006:13:00:00
    exit

  schedule daily4am
    description "two hours between 4 and 6 AM"
    duration 02:00:00
    every 1 days
    start 09/04/2005:04:00:00
    stop 01/07/2015:04:00:00
    exit

  schedule hourly
    every 1 hours
    start 04/10/2013:10:24:00
    exit

  schedule weekly
    start 05/06/1995:02:00:00
    exit

;============================ global policy ==============================
  policy-filename-fileset website
    path exact /www/xml/
    name regexp not "\.(wmv|avi)$"
    recurse
    exit

  policy-filename-fileset hiddenFiles
    path regexp "/\.[^\.]"
    exit
```

```
  policy-filename-fileset xmlFiles
    name match *.xml
    recurse
    exit

  policy-filename-fileset fm_pdf
    name regexp "\.(fm|pdf)$" ignore-case
    recurse
    exit

  policy-filesize-fileset veryLarge
    select-files larger-than-or-equal-to 5M
    exit

  policy-union-fileset bulky
    from fileset fm_pdf
    from fileset veryLarge
    exit

  policy-age-fileset dayOld
    select-files older-than 1 days
    Last Accessed
    exit

  policy-filename-fileset allDirs
    recurse
    exit

  policy-cifs-attributes-fileset online
    offline not-set
    exit

  policy-intersection-fileset onlineDayOld
    from fileset online
    from fileset dayOld
    exit

  policy-age-fileset modThisMonth
    select-files newer-than 1 months
    exit

  policy-age-fileset modEarlier
    select-files older-than 1 months
    exit

  policy-age-fileset 2mo
    select-files newer-than 2 months
    exit

;============================== win-mgmt-auth ==============================
  windows-mgmt-auth fullAccess
    permit all any
    user juser windows-domain MEDARCH.ORG
    user jquser windows-domain MEDARCH.ORG
    exit

  windows-mgmt-auth readOnly
    permit session monitor
    permit share monitor
    permit snapshot monitor
    user mhoward_md windows-domain MEDARCH.ORG
```

```
    user zmarx_cpa windows-domain MEDARCH.ORG
    user lfine_md windows-domain MEDARCH.ORG
    user choward_md windows-domain MEDARCH.ORG
    exit

  windows-mgmt-auth snapViewers
    permit snapshot monitor
    user juser windows-domain MEDARCH.ORG
    user jquser windows-domain MEDARCH.ORG
    exit

;====================== namespace managed volumes ======================
  namespace medco
    protocol nfs3tcp
    exit

;====================== namespace managed volumes ======================
  namespace wwmed
    protocol nfs3
    protocol nfs3tcp
    description "namespace for World-Wide Medical network"
    volume /acct
      metadata critical
      modify
      policy migrate-method staged
      policy pause backupWindow
      reimport-modify
      reserve files 4000000
      metadata share nas1 nfs3 /vol/vol2/meta1
      share bills
        import sync-attributes
        policy freespace percent 2 resume-migrate 3
        critical
        filer das8 nfs /work1/accting
        enable
        exit

      share bills2
        import sync-attributes
        policy freespace percent 2 resume-migrate 3
        filer das3 nfs /exports/acct2
        enable
        exit

      share budget
        policy freespace percent 2 resume-migrate 3
        filer das1 nfs /exports/budget
        enable
        exit

      share it5
        import sync-attributes
        policy freespace percent 2 resume-migrate 3
        filer das7 nfs /lhome/it5
        enable
        exit

      share-farm fm1
        share bills
        share budget
        share bills2
        balance capacity
```

```
        auto-migrate
        enable
        exit

      place-rule docs2das8
        report docsPlc verbose
        inline report hourly docsPlc verbose
        from fileset bulky
        target share bills
        limit-migrate 50G
        enable
        exit

      volume-group 1
      enable
      exit

    exit

;======================= namespace managed volumes =======================
  namespace medarcv
    protocol cifs
    cifs authentication kerberos
    cifs authentication ntlm
    cifs authentication ntlmv2
    cifs filer-signatures
    proxy-user acoProxy2
    windows-mgmt-auth readOnly
    windows-mgmt-auth fullAccess
    windows-mgmt-auth snapViewers
    sam-reference fs2
    volume /lab_equipment
      cifs access-based-enum auto-enable
      modify
      policy migrate-method staged
      reimport-modify
      reserve files 4000000
      snapshot directory display all-exports
      snapshot privileged-access
      snapshot vss-mode none
      auto sync files
      metadata share nas1 nfs3 /vol/vol2/meta6
      no compressed-files
      named-streams
      no persistent-acls
      no sparse-files
      unicode-on-disk
      share backlots
        import sync-attributes
        policy freespace percent 3 resume-migrate 5
        filer fs2 cifs backlot_records
        enable
        exit

      share equip
        import priority 1
        import skip-managed-check
        import sync-attributes
        policy freespace percent 3 resume-migrate 5
        filer nas10 cifs equipment
        enable
        exit
```

```
share equipSnap
  replica-snap
  filer nas11 cifs equipBkup
  enable
  exit

share leased
  import priority 1
  import sync-attributes
  policy freespace percent 3 resume-migrate 5
  filer nas10 cifs for_lease
  enable
  exit

share leasedSnap
  replica-snap
  filer nas11 cifs leasedBkup
  enable
  exit

share scanners
  cifs access-based-enum exclude
  import sync-attributes
  policy freespace percent 3 resume-migrate 5
  filer fs5 cifs xraysScanners
  enable
  exit

share-farm tier1
  share equip
  share leased
  balance capacity
  auto-migrate
  enable
  exit

share-farm tier2
  share backlots
  share scanners
  balance capacity
  enable
  exit

snapshot rule hourlySnap
  schedule hourly
  report snap_hourly
  enable
  exit

snapshot rule dailySnap
  schedule daily4am
  report snap_daily
  retain 7
  enable
  exit

snapshot replica-snap-rule mirrorSnap
  schedule weekly
  report snap_offsite_weekly
  retain 52
  enable
```

```
    exit

  snapshot rule labArchive
    schedule daily4am
    report FA_lab
    archive fileRecordsMed
    no contents user-data
    contents volume-config metadata
    enable
    exit

  place-rule busy2tier1
    schedule daily4am
    report leTier1 verbose delete-empty
    inline report hourly leTier1 verbose
    from fileset modThisMonth
    target share-farm tier1
    enable
    exit

  place-rule nonbusy2tier2
    schedule daily4am
    report leTier2 verbose delete-empty
    from fileset modEarlier
    target share-farm tier2
    no inline notify
    enable
    exit

  place-rule masterDirs2Tier1
    from fileset allDirs match directories promote-directories
    target share-farm tier1
    enable
    exit

  volume-group 2
  enable
  exit

volume /rcrds
  filer-subshares
  modify
  policy migrate-method staged
  reimport-modify
  reserve files 4000000
  auto sync files
  metadata share nas1 nfs3 /vol/vol2/meta3
  compressed-files
  named-streams
  persistent-acls
  sparse-files
  unicode-on-disk
  no cifs path-cache
  share bulk
    description "new server to hold big files (such as xrays)"
    import sync-attributes
    policy freespace percent 5 resume-migrate 6
    sid-translation
    filer fs2 cifs bulkstorage
    enable
    exit
```

```
      share charts
        description "various medical charts"
        import sync-attributes
        policy freespace percent 5 resume-migrate 6
        filer fs1 cifs histories
        enable
        exit

      share rx
        description "prescriptions since 2002"
        import sync-attributes
        policy freespace percent 5 resume-migrate 6
        critical
        filer fs4 cifs prescriptions
        enable
        exit

      share-farm medFm
        share rx
        share charts
        balance latency
        enable
        exit

      snapshot rule rcrdsArchive
        schedule daily4am
        report FA_rcrds
        archive fileRecordsMed
        no contents user-data
        contents volume-config metadata
        enable
        exit

      place-rule dailyArchive
        schedule hourly
        report daily_archive
        from fileset onlineDayOld
        target share bulk
        no inline notify
        migrate close-file
        enable
        exit

      place-rule masterDirs2Rx
        from fileset allDirs match directories promote-directories
        target share rx
        enable
        exit

      volume-group 2
      enable
      exit

    exit

;====================== namespace managed volumes ======================
  namespace insur
    protocol cifs
    protocol nfs3
    protocol nfs3tcp
    cifs authentication kerberos
    cifs authentication ntlm
```

```
    cifs authentication ntlmv2
    cifs filer-signatures
    description "WW Medical insurance claims and records"
    proxy-user acoProxy2
    windows-mgmt-auth readOnly
    windows-mgmt-auth fullAccess
    windows-mgmt-auth snapViewers
    volume /claims
      cifs oplocks-disable
      freespace calculation dir-master-only
      modify
      policy migrate-method staged
      reimport-modify
      snapshot directory display all-exports
      snapshot privileged-access
      auto sync files
      metadata share nas1 nfs3 /vol/vol2/meta2
      no compressed-files
      named-streams
      persistent-acls
      no sparse-files
      unicode-on-disk
      no auto reserve files
      share shr1-next
        ignore-sid-errors
        import sync-attributes
        sid-translation
        no strict-attribute-consistency
        filer nasE1 nfs /patient_records cifs patient_records
        enable
        exit

      share shr1-old
        freespace apparent-size 2579496960
        import sync-attributes
        sid-translation
        import rename-directories rename-non-mappables
        no strict-attribute-consistency
        filer nas1 nfs /vol/vol2/insurance cifs insurance
        enable
        exit

      filename-fileset images
        path exact /images/
        recurse
        exit

      snapshot rule mpHourlySnap
        schedule hourly
        report mp_snap_hourly
        enable
        exit

      volume-group 10
      enable
      exit

    exit

;======================= namespace managed volumes =======================
  namespace labResearch
    protocol cifs
```

```
cifs authentication kerberos
cifs authentication ntlm
cifs authentication ntlmv2
proxy-user acoProxy2
windows-mgmt-auth readOnly
windows-mgmt-auth fullAccess
sam-reference fs4
volume /labMice
  modify
  reimport-modify
  reserve files 4000000
  auto sync files
  metadata share nas1 nfs3 /vol/vol2/meta8
  compressed-files
  named-streams
  persistent-acls
  sparse-files
  unicode-on-disk
  share Mrecords
    description "records and charts - white mice"
    import sync-attributes
    filer fs4 cifs white_mice
    enable
    exit

  share Mresults
    description "test results - white mice"
    import priority 1
    import sync-attributes
    filer fs3 cifs mice_lab_results
    enable
    exit

  volume-group 3
  enable
  exit

volume /labRats
  modify
  reimport-modify
  reserve files 4000000
  auto sync files
  metadata share nas1 nfs3 /vol/vol2/meta9
  compressed-files
  named-streams
  persistent-acls
  sparse-files
  unicode-on-disk
  share Rrecords
    import sync-attributes
    filer fs7 cifs lrTests
    enable
    exit

  share Rresults
    import priority 1
    import sync-attributes
    filer fs6 cifs labRats
    enable
    exit

  volume-group 3
```

```
      enable
      exit

    exit

;====================== namespace direct volumes ======================
  namespace medco
    protocol nfs3tcp
    volume /vol
      direct
      share corporate
        critical
        filer nas1 nfs /vol/vol2
        attach vol1/notes to notes
        attach vol1/corp to shr
        enable
        exit

      share generic
        filer nas3 nfs /exports
        attach vol2 to data
        enable
        exit

      share sales
        filer nas2 nfs /vol/datavol1/direct
        attach vol3/sales to export
        attach vol3/mtgMinutes to mtgs
        enable
        exit

      volume-group 9
      enable
      exit

    exit

;====================== namespace direct volumes ======================
  namespace medarcv
    protocol cifs
    cifs authentication kerberos
    cifs authentication ntlm
    cifs authentication ntlmv2
    cifs filer-signatures
    proxy-user acoProxy2
    windows-mgmt-auth readOnly
    windows-mgmt-auth fullAccess
    windows-mgmt-auth snapViewers
    sam-reference fs2
    volume /test_results
      cifs oplocks-disable auto
      direct
      compressed-files
      named-streams
      persistent-acls
      sparse-files
      unicode-on-disk
      share 2005_charts
        policy freespace percent 5 resume-migrate 6
        managed-volume medarcv /rcrds
        attach 2005charts to 2005
        enable
```

```
        exit

      share chemistry
        policy freespace percent 5 resume-migrate 6
        filer fs1 cifs chem_results
        attach chemLab to .
        enable
        exit

      share hematology
        policy freespace percent 5 resume-migrate 6
        filer fs3 cifs hematology_results
        attach hematologyLab to .
        enable
        exit

      volume-group 2
      enable
      exit

    exit

;================================ archive ================================
  file-history archive fileRecordsMed
    description "archive share for ARX file histories"
    exit

;============================== global-server ==============================
  global server acopiaFiler
    virtual server bstnA 192.168.25.12 255.255.255.0 vlan 25
      enable
      exit

    enable
    exit

  global server ac1.MEDARCH.ORG
    description "CIFS and NFS server for hospital insurance claims"
    windows-domain MEDARCH.ORG
    active-directory proxy-user acoProxy2
    virtual server bstnA 192.168.25.15 255.255.255.0 vlan 25
      wins 192.168.25.102
      wins-name INSURANCE
      active-directory alias fs1
      active-directory alias fs2
      active-directory alias fs5
      active-directory alias insur
      active-directory alias fs1.MEDARCH.ORG
      active-directory alias fs2.MEDARCH.ORG
      active-directory alias fs5.MEDARCH.ORG
      active-directory alias insur
      active-directory alias insur.MEDARCH.ORG
      wins-alias INSUR_CLAIMS
      enable
      exit

    enable
    exit

  global server labtests.MEDARCH.ORG
    description "CIFS server for lab animal tests"
    windows-domain MEDARCH.ORG
```

```
    active-directory proxy-user acoProxy2
    virtual server bstnA 192.168.25.16 255.255.255.0 vlan 25
      wins 192.168.25.102
      active-directory alias pubs_test_fs3
      active-directory alias pubs_test_fs4
      active-directory alias pubs_test_fs6
      active-directory alias pubs_test_fs7
      active-directory alias pubs_test_fs3.MEDARCH.ORG
      active-directory alias pubs_test_fs4.MEDARCH.ORG
      active-directory alias pubs_test_fs6.MEDARCH.ORG
      active-directory alias pubs_test_fs7.MEDARCH.ORG
      enable
      exit

    enable
    exit

;================================ cifs ================================
  cifs ac1.MEDARCH.ORG
    description "insurance-claim records"
    kerberos-creds ac1.MEDARCH.ORG MEDARCH.ORG FU3pp8uGAkuY0lztnjd7bg== ac1$ HOST/ac1.MEDARCH.ORG
2
    signatures
    dynamic-dns ac1
    dynamic-dns fs1
    dynamic-dns fs2
    dynamic-dns fs5
    dynamic-dns insur
    browsing medarcv
    browsing insur
    export medarcv /rcrds as ARCHIVES description "2 year-old medical records"
    export offline-access ARCHIVES none
    export medarcv /rcrds/2005 as Y2005 filer-subshare
    export medarcv /lab_equipment as labs description "lab equipment"
    export medarcv /rcrds as bulkstorage description "big share, now merged thru ARX"
    export medarcv /lab_equipment as xraysScanners description "scanners and xray machines"
    export medarcv /test_results as chem_results description "chem-lab records"
    export medarcv /rcrds/VIP_wing as CELEBS filer-subshare hidden
    export medarcv /rcrds/2011/mp3downloads as MP3S filer-subshare
    export medarcv /rcrds/2004 as Y2004 filer-subshare
    export medarcv /rcrds/2010 as Y2010 filer-subshare
    export insur /claims as CLAIMS description "insurance claims"
    export insur /claims/specs as SPECS
    export insur /claims/stats as STATS description "claim stats"
    enable
    exit

  cifs labtests.MEDARCH.ORG
    description "tests and records for lab animals"
    kerberos-creds labtests.MEDARCH.ORG MEDARCH.ORG oAsaIYq1YtBCcY70f0n8RA== labtests$
HOST/labtests.MEDARCH.ORG 2
    dynamic-dns pubs_test_labtests
    dynamic-dns pubs_test_fs3
    dynamic-dns pubs_test_fs4
    dynamic-dns pubs_test_fs6
    dynamic-dns pubs_test_fs7
    browsing labResearch
    export labResearch /labRats as RATS description "tests for lab rats"
    export labResearch /labMice as MICE description "tests for lab mice"
    enable
    exit
```

```
;================================= nfs ==================================
  nfs acopiaFiler
    export medco /vol as /vol
    no nlm enable
    enable
    exit

  nfs ac1.MEDARCH.ORG
    description "insurance records for WW Medical"
    export wwmed /acct as /acct access-list eastcoast
    export wwmed /acct/wksheets as /acct/wksheets access-list eastcoast
    export insur /claims as /claims access-list eastcoast
    offline-behavior deny-access
    offline-behavior namespace insur retry
    offline-behavior namespace wwmed volume /acct path wksheets retry
    enable
    exit

;========================== auto-diagnostics ============================
  auto-diagnostics
    additional-command "show global service"
    additional-command "show share status"
    mail-to juser@wwmed.com
    schedule daily4am
    exit

;============================== statsmon ================================
  stats-monitor  ; This is a terminal beta configuration option.
    notify cifs-service
      trap response-time
      trap errors
      exit

    notify nfs-service
      trap response-time
      trap errors
      exit

    notify filer-share cifs
      trap response-time
      trap errors
      exit

    notify filer-share nfs
      trap response-time
      trap errors
      exit

    exit

  exit
```

Figure 33.2 *Sample Output: show global-config filer*

```
bstnA> show global-config filer
;========================== external-filer ============================
external-filer das1
  description "financial data (LINUX filer, rack 14)"
  ip address 192.168.25.19
  exit
```

```
external-filer das2
  description "Solaris filer 2 (rack 16)"
  ip address 192.168.25.22
  exit

external-filer das3
  description "Solaris filer 3 (rack 16)"
  ip address 192.168.25.23
  exit

external-filer das7
  description "Redhat-LINUX filer 1"
  ip address 192.168.25.24
  exit

external-filer das8
  description "Redhat-LINUX filer 2"
  ip address 192.168.25.25
  exit

external-filer fs1
  cifs-port 445
  description "misc patient records (Table 3)"
  ip address 192.168.25.20
  exit

external-filer fs2
  cifs-port 445
  description "bulk storage server (Table 3)"
  filer-type windows port 80
  ip address 192.168.25.27
  proxy-user cifs_admin
  manage snapshots
  exit

external-filer fs3
  cifs-port 445
  description "Hematology lab server (Table 8)"
  ip address 192.168.25.28
  exit

external-filer fs4
  cifs-port 445
  description "prescription records (Table 3)"
  ip address 192.168.25.29
  exit

external-filer fs5
  cifs-port 445
  description "docs, invoices, for scanners (Table 7)"
  ip address 192.168.25.71
  exit

external-filer fs6
  cifs-port 445
  description "records of lab animals - lab rats (rack C)"
  ip address 192.168.25.30
  exit

external-filer fs7
  cifs-port 445
  description "lab animal test results - lab rats (rack C)"
```

```
  ip address 192.168.25.41
  exit

external-filer nas1
  cifs-port 445
  description "NAS filer 1 (rack 31)"
  filer-type network-appliance management-protocol ssh
  ip address 192.168.25.21
  ip address 192.168.25.61 secondary
  ip address 192.168.25.62 secondary
  nfs tcp connections 16
  proxy-user nas_admin2
  manage snapshots
  exit

external-filer nas10
  cifs-port 445
  description "NAS filer 10 (rack 38)"
  filer-type network-appliance management-protocol rsh
  ip address 192.168.25.49
  proxy-user nas_admin
  manage snapshots
  exit

external-filer nas11
  cifs-port 445
  description "filer 11 (rack 38)"
  filer-type network-appliance management-protocol rsh
  ip address 192.168.25.48
  proxy-user nas_admin
  manage snapshots
  exit

external-filer nas2
  description "NAS filer 2 (rack 31)"
  ip address 192.168.25.44
  exit

external-filer nas3
  description "NAS filer 3 (rack 32)"
  ip address 192.168.25.47
  exit

external-filer nasE1
  cifs-port 445
  description "NAS filer E1"
  filer-type emc management-protocol ssh
  ignore-name lost+found
  ip address 192.168.25.51
  ip address 192.168.25.52 management
  proxy-user emc_admin
  manage snapshots
  exit

external-filer smb1
  cifs connection-limit 500
  description "Samba filer"
  ip address 192.168.25.48
  exit
```

Figure 33.3 *Sample Output: show global-config namespace wwmed*

```
bstnA> show global-config namespace wwmed
;======================= namespace managed volumes ======================
namespace wwmed
  protocol nfs3
  protocol nfs3tcp
  description "namespace for World-Wide Medical network"
  volume /acct
    metadata critical
    modify
    policy migrate-method staged
    policy pause backupWindow
    reimport-modify
    reserve files 4000000
    metadata share nas1 nfs3 /vol/vol2/meta1
    share bills
      import sync-attributes
      policy freespace percent 2 resume-migrate 3
      critical
      filer das8 nfs /work1/accting
      enable
      exit

    share bills2
      import sync-attributes
      policy freespace percent 2 resume-migrate 3
      filer das3 nfs /exports/acct2
      enable
      exit

    share budget
      policy freespace percent 2 resume-migrate 3
      filer das1 nfs /exports/budget
      enable
      exit

    share it5
      import sync-attributes
      policy freespace percent 2 resume-migrate 3
      filer das7 nfs /lhome/it5
      enable
      exit

    share-farm fm1
      share bills
      share budget
      share bills2
      balance capacity
      auto-migrate
      enable
      exit

    place-rule docs2das8
      report docsPlc verbose
      inline report hourly docsPlc verbose
      from fileset bulky
      target share bills
      limit-migrate 50G
      enable
      exit

    volume-group 1
```

```
    enable
    exit

  exit
```

show running-config

Purpose The running-config applies to the local switch only; these parameters are not shared with the redundant peer. Use the show running-config command to view an ordered list of CLI commands required to recreate the current running-config state.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show running-config`**

Guidelines The output is interspersed with comments that break it into sections. Each section is labeled in the comment. You can use copy running-config to save this text off to a file, or upload it to an FTP server.

The running-config parameters all apply to the current switch only. For parameters that are shared between redundant peers, such as namespace parameters, use show global-config.

Sample `bstnA> show running-config`

shows the running-config. Figure 33.4 shows a sample report on an ARX-4000, Figure 33.5 on page 33-50 shows a sample report on an ARX-2500, Figure 33.6 on page 33-55 shows an ARX-2000 sample, Figure 33.7 on page 33-60 shows an example from an ARX-500, and Figure 33.8 on page 33-66 shows an example from an ARX-VE.

Related Commands
copy ftp
copy scp
copy {nfs|cifs}
copy tftp
copy running-config
show directories
show chassis
show global-config

Figure 33.4 *Sample Output: show running-config (ARX-4000)*

```
bstnA> show running-config
; ARX-4000
; Version 6.02.000.14335 (Feb 29 2012 20:11:44) [nbuilds]
; Database version: 602000.31
; Generated running-config Fri Mar  2 01:35:21 2012
; System UUID  d9bdece8-9866-11d8-91e3-f48e42637d58
; ip private vlan internal 1010 metalog 1011 subnet 169.254.68.0 255.255.255.0
;
terminal character-set unicode-utf-8
;================================ vlan ================================
config
  vlan 25
    description "personnel dept."
    members 2/5 to 2/6
    exit

  exit
```

```
;============================ config-if-vlan ============================
config
  interface vlan 25
    ip address 192.168.25.5 255.255.255.0
    no shutdown
    exit

  exit

;============================== system ==============================
config
  clock timezone America New_York
  hostname bstnA
  ip domain-list wwmed.com
  ip domain-list MEDARCH.ORG
  ip domain-list bigorg.org
  ip ftp-user juser encrypted-password eX1Jci+dkoPRvYvA021HzPOqUtD0FKF9evZPY/lmFTCCjNunSe2tsA==
  ip name-server 192.168.25.102
  ip ron-user admin encrypted-password euTb0S+dkoPRvYvA021HzPOqUtD0FKF9H0nfBBv/vYo=
  ip scp-user juser encrypted-password eX1Jci+dkoPRvYvA021HzPOqUtD0FKF9evZPY/lmFTCCjNunSe2tsA==
  login-banner post-auth "Running-config last saved 1/7 by J. User"
  nsm warm-restart
  raid rebuild-rate 85
  raid verification-mode manual
  raid verification-rate 50
  snmp-server community public read-only
  snmp-server community private read-write
  snmp-server contact "jpublic, jpublic@mycompawells.me.org"
  snmp-server host 172.16.100.183 public
  snmp-server host 10.1.1.68 public
  snmp-server host 172.16.100.101 public
  snmp-server location "2nd floor lab, row 3, bay 4, shelf 5"
  snmp-server name arx1
  snmp-server trusthost 10.1.1.68
  snmp-server trusthost 172.16.100.183
  ssh-host-key dsa encrypted-hostkey
u9oFngi4v7vM94WNlfxoUKFY23fhmZ9JQuS6aEgd6cZip2E/2pxymSSnGfazI1hhJmiMq0IZPzk8GT21B01kcsXumKYR9kvXcm
RSgz01EezA/5HN3YJXTksdLSFXTHQBmpe+FmGOZScDVIXfvTweqXCEhjmCbqO2z+3Pog8sWlHKMHu06YhlvCat4mRJAOScswlS
1r4ofzBhTU+/JjZ9CGtUWadI0qX8eFq+JLXiBoEIAH8CNGBLvunHK5xjZComS7EP/Foi+21cC9QNe48oIkOZaTm7/VgoqH4om7
bZv0xT/d9WaBHoi5XSPngBt1TszDtNbAWQU6bhron6epwNrsQ1S2kbXgdcb/e5AS+SCpsB+RjBNChG/KfIPC3bzrQ72luiYtJj
wBEe+w1MEIuzw/X6s95k+ITvO85Ull4NS49DOJIac/QVNvIut3lbK+WxKZX1m4TbYgxkYKAjeRuW3+MXMrnpUiwbhwCxduhDBF
k8fnkfSCjFQ0D6X6CNMEz8lfab2MiUUeY1QggCBu7iqRDPo7J6+cRiGpORxBoII9XnBOxYS09cy0oq5kZp5GKcnNqDgTTSuKF2
E6Ztkr3SdRzltLpci7fnTp88JWhk9XDKfE+5VR49Yo1vfo4fZzHyUg2ujUDd/XrenaeSeoABuZJwSW2GnAsh04133T3088Zqas
bKP/4GF8+08xJACOe5gIhSGT2y2u0jmbx/9DhcIhKtrX7mgAy2h4Fyb5Bxy/X0cwoeKXa7r8/RzjHwMDne7DuXPy7oa2Yp8orQ
ABZlhgr5GZY5w9o1ljrFnfPbBzr07jRfJFcjMtuOz2IzrZc0N2yPLxqM49v8q8Tj/ZoMxXaCpc5p2lYq+Km+cQzyYpLvIYo5lt
vLe1vY0IBOc6Gk
  ssh-host-key rsa encrypted-hostkey
u9oFngi4v7sheeATbx/2B1sFeka5+7R2rmsI4ZG504gQX2MkfHCXvPlsavRJb3wH0NfTBOKW4cOSYFMz1xokz3VZyuH37uA6fA
ArnWT4Ks0uy16VFI3P0vRk+aRwowHnhMFWbJr66CMkPBbB1YPL7ceyxMJqpJZboWbOwUT+BWuiObnFJ0xcSxVrHEK9soJWhKIT
RaTmuhIkdlQ0yR0V7NgSRMrUuanvgnm0NIEhfeLmNum2vlUA5LrPUEGM/S8YrmDCDmaXkm3eOc831c++i3p51Hy0p1QE+UpwAn
GjR1V7WE/OKtiAzpzgorIGCmhN5R2/bpS9A9jmDmkjy2H5B7HejgyjOfKmMqFVqeuakX3JmrZcwpqDHyF1xyYqSwCxhqeyJJj3
4up19Xnmx90jD9XED5OjGB18V81HGATGs4AC6YEk+tsF6GTNdsNEY+dhiks2d6FbsqK6kiVvlhxxwQv3XlZINdQLb1csfaWPQd
93fnUFjeDRyWpcaXiPQzS9UPtcBLpUASyR3ow/NDbKw5X2JatugUhRDSoKgHplIxH+RutAsYLpsCs9jLfCN7XCrJhyZ4z7Es/I
kfnc2jZ6cXNedfc/0zQLUCR02ag0jHH05CfiXNdvpnMZnMt+6pQjHaZ0IbM9fANIcf5Tju9czukySHnZCikT1qxnwMxJyLsyyL
yfjEA7YinlgqWXxqKXUX1qbBQvd3qu7EuRyffhm38C0YIJy178vO7WR5Al1ssfytojd2dP4r1XnBZZUJmcIytYsI24R+6z/L5R
dxcrKRQfQdefS3CfJuAMFvbA7v6a8mScmtJ60pzgpJGJttpM0TWcaDyDSdxgBdJibl7Z8bQFh/vyqpFtxZyFCDRYAsA5ypAyj8
U6mKsklQrbCGp7NHOlr6xfDFjGBo4cBh6VHDHqbZ9UeDXbebVHSEPpEUd4pb+HHQkNkPDcN1u+E01P29vnmSQnTm3TpLzo6DLD
gccLUdNDzMroyNfydmBHlkqjR6Yy+RLb6bvN/VF/Q4FSfAutJznnI5IEddLEO/eJc/CgOx+PQVSbh+KdjejGs4hjwfqv5EiI+0
0tfg8caiFLhBkpX0ovgT3RjFu2YAmiXMArC/X/8t/5V3PKgykMEDrMogYbP8/GpuEUMed4JROVpJdSCfwxH/cRprCv3sYVALJd
ffXkkt0OQviJVtWk+/S5amajaW6zq3V8x1I+wRVwgYY9IzHWSf0y1CcpTSAfp2Z1FFB4IF9cXOJvvFfz638wnkscLo8sqpkB+J
AfNXo6GxqSWkcWVZxW4qXz5oJZDLVRK2H5BCdWqWlrU30sXA5kDM8UyyZrek/kBzRprN5TyZmbHFJJPKrjR2Nx/Oevt60bVEzY
```

```
wkr9HWbpx7WqSzEQn+40XzpTreXcL4SzdAWiONJGKqLSFxMx2mVtCZ8qbHjiao5/ABdFSpCmzhW8qiJv5DSiqvPFZ0k/m+B0bP
s+QS8xSwcEO08aOZN2ck/ES2lmoKSoyTeprb/RC3L/g4RitCiOFt1rFIU13EuC+hxF4lIwBJCCoW77gx1Ec58BS+NCxkT1LZaV
xqX/Ko2s9pBztck54fwBEbtrv3zOzlyPEiRgJk6zcfPDdFtsijsNBWJzr8nxjTKZ10mWnQXRl5IUHpwZhFdub1V3jtTKQYriEP
ObocQLd+lBavxyAJsjMpymdyEteuN57vhPnbxxrFmWWqV3LgBUsVsrpcHMfCnKYMAQU/KvKAKIyP4K57dgpDFzxXbk3D98eSFm
4HOzaQqi/VzfKTTVV1dR/Rfj3ckLiVA+WNX8cMm5GZRKEP7GhhTzXKvMFJfV7jPsCEq4jAYY846WtHkfgOv9Ku5SA3E9+1fXmV
MOfugrIDyjKM8/GsQZ4Gs+jdOuM0eUt031179xNh0Psw3X7xk3qmnI5vFaqf6x/G55pI58Gt1CC3d+JxaZYLyCc3onQBlk3gfc
XiuZH5eTF4+HqmKe71lAkjXEogMyo3Sxr3eXp8RKPVl58OlRrf0fNjwsdoTeU9ToIJropaFwP1N+nFxzEBT+DsIiYkDxXd1PKf
28fl+S1KTczc1iTVRDT4XKGOIZaWNFsFUVabcL5qGl8RVUpwbIpwOVqAo71ZQdZzMr7ljoFk/x4vkj1cGJNrrgS4AiS+bY8gX/
jY6axRcPxNhShe3DJOKjeMmk35/jdiJVMtat4rbk00q1nSYdB6dMOvB6OuGR4Y/rz/m17rF5U2JgyTSzRrLZ
  ssh-host-key rsa1 encrypted-hostkey
ufEEbrBkXmixIPlmxQ7i5umlhaKT0EL8QoodX9tvS6bl/NNsX5YmscWLkeDSuw9vX7+MrYp2l1Vz20G/qbKEHAoedUzNE38xkS
OixY9rzQRdKV/Q5ffg2kw2PGUsoQcxY0o5A7wAq1Kf+tvQLhmwpbyoROrcXI01mQisK/eQ13kC7ffTQR8axb9NBlyJRs3CdLeJ
aqDJ/ob3+9v4f0szgmtoeVDgQKrkHRE+QUCn04aqWJGWP0y9JEIyWDHa7lg1i/ocM+1pmAxgHTdlba7hBVyeetxKBAQoUtnFKI
jL5h9zX/2LUhH4xogpRJ4pdHkDOWrhdoN439X8Mbb1/GRBdhfxgojhWlWNTd2cFjJwz3AJxi7kHe78QUXJD9p666Q57QY/d4EF
SLA8Rh1siDhSXYgMN9cYrIORBAkP4ExJxtV71dGED5zTgg7mqdBChe4yZQKGWj548Buj1BnpmKZiZ2DkhjM5LYconG0xFHJcSQ
haf7miYp1D0YIHs13DzfGFxsON13CFYvu/RIyD4Qs94pk62NKeLBArtwIK/PWzpTWSXpbtZd60RtAAM4qRHgeLcFfLYe7xdfaA
gWKY2uwKzRTnM6rQdmjlOv9JbijkVoegaTItTNcvC2QPKXI56pw0zqwzS3/6x0iXucC5GjneYADigpEs+nNzHEgORjj3xV94MO
unbq2l8iAOohNW6JdZRlYmpSt4xQNu07QNjH5hZXdMkScVNsaTLtRY8p7CTenys+Jt2AVVj/cb/T7/t9/pIh/8dlqOIrusvocw
h2nX+rFRsj4J/rdffWzG2fuwZYtv1MOe+py5+Yn3focG8yvdp6l2COgIAU7j3IS7CIup1DYCY8vXRo+qA/Y4T6eQ0xyy8tV4y6
UtyKPOq5ffPw7bFZnKfXM7nBDGGyCBZ2SA9FFxa6nxVCCmSjjs+giamEn7wykyxtQmat9aHI6zXvaxap2ugNvNtOHH8fa43ocl
uMNuB4sOF+SxYasrFK5kaW0YHIvG6651j6xnZdZS/SaN1rEePbg39opEjUrykNYuev1sO7EHfbFOXH+rMH1sBnRp7SG91qfT6r
982GQNedR9PxFOQGuE1uBwBWbGiB01KAOEsXaTTPlWpam+5ZU9+t9NAei3w6y/5Nzua3fLcLbhYbfIA83uqYvFwdNLXHukdupg
1UJRjBxQKBvuCzgAz+Q5BmfeXfJjJ/PvMcbpsELvVxsg9CDABlrsfR3Te+5pK8XtnSkYMNqTuXAIlrxs03Xbt+hRAVKPj3ns1W
ZwG6mu0+oNfCIfiD8ne2yZbElH7w==
  ip proxy-address 192.168.25.31 255.255.255.0 vlan 25 mac 1 processor 2.1
  ip proxy-address 192.168.25.32 255.255.255.0 vlan 25 mac 2 processor 2.2
  ip proxy-address 192.168.25.33 255.255.255.0 vlan 25 mac 3 processor 2.3
  ip proxy-address 192.168.25.34 255.255.255.0 vlan 25 mac 4 processor 2.4
  ip proxy-address 192.168.25.141 255.255.255.0 vlan 25 mac 5 processor 2.5
  ip proxy-address 192.168.25.142 255.255.255.0 vlan 25 mac 6 processor 2.6
  ip proxy-address 192.168.25.143 255.255.255.0 vlan 25 mac 7 processor 2.7
  ip proxy-address 192.168.25.144 255.255.255.0 vlan 25 mac 8 processor 2.8
  ip proxy-address 192.168.25.145 255.255.255.0 vlan 25 mac 9 processor 2.9
  ip proxy-address 192.168.25.146 255.255.255.0 vlan 25 mac 10 processor 2.10
  ip proxy-address 192.168.25.147 255.255.255.0 vlan 25 mac 11 processor 2.11
  ip proxy-address 192.168.25.148 255.255.255.0 vlan 25 mac 12 processor 2.12

;========================== management-access ==========================
  management access console
    authentication primary local
    exit

  management access telnet
    authentication primary active-directory
    authentication secondary radius
    authentication tertiary local
    permit all
    exit

  management access ssh
    authentication primary active-directory
    authentication secondary local
    permit all
    exit

  management access https
    authentication primary active-directory
    authentication secondary local
    permit all
    exit

  management access snmp
```

```
    permit all
    exit

  management access http-api
    authentication primary local
    permit all
    exit

  management access https-api
    authentication primary local
    permit all
    exit

;============================= spanning-tree ==============================
  spanning-tree
    shutdown
    exit

;============================ config-if-mgmt ==============================
  interface mgmt
    ip address 10.1.1.7 255.255.255.0
    no shutdown
    exit

;======================== interface ten-gigabit ========================
  interface ten-gigabit 2/1
    spanning-tree shutdown
    no shutdown
    exit

  interface ten-gigabit 2/2
    spanning-tree shutdown
    no shutdown
    exit

;========================== interface gigabit ==========================
  interface gigabit 2/3
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 2/4
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 2/5
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 2/6
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 2/7
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 2/8
```

```
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 2/9
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 2/10
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 2/11
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 2/12
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 2/13
    spanning-tree shutdown
    exit

  interface gigabit 2/14
    spanning-tree shutdown
    exit

;============================ event-severity ============================
  email-severity auto-reboot level critical
  email-severity cpu-failure level critical
  email-severity server-offline level critical

;============================== ip route ==============================
  ip route 0.0.0.0 0.0.0.0 192.168.25.1
  ip route 192.168.78.0 255.255.255.0 192.168.25.2
  ip route 0.0.0.0 0.0.0.0 10.46.16.218 255

;======================== ip route management ========================
  ip route 0.0.0.0 0.0.0.0 10.1.1.1 mgmt

;============================== cfg-smtp ==============================
  smtp
    from admin@wwmed.com
    mail-server email1.wwmed.com
    maximum age 30
    retry interval 10
    to juser@wwmed.com
    exit

;============================ email-event ============================
  email-event chassis-monitor
    description monitoring-chassis
    mail-to alerts@lime.wwmed.com
    group chassis event disk-failure
    group chassis event temperature-failure
    group chassis event disk-control-failure
    group chassis event disk-control-status
```

```
    group chassis event fan-failure
    group chassis event fan-status
    group chassis event kernel-nmi-error
    group chassis event module-failure
    group chassis event module-status
    group chassis event nvram-battery-degraded
    group chassis event nvram-battery-failure
    group chassis event nvram-ecc-error
    group chassis event system-bus-error
    enable
    exit

  email-event noc3
    description "support team at NOC3"
    mail-to juser@wwmed.com
    mail-to jqpublic@wwmed.com
    group chassis
    group metadata threshold-counter 5
    group metadata event online threshold-counter 2
    group storage event share-online threshold-counter 3
    group storage event share-remove-complete
    group redundancy event ha-pair-qd-offline
    group policy event shadowmetadatasharefreespaceerrorraise threshold-counter 5
    enable
    exit

  email-event tech-support
    description Built-In
    mail-to e-support@wwmed.com
    group chassis
    no enable
    exit

;=========================== ntp server config ===========================
  ntp server 192.168.25.201 version 4

;========================== logging destination ==========================
  logging destination 172.16.202.8

;======================== at/scheduler commands ========================
  at date 03/02/2012 00:50:00 every 5 minutes  do "show sessions" report adminSessions
  at date 03/03/2012 00:00:00 every 1 days  do "copy startup-config
ftp://root:rootpw@172.16.100.183//tmp/acocfg.conf"
  at date 03/02/2012 05:00:00 every 1 days  do "active-directory update forest WELLS.ME.ORG
proxy-user ny_admin"
  at date 03/02/2012 03:45:00 every 1 days  do "active-directory update forest MEDARCH.ORG
proxy-user acoProxy2"
  at date 03/02/2012 00:43:00 every 1 days  do "nis update"
  at date 03/02/2012 03:30:00 every 89 days  do "cifs rekey all"
  at date 03/02/2012 01:16:00 every 1 days  do "copy reports mp_snap*
ftp://ftpuser:ftpuser@172.16.100.183//var/arxSnapRpts/ format xml"
  at date 03/02/2012 00:57:00 every 1 days  do "copy reports snap*
ftp://ftpuser:ftpuser@172.16.100.183//var/arxSnapRpts/ format xml"
;============================ SSL Certificate ============================
  ssl
    cipher ssl-dhe-dss-export-with-des40-cbc-sha
    cipher ssl-dhe-dss-with-3des-ede-cbc-sha
    cipher ssl-dhe-dss-with-des-cbc-sha
    cipher ssl-dhe-rsa-export-with-des40-cbc-sha
    cipher ssl-dhe-rsa-with-3des-ede-cbc-sha
    cipher ssl-dhe-rsa-with-des-cbc-sha
    cipher ssl-rsa-export-with-des40-cbc-sha
```

```
  cipher ssl-rsa-export-with-rc4-40-md5
  cipher ssl-rsa-with-3des-ede-cbc-sha
  cipher ssl-rsa-with-des-cbc-sha
  cipher ssl-rsa-with-rc4-128-md5
  cipher ssl-rsa-with-rc4-128-sha
  cipher tls-dhe-dss-with-aes-128-cbc-sha
  cipher tls-dhe-rsa-with-aes-128-cbc-sha
  cipher tls-empty-renegotiation-info-scsv
  cipher tls-rsa-with-aes-128-cbc-sha
  ssl-key-store /u3+7QAAAAIAAAABAAAAAQAGdG9tY2F0AAABNGeqmKwAAAUCMIIE/jAOBgorBgEE
  ssl-key-store ASoCEQEBBQAEggTq3QIWD1afIUWhI5w8HSsNeu1lMqwwIY4nulAyAZ3bnJt84+L7
  ssl-key-store hhxtWA4qt/WC5ma7t+4D/7XlDSIGiHTwEO7plVkuMrq0lYdmrS33NLckZO/hToDP
  ssl-key-store U+XS2fc2/iPxQN+9CyW5t7LYO7yjaCb+FJfbcXdmZcmIXDwR6OM1dG/kECKvztry
  ssl-key-store E+uWeeuhDo2TTv3ORG106sAX+vTFqaddqmhosh4Y8UV270v4+67KYrXbQDcXlEbB
  ssl-key-store 6l+nnYzTRSDTXQ59HtwddJ6+YOGOAWiO36p0aBquehlQh8KhaENKdvEg7RAAONMD
  ssl-key-store 4zs+8fF3rSdi/SG4MNObPGLlynst/dKXy7rSC0jwNvn7QZ6XnOjl9B8yg1lKwIcJ
  ssl-key-store Et+biIsWf1//uSI8u4OUFE6/79blICQp6egFFWyZQ42LsVNLidzjJYfaChSKyDqh
  ssl-key-store r5ptoGxjBqfp81Ahr6JMEq0r0lhK+KHHBNDAWQv59KwqUD5Z6E6VRHP7Gj5fphJX
  ssl-key-store IRaGMe0PYl9pgzUKBKb+T425HRTI/5uhm/ZQurFOVuwtXobACAyTBgBlmiMsV7rn
  ssl-key-store XrlVul3EJ9dBVlo8JpS2TRXE/7g1BcnOkKfeje6Wr/C+++9i8th3JN8WLGLwdAH+
  ssl-key-store EvjLe0voceKL/z/6Vkk8ytaVLpY0l7q2c70iImGs+ioArbIR5eG23rfbQFbcZyhW
  ssl-key-store mTf/R6vlsMYFBuIooQTbWvoqIA+9mSwjvedwFvGXlnqxULgra1Ys2Ido5tXt4LPr
  ssl-key-store oE/Pw2aFV7Ysa4850wycRboK+vIK4DFSFVR9yNpP4kcOS05JCHt0uzhOqjbQC77Q
  ssl-key-store eVioDLFe9LDgXIBdfX7owU/HsmShfF+SZSIsddAYob1/3bN2BtqGUlUpXivwKUXA
  ssl-key-store 8fZSjBQyvjRbm+cZPfYBS+anxknBoic93bVCX9J5jXirStxLWkqkWBxSPZhTwyRi
  ssl-key-store lGAX32h5aJvWukn9jvM4CTsX9MPa9iRCZVpty7Q60EFgi9YQn1YvWvud1n2gLmkR
  ssl-key-store Z2HCOQ/iqTesQFSa7BL8W+419ILno0++WcTOogzxJAu96di1k3aXU1TXYy8lEdml
  ssl-key-store aOSVAGrLEpD7EnBvLlf6BNakbMRB96K2tg0H3d75nsw/Ay9gMcL7Ltqoh8y9LRkH
  ssl-key-store z76idDynoPzptssroHhCvwhBtVmwHUBEcDN2CCqIhbOCqAR8FXc4CRkDoBbEYanZ
  ssl-key-store S8CIGtNrtwwxhNmSd8sKuY2vnGq1nlJb2+hlRZSFwHzYUX+7HsIjPK+iXDWHlmZc
  ssl-key-store U2FMAJydxRp3aYSgzxUIO7mgt9iBmuUgQgTE65WsRWka9SAj9jk3gCaOc05S6aTd
  ssl-key-store M3kE3jphQhrGJ1PrKwJQPKFIjDpY13kjxDbY8tOXLJEJ1jOyf/1VLPob19bmXlPJ
  ssl-key-store exc4eRjKC6rcbkS58Iwq7wB2O8PhZtxj1rRmVJxDLu7kNnsmqywnL5IZBBwpPlrW
  ssl-key-store AwbMi8Q24G2lkitBrZu1+ev0kQZqfZJOtD3N4QFCMSMikewzUsyBzXYmxjPny4iT
  ssl-key-store AKvxqCf1kUKt/IjSAzwNiN8uPr+IE6W2pCR5dGT/i95F47/wgVWlTK1wniFRM72n
  ssl-key-store YF0Lz526MCvpSMio/C1zwNLDGOskYuQnTS4oK/9KKGetbnSPKmWRus4ql4c8Z8tc
  ssl-key-store tWqr01CdBPOZQKVpW2dAARK3wCX1sQAAAAEABVguNTA5AAADjDCCA4gwggJwoAMC
  ssl-key-store AQICBE7zoGgwDQYJKoZIhvcNAQEFBQAwgYUxCzAJBgNVBAYTAlVTMRYwFAYDVQQI
  ssl-key-store Ew1NYXNzYWNodXNldHRzMQ8wDQYDVQQHEwZMb3dlbGwxFDASBgNVBAoTC0Y1IE5l
  ssl-key-store dHdvcmtzMRcwFQYDVQQLEw5EYXRhIFNvbHV0aW9uczEeMBwGA1UEAxMVbG9jYWxo
  ssl-key-store b3N0LmxvY2FsZG9tYWluMB4XDTExMTIyMjIxMjYwMFoXDTEyMTIxNjIxMjYwMFow
  ssl-key-store gYUxCzAJBgNVBAYTAlVTMRYwFAYDVQQIEw1NYXNzYWNodXNldHRzMQ8wDQYDVQQH
  ssl-key-store EwZMb3dlbGwxFDASBgNVBAoTC0Y1IE5ldHdvcmtzMRcwFQYDVQQLEw5EYXRhIFNv
  ssl-key-store bHV0aW9uczEeMBwGA1UEAxMVbG9jYWxob3N0LmxvY2FsZG9tYWluMIIBIjANBgkq
  ssl-key-store hkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAqIs0o/HvOO39qQekTYTqBu3XKLXQFGI1
  ssl-key-store Qm2oCvpanp/LtAgxUSIDCTH6CZlemtNBG/TOHpYq3VnFYB7MdcQvGJyurtudPHLs
  ssl-key-store XUUS616FDX52URQxg83kDpuTninYJVzh5TVQQVfnhVO0czLHcGZmtKQy28ZgKsFg
  ssl-key-store +H66GHDbaVQImaj0haFNtnoZ0VkPX+RJFwvie/sSsbG46F2tVlTlEf7qEBUk6jxE
  ssl-key-store kZF38yGhwxF/snSsM2tIZJxGe7AWoSTZY+d5ndAUMQ9zJ+6mAdedwNAqUQlS07ga
  ssl-key-store OXCLzOVWju75xR+/VWp31PkUco0G7pckoO/2v+1//2oHoywGLfDzdQIDAQABMA0G
  ssl-key-store CSqGSIb3DQEBBQUAA4IBAQCRQES3yS9t1j0HU+szSluEoYfCEGa0XwsrvuB04Uvl
  ssl-key-store CvCZQRJgDYRsG6qVB2KlKL5ljLg4q+9KnNHDCoXlbvksvgg85+i6JQR+GJ7qSWAu
  ssl-key-store 41FuTciq3WQYkc3a+G7UDL/LCQZbb8yNFi5h/Doj2hUMQfWe7SxSUNbu8BwmosOV
  ssl-key-store ckk8d0PZ3TkFVL0QRuaA9aTtA/OGNLnxXoYA+191kWOw5uuLa1tyUcIXm+BGl7Gw
  ssl-key-store mO1JPvpDACWXKXmQ0kydG9HNpByEXEOHT1Q5XZ9Z3f8pTdpdo8HqwVigalINHJPH
  ssl-key-store LiPb0P9nGZMd8qjFnr9k5kHlBPYVV5rOqw/fYKM38Dm/BvHulvmVy9cofbYxKedP
  ssl-key-store Ko3syUM=
  exit

exit
```

Figure 33.5 *Sample Output: show running-config (ARX-2500)*

```
stoweA# show running-config
; ARX-2500
; Version 6.02.000.14335 (Feb 29 2012 20:11:44) [nbuilds]
; Database version: 602000.31
; Generated running-config Fri Mar  2 03:21:11 2012
; System UUID  05d5a0fa-f2fb-11df-8daf-af50d57e388e
; ip private subnet 169.254.98.0 255.255.255.0
; Resource Profile optimized
;
terminal character-set unicode-utf-8
;================================ vlan ================================
config
  vlan 90
    description "personnel dept."
    members 1/2 to 1/3
    exit

  exit

;============================ config-if-vlan ============================
config
  interface vlan 90
    ip address 192.168.90.26 255.255.255.0
    no shutdown
    ron tunnel toBoston
      heartbeat failure 3
      heartbeat interval 10
      peer address 192.168.25.5
      no shutdown
      exit

    exit

  exit

;============================== system ==============================
config
  clock timezone America New_York
  hostname stoweA
  ip domain-list wwmed.com
  ip domain-list MEDARCH.ORG
  ip domain-list bigorg.org
  ip ftp-user juser encrypted-password eX1Jci+dkoPRvYvA021HzPOqUtD0FKF9evZPY/lmFTCCjNunSe2tsA==
  ip name-server 192.168.90.18
  ip ron-user admin encrypted-password euTb0S+dkoPRvYvA021HzPOqUtD0FKF9H0nfBBv/vYo=
  ip scp-user juser encrypted-password eX1Jci+dkoPRvYvA021HzPOqUtD0FKF9evZPY/lmFTCCjNunSe2tsA==
  login-banner post-auth "Running-config last saved 1/7 by J. User"
  snmp-server community public read-only
  snmp-server community private read-write
  snmp-server contact "jpublic, jpublic@mycompawells.me.org"
  snmp-server host 172.16.100.183 public
  snmp-server host 10.1.1.68 public
  snmp-server host 172.16.100.101 public
  snmp-server location "2nd floor lab, row 3, bay 4, shelf 5"
  snmp-server name arx1
  snmp-server trusthost 10.1.1.68
  snmp-server trusthost 172.16.100.183
  ssh-host-key dsa encrypted-hostkey
u9oFngi4v7vM94WNlfxoUKFY23fhmZ9JQuS6aEgd6cb89iO5/+2ZI5aKIESNjbGvVvuhbPuubiNcCU8YcIN+r2PchCva8yjJYS
PNyd6v4F4FMbkDZ2LlEWSKUis3oKl7Sr8dW5RApWXFQdDHrYuk+Ukf/iahRILxAOeDddrS/2ILqCxxykDTKJbex9vCVp5AEwPx
i4EbnMYib9u2Xey5piqG5B+oc+Ul26YtnyagZyo1riv8ER3eh/umGPebXO9L/b9BQGnhfphpyDErFu9bV9LiH5dLDkX0n+zTiX
```

```
VsHcG53/JS5N4RIATsKtPCwBBg6Qa9YcvdDsE+CzmmAv2dy8r1DPeeiZE3IiClD/GjLb7kJmXfuVqgYUUiEUgWjHT7ZdfJ76VT
Y32AeCUZ18jNmtb9kb0OVdqljDwjm+8seRvnD2nY7LVG4aQV8q0spy1P3D377mcIx7MGYWrl9Bd2j6D1VL3tSmGxsTIWGtFGq8
wL2U+++ldLsGygURe1Pc6xMBDIqvB7hTL6g+P3xByCdb16GXiqaEpHoDnVgVyMq56k/eR3YeWocuTdPR2Xt/oCEz6idF96qJt4
XC0vHfuOVTI6UZkx+LH3pARHlA+CnpsAPWfdeoRQnEXtcX0gC3WdSpv7ZSumC2tNEh0L3o5Nrhz4OuAYe10fThqlrbJXRhlhCe
jpsYFy7glL67+G5m6IyiAsogPnkUlnRUwKt2B0FSayW7M4ktzrVA3g9ncB3F1/bqGoo4moH3sclaisIRgu56oF+qgseBCDE6T3
kmghooYmBl8ddeAZx4YBOjsyTE0dt61xeJZ0B+Pd2oOuMaKxRz7ZoAPuagh6uq3dspZOWL+sKNdg5+vuBoE6cmPMgDMvi7z72K
zrtfnOF5z+xRfD
  ssh-host-key rsa encrypted-hostkey
u9oFngi4v7sheeATbx/2B1sFeka5+7R2rmsI4ZG504jqu6PjOe2iuV0J2GnFsayJ+jOhgGoHX4kt7Grm7pv/MY/xa7nCh+jBOR
ENkHbDSKz2DE9T3U9HPylaTzO1SkKlyB0TqmekcXhd8xemkLpG8JObAOY33ykWHaXf2IBA0v9KTyOkaqgsQgFrfGUbAu2c/so1
Kgj5hE52X/YSicOlcTpCYyvH6//4yF1RMs9TxQhP8e/yHODXpIR1pp/Ny6nU7wZfdYHFUZ/u1AWFCEHJh5Azfxq+bD5ZRATMRs
u5iUsmmmz0jE+UaPIw5sSm2VrWs5T8sB8Q/WzoWYNHptKdmVgtldLec/Qe6yerlRCo+hECDrtxDbro+IhXpttJZszDwNzF0g5m
3ueF4xIg+6RsGpLgJhcp1KkK4Zn9VTe0l/wd7+DpIBtZyaFC6qyIeqnSVzC2l6OGOKbkc1uqq/ADGFNu/zfQYDo++aWVlcOvEC
LDBfhXw48gTL1c2Oi1bS9GSO6CgSa/CMC0Lv1lkZH2t6iMxh1fezZsuqGgw1ts3osYrW7IzSPgftRznPoItcVzFTGP45W37UI0
DZJrEBhjwE4Up57tKObH8Lj+ggdR8F6oydJd57qOGtxQgAsHVFGQ1iGkvBZU+hf4k7kT2TVUC0w1gCpn0hmWNfFPAjdcnymn+G
2PNjGaDKuBG2WZ5H1rUjeJx/8TozWwWfRFGNs38Wj1zJgCtdC7k/Pi+iHWF9lIRmYKTwUReLWrECIg22Ex0hgK9wCpb9aHkgEg
hk4QtmpdtS2iptL0/RTrZHF5F32lHsZYhtxFHC4fQXWHYbCrgwLQAOKKZI1E5ub1UGcToefvHyfAUz2jrCO7C+wfXq3DPYNx0C
AZhBeUh3cgt+8JDx7QLZqltorugAX2eONTsTsDUWG31XEBwJk4qOQxWVVFqv0qG6nTjo8sBTeHaKiwEY9AJ+rfJxvfEM9CsR9c
zeqgyoRayFpUTZQSOzfeMVd5Bh773rJmimdGqylqqq+cAQorMKyUZmHDfC7MRDOR9ECQPXag/aWCnRRw/MduX8nQpn/cQaoF23
DxsU3c0LlugKnHh6KmeVQfj+YwLZ+EgPKKtgwzREiolT8VmOQ/TmXoffMA1j/7IIsygrZkWzC4Ohl+AOSsFuYzphda6vKkRW3i
bPwNt+RMi9573CXnlVpZ/BnFMb+SBcVvs9bDkcKlJgGGPWt/UVb5CH1/dhh8p7heXNkfBuKPFd2k2dkcHyH9/Sfy58ko6UODK4
ng6/ZN8KgGNEWOtVvIWk5PvL2JIkqQr+rg93G6OAfEXTwhSobGc/G5GP4n+TW+GnVICgRj96RRn7bh9xM90KvEWic2OqEF1gAn
zdeA/GBIrhgtCxN9psozb4rnvvcuJaZfrTs/uxEIYHefmb9rGJO01aoAaRNPX3BVlVnDCr2sFJZP157PFdyJQtCcqZyEFBB7jp
fH+c6PZE60A594IuBchZmoK4Asyrb879wRBZoLajJXq/tE97i4r75EzfPYdzRuUbxBE4aZ+XD0tIbqRqYnvhIfuUueWwA2bfGw
/MiuJYNh/TXSTLIAKK15VU+M1CG6Vw7t/DKpQlBRbVgC4k5P90ddJyTpCDrJKhH22cNgXyHs5yFeZGFJIitVzjZ5uAO4eX2xu1
ZQTot4eSjQ5T3wqbdRt/ztL4SwZXdYbeqsfj5Orinso7PtOnSf6PzwV1q8AXhYNUJ5rr556wc1VrJbzF3cNMisXpgn+faPf/4d
T0JB05C0gO4EY5wSS+pk8j3BUnHmVUDXYxKI+s6Vb9pFNcxzS7ZdaQDZHSIZ71SPW6TT7zCQ6rJ6RBV9YNyfly4tQq33giwBQU
uWelknkLbngEu7AYU/UbWd6wSQTihZIVi2YyehIWkG3HQ73+4PrGDktv7f+mQsJ/baBHwHn1ORj9GcqEth+EDHYq+JwTwZRXUY
pFbitN9pD4+YnJ/vuHK37vGWmR8meCOOBbL0Mf6kQRTe5sSEL64PgioGHyUg58i7yOHD46dbW6QX59f+CC/LEYN2YaoJDJA/BC
SoIPAb3/tsNUAJDiTJPsg/M+JQq+DndhzdbvKfCF6vTag9uUVogUYE9O1hxb10l3Ryaojk6Gb6jB1ulb1HCkNVZwRvoxkOxce6
/+XuZ/8InjBIP6Zk3oyb33hVoqqi/pOcVxi6jR76zLyUaXBfmMu6MjkrjM1Xl6BbckeiTkOu+lsLRzM4mZBg
  ssh-host-key rsa1 encrypted-hostkey
ufEEbrBkXmixIPlmxQ7i5umlhaKT0EL8QoodX9tvS6bl/NNsX5YmsYwV36+4pvANhRlEMiolqO/J+vsEJZ4YFwyqEjV0P5nDN5
kme+1NarjEmnjISdaOW3ohsQBl9zSw1dMfj0TGW+uVqkEwLIlOpkw1U9+dom2hE92N1LLNPWdQiVABnyXg3ShJ4XTX3irNHqdM
kwVE5Wo+UHXWlLfCYKoDtV3xQnDqu9igeuc6OYxwqMW/l1uWCcCAmszkxaPTcmwcP3Yf4m/mZDssgdIGbdt4y9fbVYxgUKi051
koUMcUNjC1JPXT8tfCtAMjC05SfXNNKCb8K+SQW8PIvC6kB4jPoKSUcXqtTTZxskIcb5PEW1UcSpJ9fbw719hY+aKPb941PZ8L
wGtESYQbl0ALW9DMrWSBQ7YNXSq8ZJ3HBtho8iRG3jxzhIfdX+UniP1Xh8V2F8l4y1alJ2r2i3BUBebPeJTDGA3nW4YHIyW16p
AIeRB6SKQM1BJxUBP0yE1JOMWhwxcu/lzeICP7pibBQrRBpXdwjI7bp27brqrjtRNLF1bIcbDn3bRDA1CALD+GGjjyx35z2HVD
C3LVOD3iFDpfH0Mp7f5b6paTVGelDV9lmhJt6aUR4xDTxIUj9kac3LOlhy+qV2hM4N6OUI5tB32+rEzRSie2rcKAtGiWlQzSTS
NQ0alJ1regt9jumyiLH4TQ64slrAb/NZffV4RJn6CgHA+8s6rBj7YVfRK2/LNGEGo4BS8sNNGmD79w8RrWN0+z3KQQat5te3aN
jRuflvRvF6XV/mRJ7CAU7PA855OV+rRCvgZsA7h+PlO/CmFJPmCzVkSEcglglphIWohr6+zjbdd952DSYG9Nfj0Xj3/r2EYp/p
sDfO+qLCsOFY+bKXriix5MjO2Ib+r9gUP/+m8hc/L/3nqgiM3GUDN0Xfm2lOWE5m+qCous+hdX8+cKWzOT02RUssu22d5XJ45H
3jnXnUdwBcslAKeNuPzscoALprylLyCF0mP/lUnw66MbsSKwT3pW+xQgynaJdi1fD+y65cy64iMZr5O2bqn2fTSOylffioczgU
z7vhGwLr/kT8H4MN0czO3jOKTPUltS+H0nxz44GvYD/4SKRUEt3oe0R18YJUrAPrCM9MWCtUTQd9+AIIsmniTocXtuaW9qVO00
mOdYqwUz13WzN8dnYXvX7enYlSGt91HpwoMdpe5TRGK7f4fMmDvh55ynN81Nop7H0Y5CZj/WxhGEDWEen4mxzvmlh5GNPD/F8f
M86ZTuHikWwXrdxebN81BamENKWQ==
  ip proxy-address 192.168.90.31 255.255.255.0 vlan 90 processor 1.3
  ip proxy-address 192.168.90.32 255.255.255.0 vlan 90 processor 1.4
  ip proxy-address 192.168.90.33 255.255.255.0 vlan 90 processor 1.5
  ip proxy-address 192.168.90.34 255.255.255.0 vlan 90

;============================ management-access ============================
  management access console
    authentication primary local
    exit

  management access telnet
    authentication primary active-directory
    authentication secondary radius
    authentication tertiary local
    permit all
    exit
```

```
  management access ssh
    authentication primary active-directory
    authentication secondary local
    permit all
    exit

  management access https
    authentication primary active-directory
    authentication secondary local
    permit all
    exit

  management access snmp
    permit all
    exit

  management access http-api
    authentication primary local
    permit all
    exit

  management access https-api
    authentication primary local
    permit all
    exit

;============================ config-if-mgmt ============================
  interface mgmt
    ip address 10.1.14.76 255.255.255.0
    no shutdown
    exit

;======================== interface ten-gigabit ========================
  interface ten-gigabit 2/1
    exit

  interface ten-gigabit 2/2
    exit

;========================== interface gigabit ==========================
  interface gigabit 1/2
    no shutdown
    exit

  interface gigabit 1/3
    no shutdown
    exit

  interface gigabit 1/4
    no shutdown
    exit

;============================ event-severity ============================
  email-severity auto-reboot level critical
  email-severity cpu-failure level critical
  email-severity server-offline level critical

;============================== ip route ==============================
  ip route 0.0.0.0 0.0.0.0 192.168.90.1
  ip route 192.168.78.0 255.255.255.0 192.168.90.2
  ip route 0.0.0.0 0.0.0.0 10.46.146.218 255
```

```
;========================== ip route management ==========================
  ip route 0.0.0.0 0.0.0.0 10.1.38.1 mgmt

;============================== cfg-smtp ==============================
  smtp
    from DREW@wwmed.com
    mail-server email1.wwmed.com
    maximum age 30
    retry interval 10
    to juser@wwmed.com
    exit

;============================== email-event ==============================
  email-event chassis-monitor
    description monitoring-chassis
    mail-to alerts@lime.wwmed.com
    group chassis event disk-failure
    group chassis event temperature-failure
    group chassis event disk-control-failure
    group chassis event disk-control-status
    group chassis event fan-failure
    group chassis event fan-status
    group chassis event kernel-nmi-error
    group chassis event module-failure
    group chassis event module-status
    group chassis event nvram-battery-degraded
    group chassis event nvram-battery-failure
    group chassis event nvram-ecc-error
    group chassis event system-bus-error
    enable
    exit

  email-event noc3
    mail-to juser@wwmed.com
    group chassis
    group metadata threshold-counter 5
    group metadata event online threshold-counter 2
    group storage event share-online threshold-counter 3
    group storage event share-remove-complete
    group redundancy event ha-pair-qd-offline
    group policy event shadowmetadatasharefreespaceerrorraise threshold-counter 5
    enable
    exit

  email-event power-supply-monitor
    description "NSB power-supply monitoring"
    mail-to jishan.zhang@wwmed.com
    mail-to howley@wwmed.com
    mail-to m.daniels@wwmed.com
    group chassis event power-failure
    enable
    exit

  email-event tech-support
    description Built-In
    mail-to e-support@wwmed.com
    group chassis
    no enable
    exit

;========================== ntp server config ==========================
```

```
  ntp server 192.168.25.201 version 4
  ntp server 192.168.90.18 version 4

;========================== logging destination ==========================
  logging destination 172.16.202.8

;======================== at/scheduler commands ========================
  at date 03/02/2012 03:29:00 every 5 minutes  do "show sessions" report adminSessions
  at date 03/03/2012 00:00:00 every 1 days  do "copy startup-config
ftp://root:rootpw@172.16.100.183//tmp/acocfg.conf"
  at date 03/02/2012 05:00:00 every 1 days  do "active-directory update forest WELLS.ME.ORG
proxy-user ny_admin"
  at date 03/02/2012 03:45:00 every 1 days  do "active-directory update forest MEDARCH.ORG
proxy-user acoProxy2"
  at date 03/02/2012 03:30:00 every 89 days  do "cifs rekey all"
  at date 03/02/2012 03:27:00 do "show cifs-service user-sessions all" report cifsUsersGgh
;============================ SSL Certificate ============================
  ssl
    cipher ssl-dhe-dss-export-with-des40-cbc-sha
    cipher ssl-dhe-dss-with-3des-ede-cbc-sha
    cipher ssl-dhe-dss-with-des-cbc-sha
    cipher ssl-dhe-rsa-export-with-des40-cbc-sha
    cipher ssl-dhe-rsa-with-3des-ede-cbc-sha
    cipher ssl-dhe-rsa-with-des-cbc-sha
    cipher ssl-rsa-export-with-des40-cbc-sha
    cipher ssl-rsa-export-with-rc4-40-md5
    cipher ssl-rsa-with-3des-ede-cbc-sha
    cipher ssl-rsa-with-des-cbc-sha
    cipher ssl-rsa-with-rc4-128-md5
    cipher ssl-rsa-with-rc4-128-sha
    cipher tls-dhe-dss-with-aes-128-cbc-sha
    cipher tls-dhe-rsa-with-aes-128-cbc-sha
    cipher tls-empty-renegotiation-info-scsv
    cipher tls-rsa-with-aes-128-cbc-sha
    ssl-key-store /u3+7QAAAAIAAAABAAAAAQAGdG9tY2F0AAABNLiJr38AAAUCMIIE/jAOBgorBgEE
    ssl-key-store ASoCEQEBBQAEggTqRf48agFQBaoSoFuujYxV7qTKrWvxWC3FHhfgAKQshmxYvHqT
    ssl-key-store 14z2r+KLpl7D5fOqkE+VeIaV2goXZl9LqNKcXVrF4/0tWa3af3zReFdWsq1I+0pg
    ssl-key-store 0kcNzBRzDFEEY+WVQasjhyCTfv/s/FAsUSR/55veG44p51gA0Nqo4g9ypK0RKh+i
    ssl-key-store plk1NLADyMadhhwgT+dBMtR1OakmjbTfK6ntrpWFe737DI+iKIErvasI1wHbfQ4i
    ssl-key-store y13Kme47X7u2+p05BrfjZofKCqJH6dnVSj9VRqDDOHrNOPQ4jyi/Fqg3yEkKQaOb
    ssl-key-store 9h7Y7ATkFhGoxTdN69VSQtSFXzh7ndQSWO1pWiMcjz59E4rxAysFKuK4Q+uhdpck
    ssl-key-store a/RXVe9lZ4B85MeBH7jI6IxV3LJSZwWxoAU+v1SycDOnKf85OVUdS4wKQujfak/0
    ssl-key-store 3lvJLBHH5sSSljPYjWdSaBYRgRpbw3CTVg0nPfZFzUqU35sfka3GIBvNE28sMliB
    ssl-key-store VpoiVrU15oEC76fJxdzt5flY8PTPnPh/J8w76kslfIEcFr+uKsHzf5K2ddxYHdTt
    ssl-key-store MGsT6t/R36aFPNWTkS742R7w7MxNm5dH28YV5icQJtBABcfYi9i5duVRg1eJ1n0K
    ssl-key-store Obp6xnXgQx14kK3DBsWTCwvDafYamybLUZB+PByuU/Bq20LZqc+uyQwbXeCk+QxV
    ssl-key-store sHUtCabQr6yCGsYwSKXcXVO/+N8X9Q3oVRR7KTrgwHJGMapx3LVjOH/RO5jNGQ0f
    ssl-key-store 6BY8Y32LpXsPpw+g2SeeLwEH4nRi08lKZdoDJFCpu6QqHpjS0M/dSIkpl+7RttjH
    ssl-key-store z0/URx9GFn56lq7mv3/xs+iVWwmb779HtUEvIJDcOd6uJNpSLPX7tFRfxQCYymWG
    ssl-key-store /demhpL9mHwsqTDrXREIzogWUi4jhK4vx6x5lsbe6TsdDyDkYfQWVIpcLDzEyEpV
    ssl-key-store AfzbmocfoWWbgZaup4HgAHk4sVV26gp+hlfFoqAD0aPDGtVGv8z1W4rj9ewXBaBU
    ssl-key-store OKDFOaFSib3/OjzqAL6RltgocyQN5vAIBW/7Wytl+b8zCojYeRG1pYkWqE2LjhKF
    ssl-key-store jUk212lzB+sz2i2sYyu0IfdXCGW3ptibjYghnlPc+Cmm5g9dB5biMSEtxpSZVKur
    ssl-key-store /Eo7AwnV9zzPM1SE6OgUQczkbS2a5zr2BJ/ExzWSm6Bj/WYmrnVIeZae54BBDQkW
    ssl-key-store xy1PWd3KMQrxCwEaYbPUxXdnyqkkF2jM8rDu3FcLnHF+0VYK4wutLAYZ1aNtEk3v
    ssl-key-store ehNuXVA8BgLzByWlb2LlJPcGSt6gNzlXehEn77ECtVzOIPc86ybs+R320jVgj2zH
    ssl-key-store 2E6JU/n1BHH06J8PRoVjjO3fYZBZOECE9Uv3e4HjCctOwzQu9YGE+UMnAFN7hfX7
    ssl-key-store rg7ZyaKvCZspoNZ/HkEImuTmNAVrO5Zd6M6DANQ6ORkseO/dN54ZslnmO1mLMxys
    ssl-key-store 1GNKLwOPcaCDMma8u9hrjfAoNHgOXSM2cytgzaY9aFBXxt4MHkIPp5eMzoThdKuD
    ssl-key-store wErJYxveFxLFWdAa2VQc9B2gBo1SEw0g/coi59PsPtKgs5O3VGJvloC30xUmTFMD
    ssl-key-store 8SJpHbG4G93mSDyE2J4JZXSW5z9bYjY8YoD8RoR8MiUX7W/B4B4q+nuauBW7tNMH
    ssl-key-store CBCBRbOnmLCgWNhepHKje40cFpspOAAAAAEABVguNTA5AAADjDCCA4gwggJwoAMC
```

```
    ssl-key-store AQICBE8IVGUwDQYJKoZIhvcNAQEFBQAwgYUxCzAJBgNVBAYTAlVTMRYwFAYDVQQI
    ssl-key-store Ew1NYXNzYWNodXNldHRzMQ8wDQYDVQQHEwZMb3dlbGwxFDASBgNVBAoTC0Y1IE5l
    ssl-key-store dHdvcmtzMRcwFQYDVQQLEw5EYXRhIFNvbHV0aW9uczEeMBwGA1UEAxMVbG9jYWxo
    ssl-key-store b3N0LmxvY2FsZG9tYWluMB4XDTEyMDEwNzE0MTkxN1oXDTEzMDEwNjE0MTkxN1ow
    ssl-key-store gYUxCzAJBgNVBAYTAlVTMRYwFAYDVQQIEw1NYXNzYWNodXNldHRzMQ8wDQYDVQQH
    ssl-key-store EwZMb3dlbGwxFDASBgNVBAoTC0Y1IE5ldHdvcmtzMRcwFQYDVQQLEw5EYXRhIFNv
    ssl-key-store bHV0aW9uczEeMBwGA1UEAxMVbG9jYWxob3N0LmxvY2FsZG9tYWluMIIBIjANBgkq
    ssl-key-store hkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAiahtKdXZ1iq3ZcJ1++NHr+B7mIQ9+aJZ
    ssl-key-store 2cjcN/pTAph4L2Jm2fd4fuSleL/qLrSCIsUCSynE9808hWRSJxLo6gtNB2sd4bzq
    ssl-key-store quzkXTDL4afeOUqXRd1nqKGx+Zl9yuPO0qSFhb558LKq1zbMezCME9Rv6Vi9iLoc
    ssl-key-store 7x/UMQftl5jY0rUBI+sxnO51m2fdh7c07G+OX07YsTsn5H5DDLdDybbKD9GzcHTw
    ssl-key-store 9hQKz31HBux5WnzmB+JP+2fjLZs2JuW5SOZ885T0NWm3gbiwRIEu+s3c6ciKPS2B
    ssl-key-store DvFb4G0Y/vwNm4v8Guhz6DqqAKU4EroDB7yfE6XoBiTnbHD49qYzswIDAQABMA0G
    ssl-key-store CSqGSIb3DQEBBQUAA4IBAQBkEDAyVFRkhoEOJ5SyJmFRbj5z47ObpP7jRyLvut2e
    ssl-key-store 2r63k+osd0BRtSqS69uX/bppOpmg2izFyGVstb67mxsQk+z85co5LepN84924bdz
    ssl-key-store Fxsn6yewizygKP3jz4U1trQD/S0tRjZmXHfSaBrBUX7K99uSSlThYkNSRtkz3sk3
    ssl-key-store 3vhXXO7Q2nEk5IZRcXebiTYp1n3c7Tq5mSZpAf3cJ8us2H2ti/DS1WQL/n9x8mJz
    ssl-key-store u4kMrsee35gJZnqxGneajl0boJAhREMn6jFrmPbcQDlcC7TtIHWXa9VLMYZDOwGx
    ssl-key-store iBlexYFgEtZu3FvKlKRjoAQZkcEVSu0eCh33NWkxBYi2zCnfxsg7Q9Whkx8RIgmU
    ssl-key-store jGEFvIM=
    exit

  exit
```

Figure 33.6 *Sample Output: show running-config (ARX-2000)*

```
prtlndA> show running-config
; ARX-2000
; Version 6.02.000.14335 (Feb 29 2012 20:11:44) [nbuilds]
; Database version: 602000.31
; Generated running-config Fri Aug 12 02:27:58 2011
; System UUID  876616f6-79ac-11d8-946f-958fcb4e6e35
; ip private vlan internal 1008 metalog 1009 subnet 169.254.58.0 255.255.255.0
;
terminal character-set unicode-utf-8
;================================ vlan ================================
config
  vlan 74
    description rtTestVlan
    members 1/5 to 1/5
    members 1/6 to 1/6
    exit

  exit

;========================== config-if-vlan ============================
config
  interface vlan 74
    ip address 192.168.74.66 255.255.255.0
    no shutdown
    ron tunnel toBoston
      heartbeat failure 3
      heartbeat interval 10
      peer address 192.168.25.5
      no shutdown
      exit

    ron tunnel toMinturn
      heartbeat failure 3
      heartbeat interval 10
      peer address 192.168.81.22
      no shutdown
```

```
      exit

    ron tunnel toGoffstown
      heartbeat failure 3
      heartbeat interval 10
      peer address 192.168.158.147
      no shutdown
      exit

    exit

  exit

;================================ system ================================
config
  clock timezone America New_York
  hostname prtlndA
  ip domain-list wwmed.com
  ip name-server 192.168.25.201
  nsm warm-restart
  snmp-server community public read-only
  snmp-server community private read-write
  snmp-server contact Dude_your_gettin_an_Acopia
  snmp-server host 172.16.100.183 public
  snmp-server location "Training Camp"
  snmp-server name BORGATA
  ssh-host-key dsa encrypted-hostkey
gNRZ3TE6lp97oVRHETyLGlTYotoz1uqVl+h2e1DveFiEV89hLnEPzIIqnzyys6Jnb7vNpoXze4LFUfuRvfjSiZz14NdYc2cNyf
RaBzcUeAm1qjxY2fhz76iUTyeK4Gq4uGH7X/Tjn7jryIFsOUpfDK1zUiVhs93pRySC1UmPYP83Cv4hxd5cAP2qD+zMvG1ODQNe
myj8L+1LIpNHagWQOtYdaBR0iND0SXbdZqdKl2waHroYgfU1iqVAFYSXMO+cXG4Ya156jKbyc/FeikRCpKLOCIq6L09GevX4fo
m+5IKQ+HOpxMF/cD0dBIW0jylTZ8yaDQNxvM8ghcPw9xS2nrMOJc/4EihXi9KUGwpcH7VDlaqo4lCV/6qx/CUhSbqXq0zEJG+s
dNpxguN3LzqdCVia6PEaDxtsJLs9OMdaeFTK1g0bx7ev9lT7aWYOnGgdNeyJu5RV2+I1/dq7Yz6/1/Pe1jPJxpbVWRZWK3ndiY
qDaUsROMsM62dLahuD3djQ+l7IVf32/AZyHMnGQeoELGddj735+aDMHyNIvwXyUQQNiSfnO9T+Ee7yv65BrPkIiVe9OLJXSFQ1
xDhIP0uyoh8V2E7J9h8ShyjDSamTJJX2b30KhIgyxevsZTchFLE8j8budnSMe63ybTY4ftEMyqb+5YBGwpuGrBUoSWY5KCYs16
5r/+F64chD9J8I5FroL8pisaZQ3XIyB8qcorccsdcnTHCy1tK261IfdMI2Cf1+QOxPVQOIB4mv5dfL4FT0AsoP1NJjSrwssj6d
Y7/YUcNFpb5Zh7wRxl2QgaMtYQJbSx6tCs8WTwerg2FO0PIyZdGJIJUIbHdxHMZM/lw6/+lcmsXAqrD7tg9vX/mbTSRJ9zmFJI
0oApws0Se32RPo
  ssh-host-key rsa encrypted-hostkey
gNRZ3TE6lp9ZY4xuVz//ReoemfZlmMf5L6WyVUO2dLQ4uas462uU5K0zQqTcOYzMhMRYVXiJEVJ93ae/1+ykLFwTc7TCi/1NBJ
Gv80bf+ba2eLzl0FSfQ+kzAcfnz961G87WifW3QosII9cJ3cVK6u8WwIM2QkeefWAYmgyviVXBk7IZwrd9nB9HGHPe5BP3vQ++
TCqTRdpvJ2EUCn4yaLFUycDPjBbKhSSHWZ5W/Ba7cEG2bjtxJS0XPk6axPCaEuxpgF8OJLxLlISueXnSZJWIX1z1p/vzMdPLuV
Oded5HSrTIRI5cpeHGZZk1CdrSQJEoJo/9Lj6zZPL/ojrEIdCL+Xw2Z67ECSeNuP9ic0S2SehLIxwEhclvPiyYHudp7CASLhIh
afgVZ4ArHeWU+gwBmXp/uTeKp8jnElmd6WF+7LA0suFYu5ncdDu2JjgRirrWLY4ofC7fbgEOrFS6eawrYly3WyujeSi1TlvP27
oprlOjHNlinZ6tHf/xcI4k7LmzkZjCtIsnsz784CpKw8j7ktVbWMG+AdcytHlf1cJ3mJcfA8FMHvsQUS7vhOcfLuCYvjP1odb8
190Uh938GrHyK67BU8bioh3PF0eMUsbvN8DNEmH9tkyTSkIHpaoDnLKO5PO9+t4EbXSK1i8tkPiAVi0Z8DM2RIpOFOgQdJfYWA
hOp3J5VRCNG6x6AIg7vwZQRgQdZoQVScVVV+3ayikkUCh1ZRSorxQ1vUi6cAsZSXOz+nDLAJnpK5HIJ/wGNvLkY18ZvKXvimYR
B6dfRB3EnWRpkDsXA4aEb6srJlhLfphwJ5tOFwi62c/LEpqc3NUYsE7J/B944MXTSmxddIzYt4GQ9w5ZWfF19ZaHHoGeMTY9iW
MIKn9Z6lSGEnOAlohcSn30zxUfXK6wjPEhIgdRESW0xDdD+6eshbqyQur8TdrIbReUWfgp4CUiFaqclpZeDjP7dJQIUzwApUnV
3yGTPuTUXRalTc5s7dvuPgenKQ5k+GoPlrZRdb1AR9hp6bQHU8Shus1ljcG5U4hcvHQSoBbFd/nyNJK1TlDdDxnJ3WVD7ZV7g6
DFekuKz5fmL2GhKcRDn6OmJxsunsxgzEMr+QOpB2mQJW3V6HQLf46/Rmz1P78cLrQtwozilDwqxDaNV4Qo4pWskGHUKuf8j2eb
50FYqNZ3zFo95uWI8mgqLUwb7ct9bcoAUY+4Kw6zNr++AWOtjcUhk/u93aPHc/5DlAPhPq38RJP7uXajImQ52he007rAMurUP/
TJgCaKrOIFP/PPjV9J7+GWCJgOuh4HjzfHSN8unoFxh/XAvZCvh9W1lABF08tg6kD1tQj9qLtm2rfNTHkfhlaicvgK+kzHdaR6
s6nL73qXLnnrHjrHYF25gvNRR7ViHKSs3u8Lk8qL4ZinHQW0G2+22dcnVJQFW3i0pHzv9lClJ3t5F3RAL5K4H069EytiIpdpAt
Bb2lcgB9hkM6od7ipeCZCCO6aLskU1Lo9xPplj/68Vo17rohReORRsq9XntgC/IRh+PxV2qXNMijbUkaaniSVcMUVgVobhRBHa
K3duIf/SCjm4tE1sJqS5TNllLAuECh8Eqb06Rt7ZbBGqL3gxEsyZ1/vzyJ/PFT3VCiMZHg24bPgbe9o7FtGhYBqjk2zMAP4wQY
NTy9wVlXiU4B7xLbHlpcRpVUaMhjDMWplk/mJOhxy1AYgvhfpFPH+2hhQb37zfPjmvscwV929mky6kXL+yctEBT9VwB4q/EgqU
AqBdFmiwuP5MvsNQ11a9KTwkvrePdx/EaxgJ0BVLSNqGtS/NCwgcF7CKNjNtlXYZ4qOmY5v4CMaJl87sEnl8V+GmlURlS5LLeW
mHLDaeg+vrQ3pP9j1aUCPIT9M0O2kulu3CCDq4ks+4PkpGkotL6fMXi+WD1FSfFricIr5ek9Od/pGgckIuM5RNSkSsVasjgwIl
27TLgRCRu2/TnB4NpAQS99Y4DIplnwy82bMRz9/wNKSm/wTiOKg3/dOW9F5wIhlspPe6/agzn0hqd/jTTEGKtYqWjp0Xv0jJyC
iYh5OuT6oVs0Yo/RxQvvyeS9eyBbuLuSI1DYYGWJLJ0awR5BCGS5TlP9jqKf/6scj/GLLZ9n3+doZOLVHZVfP3JDi+qpNsd5bG
Ep3LVTUQ5DiqoLTLGkDHOx/ZQPcIGdoa2jAbvBpJF4q8HnXDMPdCR9Vtxqvbq+rhAR4dKEf8GJ7pAd0DsrU+
```

```
  ssh-host-key rsa1 encrypted-hostkey
3bssB6t1aQ+hhd/J/koLVHL8ww9iofch87mbK0RnYTxtRSz/OZv4HwpqM3zfrD9pPnLBRpnNYM5ODq2xOtCTy2BO4YOBNSQNnC
Jja1+9AGEk9YDL0DIPQ/WcPSKNcZPJFt869TH71JbXeRYM62WmXG6uJXS+UBgWts4kSEFEp8EYE3Bz6oXAh/g5fhag3W2+IyUJ
QpDYnej77ZKq52p0tggPVC2XvsscCbkc58Ou+SsH++kbo6keI8/QpZcQ880n9h68QipAnr7HAsulEMNLh2HuKOfyX0khdCHGxm
fsMXSD0CMWDxOOa11b7j5eAqHUQE4YX2H0dXmiTqFswE/u+/uDFwlklPOvOCk97ArUEpH/wPpJIylHFfGcT9QHubn1zyEK/tNm
5VBTEF968QjKpvgdQsn3HxyLNF9U6/6wOI3XtTvmUgQSdqiGNRd9NA2hqSnxspBlIMcs+tQ4wd6xk0VvSdM5NuZFq94CdmkRt1
2vG+EjO/JxHGrz4kCYkuOhL0iR1cNkpLw2CfnbYCtV6pGtTQSEnDpDAWf8Z++/Nh/qArLxfsJOe4LpnNS3nzG3sA8zEs7ZVYrB
5RfJONICnq1U9Brbgp6xd6Pk3qESeNnazuhZ3ZXYLbHUtWFRaMoJr10CALpxUaFZEieK48kggywu6pMcbpmRgER5LUA0RKs8zr
NVWzv3msr6N7i4BPn8ANFdMUmxdbx2QyKumovVAm5hRATUYohpFNjkUIDaybPHlQHfP4JcCWliHAQ2HiW2pCwmeuXDffh76RG6
DpLlZhP9gZQC9bnfrXECuzcdGVWBkema9mRNepA4vF3D9uw2BTbaZ316J3/1NsOcrfPcsG8dThPc8YFAF/I+MfreeYjUYlcJr8
n0Fvr18SQ56avuOKMmcMMYM7vUAXtuxllRFpjdKkXu8HIPtYH6ZsKV1EyBEPeKXU5/Tek6eecuh0CMCIVY/qFdrtFJC1SXbs6w
7wl72QD3i0BXI7FcnoNhPC/heiFhstItd2gBFkW1SZz/erbS1raO5M6A9IZetKFeEqKztu48z6pwJJcdeFmxpFjWWZS7/XFfGO
LEXWEO5KHr0Ap0MqfNfSgFv9wzYfCKL5cJTGtyO/v+i0njOw/Bz9m4TUTAgofB5w5mQiS3eqdmUYFPFE+cjdhd6oBOjLsK3W7X
Ws+1c9DBJzh76dUoAHhV8ikp0Q423PSsMHr0+VFExN+B02cEoNJ+4qyYrFLT7ZX5Ko0xodYw2c7VVSwotUT7pcnrCYuAXjoC/F
O/ImcxL//PpgBy5X5DUJ2U2pNb+g==
  ip proxy-address 192.168.74.21 255.255.255.0 vlan 74 mac 1 processor 1.2
  ip proxy-address 192.168.74.22 255.255.255.0 vlan 74 mac 2 processor 1.3
  ip proxy-address 192.168.74.23 255.255.255.0 vlan 74 mac 3 processor 1.4
  ip proxy-address 192.168.74.24 255.255.255.0 vlan 74 mac 4 processor 1.5

;=========================== management-access ===========================
  management access console
    authentication primary local
    exit

  management access ssh
    authentication primary local
    permit all
    exit

  management access https
    authentication primary local
    permit all
    exit

  management access snmp
    permit all
    exit

;============================ spanning-tree ============================
  spanning-tree
    shutdown
    exit

;============================ config-if-mgmt ============================
  interface mgmt
    ip address 10.1.23.11 255.255.255.0
    no shutdown
    exit

;=========================== interface gigabit ===========================
  interface gigabit 1/1
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 1/2
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 1/3
```

```
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 1/4
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 1/5
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 1/6
    spanning-tree shutdown
    no shutdown
    exit

  interface gigabit 1/7
    spanning-tree shutdown
    exit

  interface gigabit 1/8
    spanning-tree shutdown
    exit

  interface gigabit 1/9
    spanning-tree shutdown
    exit

  interface gigabit 1/10
    spanning-tree shutdown
    exit

  interface gigabit 1/11
    spanning-tree shutdown
    exit

  interface gigabit 1/12
    redundancy protocol
    no shutdown
    exit


;------------------------------ ip route --------------------------------
  ip route 0.0.0.0 0.0.0.0 192.168.74.1
  ip route 192.168.78.0 255.255.255.0 192.168.74.2

;------------------------- ip route management -------------------------
  ip route 0.0.0.0 0.0.0.0 10.1.23.1 mgmt

;-------------------------- NSM Configuration --------------------------
  no nsm recovery

;------------------------------ cfg-smtp --------------------------------
  smtp
    from BORGATA@wwmed.com
    mail-server 192.168.96.53
    to alerts@lime.wwmed.com
    exit
```

```
;============================== email-event ==============================
  email-event chassis-monitor
    description monitoring-chassis
    mail-to alerts@lime.wwmed.com
    group chassis event disk-failure
    group chassis event temperature-failure
    group chassis event disk-control-failure
    group chassis event disk-control-status
    group chassis event fan-failure
    group chassis event fan-status
    group chassis event kernel-nmi-error
    group chassis event module-failure
    group chassis event module-status
    group chassis event nvram-battery-degraded
    group chassis event nvram-battery-failure
    group chassis event nvram-ecc-error
    group chassis event system-bus-error
    enable
    exit

  email-event tech-support
    description Built-In
    mail-to e-support@wwmed.com
    group chassis
    no enable
    exit

;=========================== ntp server config ===========================
  ntp server 192.168.74.104 version 4

;============================ cfg-redundancy =============================
  redundancy
    peer 10.1.23.12
    quorum-disk 192.168.74.83:/exports/quorum-disk/portland1 nfs2
    critical route 0.0.0.0 0.0.0.0
    enable
    exit

  exit
;============================ SSL Certificate ============================
  ssl
    cipher ssl-dhe-dss-export-with-des40-cbc-sha
    cipher ssl-dhe-dss-with-3des-ede-cbc-sha
    cipher ssl-dhe-dss-with-des-cbc-sha
    cipher ssl-dhe-rsa-export-with-des40-cbc-sha
    cipher ssl-dhe-rsa-with-3des-ede-cbc-sha
    cipher ssl-dhe-rsa-with-des-cbc-sha
    cipher ssl-rsa-export-with-des40-cbc-sha
    cipher ssl-rsa-export-with-rc4-40-md5
    cipher ssl-rsa-with-3des-ede-cbc-sha
    cipher ssl-rsa-with-des-cbc-sha
    cipher ssl-rsa-with-rc4-128-md5
    cipher ssl-rsa-with-rc4-128-sha
    cipher tls-dhe-dss-with-aes-128-cbc-sha
    cipher tls-dhe-rsa-with-aes-128-cbc-sha
    cipher tls-empty-renegotiation-info-scsv
    cipher tls-rsa-with-aes-128-cbc-sha
    ssl-key-store /u3+7QAAAAIAAAABAAAAAQAGdG9tY2F0AAABNawdrJIAAAUCMIIE/jAOBgorBgEE
    ssl-key-store ASoCEQEBBQAEggTqVatRMfGcc/1i0yqSnJBS4eVRRufRoF3ck7TP71jU+YW3Fz1a
    ssl-key-store zun0+NJJvt3lfe6GtmXJkkS/pQEQWL4laYrN9lv6o8TvJsHYTJ4nDhqsmPqyGt4l
    ssl-key-store FyS7GTTuqtCq9VJVa69MAfjx9s18g+3uD8fsBpmLQDWBBot9R1YQ5FfZqTpaJEdt
    ssl-key-store /+474QLt3kC3tqAcFF/0CJiVeWoCiaZjyf/8IvERnJlmxqT8EunYQS7i6tMHId3+
```

```
    ssl-key-store SSys/LP5RT9awWmDbA7cqnh1z7RXa6EY+RfG4WKmJqSVr+MGSXUWT0BU4nQqWd4K
    ssl-key-store V+jll5SHo+ndLn1e8uDh0IgCHgd/QaFAIZP5kENdwiln2yDffLhE7GB0VEGKHWVo
    ssl-key-store P3ZApdhCXkd7sXPXZbq8l7DzKdtq2+vk8F20/GJm+qeIDGdjcNcYsuo6nAtmT5dp
    ssl-key-store 3+dnXHtn/h6dkJvq9fO+GStsBk0152MfjVjgYsTY1TMpkw1T3OU3aY+O84c5tNUO
    ssl-key-store E4/klAV8UgOZMtr18XC+mZJwPNPmTb6CQuiybntrbfBntsX2WtgOx4mJ/QuEA58R
    ssl-key-store jlTRA1W2H3yBHWSEFcwyWKv3gppXYO7/d2awIbCS+DwNo30Shhs5nTNva+uMDsPx
    ssl-key-store 9WEAbAqkni7EZHa24JFE+uJWkPEpX05mX2M9IcVZJ07pqk0rq0Y8GjTuAaeG/sYM
    ssl-key-store 5LtVLeHnLgt9J37zItNWPiTTTuzlWQsb2lgLzUoNg5AEfm7f0PGOg62lDOByCuWJ
    ssl-key-store nuOwZCgnq616NDqLtyh8JGL9/atYodGHUwmUCDrXx45H4fk/BADZaSMiKfPStVE4
    ssl-key-store h4V0MUGk7X+0Vr2QOQWh3YRj/Mq8P9hJpdYSXf+sE13+MduoJkwxAyW+VoaboDRV
    ssl-key-store J0ioNWB6jhByJyQG0feLOFShlHJI7C9kbjOBP77l3N3TKtFwv/wbHH91dahlYdYr
    ssl-key-store WJpML9CXVNOukscggRYi7Sq9qYcorZtChs7KF55/+tmJxuDHGi1hjm5kSYEs1ib4
    ssl-key-store MPctBJWxaIAPEa5zReg2DHqfOB9/XLGq4kNlLqJpt8031+yRxCcm5JVF98nSEK13
    ssl-key-store 8j979mYf7BXRSfLjgV5GL55rf3ZXowZjGUuyTqzNKudoyUA2yNCFCHEMaBcQDerj
    ssl-key-store jcOQDTMRBVmWG7CkrfwssLityJiRNcBZLEIENdtY8Ze+TlK0jAQiYyNhR32Ahe6i
    ssl-key-store mJYxOL8rlJD4R7n/NbAf7PQ2j+IbMi3a4GPmaTTPiIFBa+B/SR2qYdzZqE42mcBb
    ssl-key-store rHtjbdvJPHH9RbSnnSEUeV+uqgitzQhRjaXaa3DS+lP3fTfR1OK6Ch+bsVbk9tN6
    ssl-key-store tt1aw/syMAv/BXhRNJ7nq/pW+BhLGenCA83E5AXrsf9ZbALTA830swEDgKbyOcum
    ssl-key-store yR9M/QgkzC/5dK27zNpfeiD4Px5w2is6kwP4ifGnm7k02j3a2mlZSBjUVU6Rbx/V
    ssl-key-store aUnEbonCVdj4ep6tf6wulIpbT0avnD47UpuiQNBljcgYx7bMQmkPfeCaryo5rxQw
    ssl-key-store A4daSBUNcIrdvqnQAKzPpeEK68P5KXiqZlOkVhKjRcIsFLingvLoH5D6X7QrMCYi
    ssl-key-store vywJjoZVOJqoGakPJd9JDyrjgUi0QEe+cnSsrlPh5D2INV7+SdfCeZD9M5TYhODV
    ssl-key-store hPDEkNCu3wTUDvdJZHMBvZhqxN4KgQAAAAEABVguNTA5AAADjDCCA4gwggJwoAMC
    ssl-key-store AQICBE9Gr4gwDQYJKoZIhvcNAQEFBQAwgYUxCzAJBgNVBAYTAlVTMRYwFAYDVQQI
    ssl-key-store Ew1NYXNzYWNodXNldHRzMQ8wDQYDVQQHEwZMb3dlbGwxFDASBgNVBAoTC0Y1IE5l
    ssl-key-store dHdvcmtzMRcwFQYDVQQLEw5EYXRhIFNvbHV0aW9uczEeMBwGA1UEAxMVbG9jYWxo
    ssl-key-store b3N0LmxvY2FsZG9tYWluMB4XDTEyMDIyMzIxMjg0MFoXDTEzMDIyMjIxMjg0MFow
    ssl-key-store gYUxCzAJBgNVBAYTAlVTMRYwFAYDVQQIEw1NYXNzYWNodXNldHRzMQ8wDQYDVQQH
    ssl-key-store EwZMb3dlbGwxFDASBgNVBAoTC0Y1IE5ldHdvcmtzMRcwFQYDVQQLEw5EYXRhIFNv
    ssl-key-store bHV0aW9uczEeMBwGA1UEAxMVbG9jYWxob3N0LmxvY2FsZG9tYWluMIIBIjANBgkq
    ssl-key-store hkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAvJ1ZzhoDcs9BHzgyQAwlP4OCe/HjbDL5
    ssl-key-store 1RATH1GiXSHdTZXGmjb0InH0RUhVP4iGO5PvJkySIZX3sqmCEtRkXq4M6lI37RHk
    ssl-key-store EokX0eut0w3xn0ByxNJkiDS+4UC40RHmPDDIrLtfJyuZ3iKKq5qjZKzYiJPoRwk5
    ssl-key-store pykBU+5IG7lbUPTVUF+HU4Ykr56m5p56pCOao6R/nzMVj5R+5S2MCZYvK8ExoGd0
    ssl-key-store 1YydE4xc5GGsrp6pR/19wYRRGQTcrflr1yWqxf/S1iJUFsuGgR35n5nZyVPjCLuC
    ssl-key-store aMJPpu8WHrt5bJ8WNTl9gcrnp8MpqgUQzeR6mmhUhSarxqMxlao1VwIDAQABMA0G
    ssl-key-store CSqGSIb3DQEBBQUAA4IBAQCRUlT+hFbSoe+/5AfUABBf07CA6X5y0GFX5dTGIqU7
    ssl-key-store vXzKOXULJS5B1LIQxzHWMhmqVAgH6/8PU9gwNO79EQU9zjkTbJAukK2Yj5WK3SFX
    ssl-key-store FjIWFGwYANlpPnXvk2yAL9WWcpsnmmZtDOczpgVkG+IXMOY6OFV0nqWdSzb6kAFs
    ssl-key-store cfPMuSHS6kEWcXyq9v8XMpVbhT7kMkRybY6fBwuFDvA9CVKqU58oniBbMT1Ojvio
    ssl-key-store /AECOnjFajgTxSiULtw4UuGz8gBX4OBNjl7LXoIwSdTL/ybZJvIfNkDMQiF5+iXQ
    ssl-key-store s/68ozKbhLYjlZiMs0bHYE9lc0eA844Aunp0GsMCGB895JLa3RajHY5C8glcpYAz
    ssl-key-store /kVkJFU=
    exit

  exit
```

Figure 33.7 *Sample Output: show running-config (ARX-500)*

```
provA> show running-config
; ARX-500
; Version 6.02.000.14335 (Feb 29 2012 20:11:44) [nbuilds]
; Database version: 602000.31
; Generated running-config Fri Mar  2 02:31:21 2012
; System UUID  db922942-876f-11d8-9110-8dtu78fc8329
; ip private vlan internal 1006 metalog 1007 subnet 169.254.127.0 255.255.255.192
;
terminal character-set unicode-utf-8
;================================ vlan ================================
config
  vlan 103
    description "personnel dept."
```

```
    exit

  exit

;============================ config-if-vlan ============================
config
  interface vlan 103
    ip address 192.168.103.160 255.255.255.0
    no shutdown
    ron tunnel toBoston
      heartbeat failure 3
      heartbeat interval 10
      peer address 192.168.25.5
      no shutdown
      exit

    ron tunnel toNewport
      heartbeat failure 3
      heartbeat interval 10
      peer address 192.168.8.106
      no shutdown
      exit

    ron tunnel toGoffstown
      heartbeat failure 3
      heartbeat interval 10
      peer address 192.168.158.147
      no shutdown
      exit

    exit

  exit

;============================== system ==============================
config
  clock timezone America New_York
  hostname provA
  ip domain-list wwmed.com
  ip domain-list MEDARCH.ORG
  ip domain-list bigorg.org
  ip ftp-user juser encrypted-password eX1Jci+dkoPRvYvA021HzPOqUtD0FKF9evZPY/lmFTCCjNunSe2tsA==
  ip name-server 192.168.90.18
  ip ron-user admin encrypted-password euTb0S+dkoPRvYvA021HzPOqUtD0FKF9H0nfBBv/vYo=
  ip scp-user juser encrypted-password eX1Jci+dkoPRvYvA021HzPOqUtD0FKF9evZPY/lmFTCCjNunSe2tsA==
  login-banner post-auth "Running-config last saved 1/7 by J. User"
  snmp-server community public read-only
  snmp-server community private read-write
  snmp-server contact "jpublic, jpublic@mycompawells.me.org"
  snmp-server host 172.16.100.183 public
  snmp-server host 10.1.1.68 public
  snmp-server host 172.16.100.101 public
  snmp-server location "2nd floor lab, row 3, bay 4, shelf 5"
  snmp-server name arx1
  snmp-server trusthost 10.1.1.68
  snmp-server trusthost 172.16.100.183
  ssh-host-key dsa encrypted-hostkey
u9oFngi4v7vM94WNlfxoUKFY23fhmZ9JQuS6aEgd6cZip2E/2pxymcXAHYlrwQj1P4BAhGD6vxTxUKXQZ4kqgKuCFJpj96npQZ
3hgvoOowZt08EawDGkLRbOM0RqeGaQa+O1sfShi2f7mPB4wY5w8vubG+OmLv26GxmtJpP6DZIXemQr8X8olJphzwJNEphH0COH
0xlcK6MueVnolMVf7Pr54b7QPiw+5hbciSuhHER48rj50GwGOfNhsGjGhYizS6uZt84cyIU5LIft2X7QQ25zcSAn44VtefsqBE
WF/q44jCbwM1fbz8E4Ichiv6vHbnn/basWK9UniOHEYdEOlo3ZGcax930CZvoFoupmkPfUEH/vK3ZPqh3BoQKyhUXZpnYqN6qT
OJalxSy3aEaJpASnKMoMEbZAYZzMJAGM7+YLX/qe4OyK8EO3QzIS4YJicYxPxx2MOTh8C9TZ7XrwtTlT1KIHv9e2gkeq1jpwLi
```

```
j5oMiw0WphyB2Vdld6k6WV8GYtXv1PfxxBTEvw5yOqSbRG5+hfDu62gwIuL5RftzbwwfK1aGqcQDPEGuLxbrb9y7XXGRccL5yR
0iCojIVwnot4glU5j2LsIFEnKZKmZ1uMO4u5rZ+XvfZA4EjyJH+VShtCedJt4n4wdlS+AkF6RM+XWP7OkC3FBct/et/nkFji2R
jm3HdU6ALRmAyBmkTm08/X31B7OFpZsPzzUq2yZ8P81hb52yF/oJC5WYg1F0H76Tv0jpTE4MpZ7BrgVoiem3aGPySwYdDY0l1G
AYKsLspH/Uo8sGseluqhgF+15SRQFS5UIzSXP151soXyHcYVB4j7g68jpFzakpSsKLEHEeGLmPK2a5pvRc1ACKm4/ZaZegg0ii
zYkyPljN5FbCh7
  ssh-host-key rsa encrypted-hostkey
u9oFngi4v7sheeATbx/2B1sFeka5+7R2rmsI4ZG504gQX2MkfHCXvPlsavRJb3wHT+XcXCsrIGEAe/QuuvHGKl9ni3RITqZBRF
u3Z/71j8uv7c3iuYQ/7yhhNAlMboxrEmhP2rcaxex+EywOYV7yNhdHKOsoF/dcH3rsAu/Yk3WM8y1OkEN5/BPwl7hck21sM19B
E+n9wpPcuJAGtifjycyk+aqSJUOv8Cxf2n4ht1e9jQcT/JqLecWvtoHG1bv6TqUXBqNpnunG1WZa3jyVRU/OMKu7zUUX5iAmQN
AncZGnNz93Lw+e+p1xVGVZhE5rREnnCjVN0de6xiCeusjN8wJBs+Nqm7lfkiQTbMDS/3QDr5dk+LX0sXreVkHK6PSDR27T1dCt
wF6Z0IBmOwzwMd/1uHco2VPLJwcBkmLIc8w1A8tbHsphocGo8ysxX/i4aG9UzqbaKHFwX9Ax4WspilK+3zE0GPyvS5eC76gvZ6
Qbcf/hb+e/HEbKaavetaHDGbL5RIrqN8qF9Ow+RdNE+oi14AcOgtVnn3FHOzpf5RIBHpUPDcuANrcNQUEu+kp6vViIw37RWr4L
tz5oBRK5sVb4b4i81THAQPw3qxHtsgDuW4G1Q9Wq0ETsHr9XpPR6AZj6/JGBpbJWPlfcVAzeKxuppwquqic72a1qCtUyyIW9YM
ImwLJBF00ogcJYm+X+AM5nOPOm+Aw3sb9cwUCZ7Ff1u9AJ00IMrH+SgZ7hOcLu1tm4Bwc/CmPvQ+TfS1JJZWZpLUmELSAeTTFF
jYB/+Udm28OTqHbrybwaxrE9AuAQ/Enae5hTdr80pjpalVoDdoTtDlvBFbGo8tOASvUZLfzqHm1PJqLeZmR+c3mM4fT36uF9mN
0ud5WiV6Xp67DEXm36t33ZZfx5uXFQOaS6wnItnQ+yDlCqLsmHPk8Ou9Uh3boVJ95fSns8o+6+Z5z6U/QzW8FDjBgU1a/1l3Us
vWT76dQXLhlFG1/YWpp7p2UHn5m6WqhxdfGf8J41Ey2TziR7roEsvUPAom40AIEadKtjSaeVaq02IflB2ceva/3ck/IVphP/ag
5ELzBtU21nb+uDLL+cdWDHO8ZN0o6N9eE16pWdr7PqsA/CJmLRfa71yDdq6bU6JGWALjVYfjnvb2y+dicag5npVa1yLMzQFBoQ
OjgGel8WuMY2W4Hnk3C+zqe2GkRZVobxy6fF113OOCJlPl8EP7IQm40Hq5gNK27xWgSo4By5wBo12eMQbqfWacGct27BERneKf
Rk4+QNti+hrRztkqscN8NJ/Tl5jDpvFtSh3fOP+JOD2aAlYV09K3mxNeMGQjigD0an83sOm8FSAtVQ4G3lYDf7U08YMDa5CT1c
V3GEGyw16l85ZFU2ZnIFHhCPHV8wvkNI9k01+hog8F8LsY60e8ISSjXzLGvgaRWQ0gN4iLNAkHAdCV6eYA2tiMX2rNldOg04qC
T4JXw7O6OybWILrngPngIqbXum4kFxx814BbR8FgaivLzc2HxOCNvulb3a4A3fYmjjLswVP651BWiWkIf/w1N5nBQt8mrHAXQr
Rp2+6b98JSzRBFymFd9ZIPeD/D05opoUaqQFPIaOEAP3LylNbBiM4RbUwHTWnszbkFi9eoeUwqPztoVyYMX2jpzcNV+Hw+wmkd
kOeTkM/W7tZ7FPqsRb33VLojrfkiXzkGEhVGsEeOIdYuiXftxLfeb4JhztqLrX55T0qxzwjI9Xm73srcg+CtOo63uggiXzH7qf
BqdCBk+gEvYtTGZTc5Bt7R4SDZMA2+VyIBNqintayJZCqlKsVZCdHDt0JZgR8TqIG9llxSBwK9/MMSlQ0rfKM7EIFPbKkdGQ4g
2i4xZsfrpEpCaZZ1+Z/KhWE6TJCuX4HWdPDNjDEtstw8+Z43blDX/tXGNyPG3IuKdh+VYU9zIr1eQj/RPKetKxfkFx/B9mpWCh
UX9/BwIu7FmeNMqaUxoVlDkNF3AXfRGzpjacia/4S+tzCZi0TCWVhfK0zIEmwaoCC0A+1WypRsUhlSV7B948MRicUisYvBuUvd
F7IYf98uZSOfgkxbaQ+OLHbIJXoz+uQUpJjq7psxWLP58h8BxyB6eBqQbQjfc+Nk0M1tilH0VG080kYoAxXm6xxyDmtWOJ8aTD
RyDCluMZVgFMqvNTTyaPRlia5LJzkCqy4P2mSXVSCEuu0a49lyh5AGCPibS7pTEhTZB5v+RvKJIVX6VK+H6c
  ssh-host-key rsa1 encrypted-hostkey
ufEEbrBkXmixIPlmxQ7i5umlhaKT0EL8QoodX9tvS6bl/NNsX5YmsWBYOzNLXh781hG5kXma/D1l9PCJcGvVfL6nTYtmmITDOr
r2g978B25D1SbixPJBaew5TpyHqd0kJHFwLiTKlcGoSUl3vUEDrwj6AHyDFOn4e+hS+n0C1rRCvVw3vC466VB3lDRQJVJ5aBA4
9DY0uRhLmlz7FIbOkyzErnEgVOzfwYDLS6BUB3mxMVH8mJ9mi0pp3H9bjwTvtwS+/u5b5caij30rN44NQFrtzBbOEGuRrznbkq
OdxgQCSH+axkdNMDQspCREAUSYnlxRrqyS+I8fpH6Nv9Vs0JD3WTt95dik+aux9fkxP5KGTCN8LtFbs0CsJjqQbP/0cfLXDVZR
eAQyFyNhgnTdeNECquEPknNiyQg64OXkVnGfmJPiASihEryX13PvYaf1VAEyWoKb0VZY6UuYAN9IDT5hFrDXxVOI6phoAVaZ6L
10GbAcmoO7Os80OTmf2IQqOqXYXD6Plwv8fewyIv+AcVyjKj3Mjd41unG+4DbyoCDbjwWojWT2m2n2IhgnRR12DU8Cx2eEz75j
mCRSWxBvYQeHMV698z8ZvllL4md0GfQv2oqZaiNil4RJw4G/7UgERykol7lxTayJi2kBlVgs4bNZo2TFtCgRBTXDAEcf7I8TA7
wMWjhSHKFmqVBZlg2e4kWOXW/fXmpzVJieUj0q0pxwN3KkjD2abH/LRmCBgucySKtg8aZ13GKCGheAkiFv07ay4cFHnxWh59PX
R+cxCCDdW5lq0EIYP6/pF8Cahn4RctcdXOL9Q5dy5L/mOQ0/A7oy3erljlZCNB06Cft6gLz/P3tdmf5cq1wG/yTrJBVjHmaEap
W0koG651hlYZAMuiFjWjLJPF2LM9TbaTOooKA8w3GB0WfDPgvEd2kIkWmZkLDr/afgKybAYo0cju5Pt33CYC7UAfP76fU72un9
3SyFKJ2qWqXGaujPXDidOzDBD0NJC4gmSfBIy1+wojdQvCjpyfOwUg1mkKQqzBGeHH3DQws/789V8wbiW2Mldi/UfGCWWKgPvS
290+mzJ9nCkJBY0efyzaG3P7Sc0NEaVo51V7vg358hphCQKKMu6FHHvvie6S3D8tGB01Q6XGrSZjlMnS0lqAtuk0BB/d0h0icq
vmPqCaNpTQzbYoP7DDUwhr6/PNVZUbwfia4IeuUd09jFS6YQW75KuxTqzwfjd2KM0QeRPArLMvYILmulKB9WdbX41ws8wXDpXf
eE1PAxux6QftYGJzcsrzHMxZROCg==
  ip proxy-address 192.168.103.91 255.255.255.0 vlan 103 mac 1 processor 1.2

;========================== management-access ==========================
  management access console
    authentication primary local
    exit

  management access telnet
    authentication primary active-directory
    authentication secondary radius
    authentication tertiary local
    permit all
    exit

  management access ssh
    authentication primary active-directory
    authentication secondary local
    permit all
```

```
    exit

  management access https
    authentication primary active-directory
    authentication secondary local
    permit all
    exit

  management access snmp
    permit all
    exit

  management access http-api
    authentication primary local
    permit all
    exit

  management access https-api
    authentication primary local
    permit all
    exit

;============================ config-if-mgmt ============================
  interface mgmt
    ip address 10.1.38.19 255.255.255.0
    no shutdown
    exit

;============================ interface gigabit ============================
  interface gigabit 1/1
    no shutdown
    exit

  interface gigabit 1/2
    exit

;============================ event-severity ============================
  email-severity auto-reboot level critical
  email-severity cpu-failure level critical
  email-severity server-offline level critical

;============================== ip route ==============================
  ip route 0.0.0.0 0.0.0.0 192.168.103.1
  ip route 0.0.0.0 0.0.0.0 10.54.245.218 255

;========================== ip route management ==========================
  ip route 0.0.0.0 0.0.0.0 10.1.38.1 mgmt

;============================== cfg-smtp ==============================
  smtp
    from admin@provmed.com
    mail-server email1.wwmed.com
    maximum age 30
    retry interval 10
    to juser@wwmed.com
    exit

;============================ email-event ============================
  email-event chassis-monitor
    description monitoring-chassis
    mail-to alerts@lime.wwmed.com
    group chassis event disk-failure
```

```
    group chassis event temperature-failure
    group chassis event disk-control-failure
    group chassis event disk-control-status
    group chassis event fan-failure
    group chassis event fan-status
    group chassis event kernel-nmi-error
    group chassis event module-failure
    group chassis event module-status
    group chassis event nvram-battery-degraded
    group chassis event nvram-battery-failure
    group chassis event nvram-ecc-error
    group chassis event system-bus-error
    enable
    exit

  email-event noc3
    mail-to juser@wwmed.com
    group chassis
    group metadata threshold-counter 5
    group metadata event online threshold-counter 2
    group storage event share-online threshold-counter 3
    group storage event share-remove-complete
    group redundancy event ha-pair-qd-offline
    group policy event shadowmetadatasharefreespaceerrorraise threshold-counter 5
    enable
    exit

  email-event tech-support
    description Built-In
    mail-to e-support@wwmed.com
    group chassis
    no enable
    exit

;========================== ntp server config ==========================
  ntp server 192.168.25.201 version 4
  ntp server 192.168.103.58 version 4

;========================== logging destination ==========================
  logging destination 172.16.202.8

;======================== at/scheduler commands ========================
  at date 03/02/2012 02:40:00 every 5 minutes  do "show sessions" report adminSessions
  at date 03/03/2012 00:00:00 every 1 days  do "copy startup-config
ftp://root:rootpw@172.16.100.183//tmp/acocfg.conf"
  at date 03/02/2012 05:00:00 every 1 days  do "active-directory update forest WELLS.ME.ORG
proxy-user ny_admin"
  at date 03/02/2012 03:45:00 every 1 days  do "active-directory update forest MEDARCH.ORG
proxy-user acoProxy2"
  at date 03/02/2012 03:30:00 every 89 days  do "cifs rekey all"
  at date 03/02/2012 02:38:00 do "show cifs-service user-sessions all" report cifsUsersGgh
;============================ SSL Certificate ============================
  ssl
    cipher ssl-dhe-dss-export-with-des40-cbc-sha
    cipher ssl-dhe-dss-with-3des-ede-cbc-sha
    cipher ssl-dhe-dss-with-des-cbc-sha
    cipher ssl-dhe-rsa-export-with-des40-cbc-sha
    cipher ssl-dhe-rsa-with-3des-ede-cbc-sha
    cipher ssl-dhe-rsa-with-des-cbc-sha
    cipher ssl-rsa-export-with-des40-cbc-sha
    cipher ssl-rsa-export-with-rc4-40-md5
    cipher ssl-rsa-with-3des-ede-cbc-sha
```

```
  cipher ssl-rsa-with-des-cbc-sha
  cipher ssl-rsa-with-rc4-128-md5
  cipher ssl-rsa-with-rc4-128-sha
  cipher tls-dhe-dss-with-aes-128-cbc-sha
  cipher tls-dhe-rsa-with-aes-128-cbc-sha
  cipher tls-empty-renegotiation-info-scsv
  cipher tls-rsa-with-aes-128-cbc-sha
  ssl-key-store /u3+7QAAAAIAAAABAAAAAQAGdG9tY2F0AAABNVY7HT4AAAUBMIIE/TAOBgorBgEE
  ssl-key-store ASoCEQEBBQAEggTpTA3cglh/TGaZ/wusnQXPilHBxc9l3CRS2VqorWi22epdyV93
  ssl-key-store EV8Ro+fn/e/083y7y/CIhwF5I4dH0LjGn+muQslPftrtA2erOFGcrjQkbF94nF4R
  ssl-key-store j87KBQiAcbGfzLOMfpywLxxfyWi2DZMofCgJneI9VHqpaaAciap1P0MD+3LskBAq
  ssl-key-store OgkLgfA/Hul7m5GrYOLF8Kr1sosfmCMbEvzpwf/a3uBueKqhxfQmQcYVAQuS13Cm
  ssl-key-store Pgu1ecdHITdyYABI+/FLU6BvOuPSpYjlEgZqKCT9GEA1jeYGgov0tS4sXXYtjaxJ
  ssl-key-store UoJRkAQ/BuAdJCWFFV5W2UNr7SpK1ed8X3enZSfBQwZiHVLAk+Kbr/8qaqbNhfU5
  ssl-key-store LfBDRFFaTXbnjHMi/077FH1FB0xTkbISBIIL/9wsVEU/H7e1lPCQM44Dl3eQCvSY
  ssl-key-store p1XjJkdQwD2IiLq/HVCev/31GVj48i+CgUmD/9f9oa9JbwVMsCdn+xRzcdA1w/i1
  ssl-key-store S7L0kM78pbqgh7o6rqJyrM2fSC1H2FT/gQdS75xHBstP4ooAnWYJWXVvMcNoBK7M
  ssl-key-store 4lfWh0V77jByjogGVd/pbxO1/2GJw7naTobOYXY2C8OyBm57hEyx+Ngux1FCe/l7
  ssl-key-store d1FHHgfLc+0sp2m9naMxHQ+9VgmVjV5HD9D7TANJ9qmImG80BEuyWQ6SNApGBB76
  ssl-key-store RC12B9m6KS6jWyOARQoQ6KLtABhwg8jZ/kP/0cjpngcc/i0Kyaw5G3waZnrGUG6r
  ssl-key-store tXLtwbjJpitKiqbsTLipK5K8d49jtv9WDgjuuyP3HOLiBX+99pPBgOlvVZCgca32
  ssl-key-store WfOH//GId73u85Xsy2Ic9ZKhoCqen8puk3L0Mm51MsyewN+wcsPI3Dokmvc/WKs3
  ssl-key-store XqjAjwtnZrdoqIuaTsWkHcgmV4q49V2U9km2foJYkETim5uOCq32/UnBCY8V4IE9
  ssl-key-store 5i4bNbzaY3Q/e33ezIg9rUHWOWRAKkgsZ66w7DZtfdLRp/NNy+uRjcJ538Knh1oR
  ssl-key-store LxbMsK/eYY2GU9YIlshr+3TB6PNMpxLE5tfuBSWMRBu+TK+eQqyRmoslI81CXbCw
  ssl-key-store ubY5GcN1Jci+SM1/t0mODGJMbQGqe+jpVNiOLXaUEe+v0M1QIMCp/8rdLukk6Epk
  ssl-key-store 8CWwWvkQEbAoy2pAcpPH0XZd+7rG9C0ht/2rcPNFv9UVv1b5zz5MqPneIDJI7oPa
  ssl-key-store 73/FDGiBActv+YWqbby+pKDxVBOc+iLu/VEUmqqLqd/8qptLeyHVpb2B8dHlaCTx
  ssl-key-store PBPqd7uBY+/17Y9/vq5tH9Zc9lhlnF/wGB7EQi5/Lq3Gs9RzEI75hKJjhi+50bAA
  ssl-key-store P6bsohjpLxM2Jw+am2fQO2pnagvSZ7+DFtjxYVm2pJOrDj+jFZNX14SHHsUd5DL7
  ssl-key-store ojLnAVKzzBlix+FUKHQds37uBL2kgHdNiJvwqHEOIlmCEGUWX7IXV8MYvrQi7EaK
  ssl-key-store t6sOTDRY9CgDkRZfUE/9KOp+dBWOJTLM5rO3qWb6ekUvcg0feN53Z/+3bweKp8js
  ssl-key-store uobOThuB+TJk2060FJhO98d0mEECoXlLkocSM25lMR8+SvUb8qA5ac6tUQ1LbPQE
  ssl-key-store 5EqPB5jDTYjsCmducWfHR3RdV9ZYvSeAfSv5x1dEcEmbpA55jg0XYA1iegIlDxAa
  ssl-key-store QxNGarg1XsmWNfmHvR3Okz9vi5xoAAAAAQAFWC41MDkAAAOMMIIDiDCCAnCgAwIB
  ssl-key-store AgIETzCy+DANBgkqhkiG9w0BAQUFADCBhTELMAkGA1UEBhMCVVMxFjAUBgNVBAgT
  ssl-key-store DU1hc3NhY2h1c2V0dHMxDzANBgNVBAcTBkxvd2VsbDEUMBIGA1UEChMLRjUgTmV0
  ssl-key-store d29ya3MxFzAVBgNVBAsTDkRhdGEgU29sdXRpb25zMR4wHAYDVQQDExVsb2NhbGhv
  ssl-key-store c3QubG9jYWxkb21haW4wHhcNMTIwMjA3MDUxMzI4WhcNMTMwMjA2MDUxMzI4WjCB
  ssl-key-store hTELMAkGA1UEBhMCVVMxFjAUBgNVBAgTDU1hc3NhY2h1c2V0dHMxDzANBgNVBAcT
  ssl-key-store Bkxvd2VsbDEUMBIGA1UEChMLRjUgTmV0d29ya3MxFzAVBgNVBAsTDkRhdGEgU29s
  ssl-key-store dXRpb25zMR4wHAYDVQQDExVsb2NhbGhvc3QubG9jYWxkb21haW4wggEiMA0GCSqG
  ssl-key-store SIb3DQEBAQUAA4IBDwAwggEKAoIBAQCTu68hWUXPnfWjSso0+vlijqEHFaTS+M0z
  ssl-key-store mS/P4LJGUTi2hjqMKqCJyVw0KKCeDFra77Gzhh+g7p4ppHPHcnF6KITy4R6ioiQu
  ssl-key-store 39OIMU9BTP9wtttlJueBT5T/dZmublC3GIKvjgjwe6viZVGK+RR5R5i77D1tavXv
  ssl-key-store vf2W6zqOhNTXyW93XzzHtpYiJR8S+urDcaKwUkkszKWH41zWzsumfl+9iG/9YU8/
  ssl-key-store mWX6MXNbnCC5gbWBpIh0nRL8okAUdMk6GST5+hiMNLoZCyIc5OamPYE637L+8QAv
  ssl-key-store dhkyqBRO+LNJMDRFpycrP2s/NS4wuRTkuREzKK7/hKxKg6iqM4NTAgMBAAEwDQYJ
  ssl-key-store KoZIhvcNAQEFBQADggEBAIsl9z6JbiKUl9ELgX+Sebn9BsDtVa3DVNso9USQ2/SY
  ssl-key-store XMjZEKwNP0bf0QTNA2ouL0LJPrLk+msKOHiKGfprx3AFMzJgjidNzK0Q4ijVKAC7
  ssl-key-store m0CHV+/BdaMEBkvZJetGIk3KCvfTX8/arx5bS6mMKF5of+BkNHgJUkWD9go/tBks
  ssl-key-store u3cl9zp68UY4JKzr1uhCgHFbDtGH5pm9nQiEZlhSmVHeXVpEqOzMg916eAnGqwKm
  ssl-key-store RhRUDSdw+aVf53cM3FW+38gwaIJ+MJRiq9gQCLrxYv6c8vCs46t7EkohBy3emiod
  ssl-key-store 9JrAuShS16qqDgajB6qG7FnUXXGvdoF2zSsG+jKYftkwK68S7rNi4klxmIL/PWOO
  ssl-key-store ++f/YA==
  exit

exit
```

Figure 33.8 *Sample Output: show running-config (ARX-VE)*

```
stkbrgA# show running-config
; ARX-VE
; Version 6.02.000.14335 (Feb 29 2012 20:11:44) [nbuilds]
; Database version: 602000.31
; Generated running-config Fri Mar  2 03:32:24 2012
; System UUID  8fa98111-55ec-d1c8-9380-8dtu78fab47d
; ip private subnet 169.254.219.0 255.255.255.0
;
terminal character-set unicode-utf-8
;============================ config-if-vlan ============================
config
  interface vlan 1
    ip address 192.168.66.62 255.255.255.0
    no shutdown
    ron tunnel toBoston
      heartbeat failure 3
      heartbeat interval 10
      peer address 192.168.25.5
      no shutdown
      exit

    exit

  exit

;============================== system ==============================
config
  clock timezone America New_York
  hostname stkbrgA
  ip domain-list wwmed.com
  ip domain-list MEDARCH.ORG
  ip domain-list bigorg.org
  ip ftp-user juser encrypted-password eX1Jci+dkoPRvYvA021HzPOqUtD0FKF9evZPY/lmFTCCjNunSe2tsA==
  ip name-server 192.168.90.18
  ip ron-user admin encrypted-password euTb0S+dkoPRvYvA021HzPOqUtD0FKF9H0nfBBv/vYo=
  ip scp-user juser encrypted-password eX1Jci+dkoPRvYvA021HzPOqUtD0FKF9evZPY/lmFTCCjNunSe2tsA==
  login-banner post-auth "Running-config last saved 1/7 by J. User"
  management source vlan 1
  snmp-server community public read-only
  snmp-server community private read-write
  snmp-server contact "jpublic, jpublic@mycompawells.me.org"
  snmp-server host 172.16.100.183 public
  snmp-server host 10.1.1.68 public
  snmp-server host 172.16.100.101 public
  snmp-server location "2nd floor lab, row 3, bay 4, shelf 5"
  snmp-server name arx1
  snmp-server trusthost 10.1.1.68
  snmp-server trusthost 172.16.100.183
  ssh-host-key dsa encrypted-hostkey
u9oFngi4v7vM94WNlfxoUKFY23fhmZ9JQuS6aEgd6cb89iO5/+2ZIxP0zRUO8NkSKqf8XWyK70IHlveGwe47iiRqhF+fGGcR+T
Qve0jkkUVUHySq+Q3YP05KLie3vVgmMGJqjr/TSzfAYvafTwktJx7+KyMSkbOB+yqbhysnF2C2ya+TFaB0sXBvIF2dqta40Hgd
QMROw9puqOjNzjpoNMGhQYq421aRwAhqfUgYFAo1zf2U9R42szdJsJkcgRpekCiWj/Ijw5jpTzG+8hG73nq3bke5vOmZjIOisn
djFVjUKb82KDjLhYQopaWX4V5eRdc2OX2zxBdi88+WVn/KG0h8nYS/Cj2gHSaGqK2XrQBX1qfNi7tswV170+RSBX/sKeoiATOC
RiECM3qqsBEKv2oFGWgSNrxzFbDDGGAXzerGvhMriSDvXPMVSfFdrXxeQ7JJtE/n5tme2vZpNMxDVXR+6h0aKGXP9Tq9eqOLgs
jxYjgF79kGOlRjibrD0BNvvKmq3GROh9wQGcgcItIjIUHv5D4/km8u/fBuQ26/Oe558geMNqQ0cimBn/b46Ag2MdlblAP4DTXO
nfsR5cpkAja5dcFgnbg/OiVVKOzOjdr34iozVlTx3FJqC6q2YdRdccCTkPaR0YV1/LxVeTgXMsU9iFFMletOkv1rf0ipFPh3fU
FzyoDJKqzHuPDhxb9Hi+LEKWVNXXfu0D3s5JV5yPmneL2R0xZF93f1QPW5g9q1DjJj8ZO8/AurZG0tm/UBPK7UaB7wDcFxJBYW
1cjakKyjc84FAbW8CDAD32lAzuwB7rHs7xCadvlLWK9LGdJordSijgFJmgG6lgqcVq37mlwtxz2IXL4fU2XzWTPewpjGfG3ooL
GaX3KQRYs0LEcA
  ssh-host-key rsa encrypted-hostkey
u9oFngi4v7sheeATbx/2B1sFeka5+7R2rmsI4ZG504gQX2MkfHCXvPlsavRJb3wHS96ebpUocUunaFv0xXu3qoOHX61p5BHrIZ
```

```
2F7DUI8bX1UkuN88Qj1ON2fyKS87QYHXunnoCuKQqKOxyYHwzB1P03n+V/j9vfx0sLefJu83Gta+KpQTfRVec+72YwNHj9Cdai
cnF/1YkiLD+4JmGFLrmkOH78JrHMehWHgwaGaG67wRa/eJV0VabGt+KBrmxM5O7Rfy4IKVfYy7vJSnrhNvzN8L3mnTEcR250Fu
ui5/x0ZSEXOAuKbTVDUmmkuUzlOP4NlR+l02D+T9mACG0JrLRbiGniz8+MJ0cY4YLz9ICM+xKUw5yusYbViTzG1JlBLeeVTzfY
fqNyJIZ0c1XgGkTIZhyYaVALj6NoFfAbJhAyPeZE9Zb9oBszFGs0NgWt2up7+gafPlPDkCFRWkKrpiEDfM4cuM7klMiWGgfeI0
hIzDBnrKIIYRyTexgk8EdLISXjSkTUd5E741+iLG95usrlCm9vhrBMaSRmxadjeVsggVlGvobvPHrCz4Tepw4EwttbCZodYXJR
1eS4hZ4ANWSyJ86MHDiSgyBbCwb1i4nlBJXUiRkHddKsrnFJz0rcEQ45uzQf5GhiQf5m73PkD8omGDw7Ud7GmOIxLzqhanYkx5
ItC0Egx/th1qbFk6ZB/sj/OGFccNVGXpQvrPnoW2Z4tezOpdOz/ikPIAc6voJUGN9mlbWmc9wOPVEL1MBSR1T2gEgDTVkN0k2p
+jUymj5qUhSQCchjmD/eA5pGEInjtzu8tykaTwO4/mcv3awahH5gVuOBGQzh9OANXVUkapVqH4Bp/x4SNIh+N0znwvipk/KN8M
8/uQF3E07MyLjACIBwuBrXi5JYmVq84bVFd+0rDg8BkVOOmgfhFhrSkt5QIFF92qJQbhkpzNPQiXERJO6CenMJ5c4r+CFOWvtv
ISk4kj7q63DQ6sPMGFepUo2ocvzdTYQWYG1kRxntXdL4pHO13ZpRnzmZYd5ajaReXmMVBe9e1X/Qe/i3k2K/vtK6os6xiFJfCr
XYi6Jn/0FXZwOw7nj59kuqci+O7Tvwgh4828VjD/JjGJXcvo0Z/eYkaDo6//eLlP7QD/gAuQtbH6j85Vnuvs0kysVHt46MT4zz
Q92mQ6sopdUVW3EG2fi+IykdffeRZdrBhUVEOofNW+cScQ1bZQFK0j9wfvEsXPu8plRCjctjqW9xla91mtacmpA7/jIvN+5gtZ
pNSyPSIgij1HpHMLShatNY56ekHEQ/1ou04yAsn7ddAfQkezw4RmU5YZCM3vAqgTelhLqCBvsg0o8rWf06goQ5rzN64mAlakri
4MTZgRTwncH1oFfATphhdghmxA0rsM/hpyjUsLTu2ykQe9RDBamvpTtlpQM1uI9rT5k1454v4geLwiP+zGUFCAM9C87d0RSen9
gWTd4lBi4erjHpLG8MuXdlWZ/EEhERnZ2NM9Iwio79KiMvqukEiqaNIDw1V0dAcSs81fuMmVPtiunpl3AqOjVwJuEcVgNJ3sPA
HHcMuUoLpk/pXp7lUi04Qs8n+fwqzXpiWSAg+6QOu1KlyEV36UX7RLrNH9T02QAAzGDf+xE0k10Ru5hsRAZLXaIZIVYMA/Yz9H
3ytL89sdU1dToTbcVCHUOjWXBmXRwgNfUNFqFONPQ6YFo6eE85KpucvmHxywSW4sqiV9+fSo9TXsXK3G+RMkREsRXiLHg+H4Bf
dyclFUs41Yl7gs8i01ImjbzVpZsFVBxKnXDE3lUKY3vSfQYGIPqrqz0hdEGUTxi+BBqEENhbErUsIYLPJn7alRIbhlMSGyGTaC
+7WhTRTo/ZVepTSEUlXnbXJf69G7lqTrqPtieC62/wVqPjyugN/ITAVUQDl05syVjg486pd3Gmx3goIwUXQ+G76Zrwy0J7l12G
aFfP5mVawdoCk9/A8QOjRMs9HV2AVm0j0RR7Eoh/1bWbPNXindPNNp1lK7mtQ5SdCEjB/LWVkFT76W8oJsWN8zs7tK+BUrdD7R
UoB8MIefVQRLb6pwuM4dxGKUoD6MMynx2/bnzmBdHIADMU4G7voRtZaKp/nhCQM7jIrX28MolDob9agcuKaT80SFSWav71tIOi
nFwwfCnwe5Bu6AIlbMRtfyCLJt/el1Avz2ZlQJpqWYwYEEHeK4l8Dnqucz0fXv0vJ02XqbFDbnzt80BQ3KeL
  ssh-host-key rsa1 encrypted-hostkey
ufEEbrBkXmixIPlmxQ7i5umlhaKT0EL8QoodX9tvS6bl/NNsX5YmsTE9+4p59VqbRdem/XPNI4KsryHTtqPWCcJSlms1OGEymz
OCO775umWuRqE4NNjwil3QXRW61UZfNS3MI4Zeq6k+y0ds73v0nMsJJhDth7TeAxHjtmoQJkGOyrmppn/I8Qw/+S54oP44E99D
pAHSdx0M8V5a7AecWHJ04moYEOTuVxbafb30NjjVHcfqyZ/n4jGfKQC+gX7436oVkubPXOirPzATBhckgjPeuo8nb2FTnHvLk+
60JQuBvJaKEQS6864jCG9sVDXy76YHWQ1g4W++0SkGLrkEC71oK39nykqkZNNOUz/CnezDsBEzQtJCgiVFRpy8NJyfyLFTPJN1
NFvV9ws6Qz2n+Jp6Qi1e/afv0e9jnU1SRs/bDmeNVB3zxMKKjVFOXn7COxXTLnkbezHsQUrwFTsVfJQIiSzCq9KRypUT1eY4Dj
SvsT9snubt3CjH2roKUBkysW+MTvUJP5mcwjW75RcipJGRNRYkelTAu8QGDgNKMp17XcUVIIbbhzDuK2qITqQOHHTIKoj5Dydt
Z0J8deidqivwdeB03v1DXvBelQ/6oPVwe7Netz1PK7g31MMgTXw8OY5iKRQ/MPIy1C8HqN/3qEiZP8G2EMECngzaTO+jyCR/4S
5q2WtSE1Ck0MtdB0pssqDICoTrx803kESCyKmocWUae2t0AueUKWcZ6WLFfSv5XDMq6meDWyJ+9RpSQtRc8PLHUzhXEwz8xSAZ
pXkHUtLmPEIa/mlxPD/FspFOHezS1zAYoPgUXXglTxUWskJixbW8Asp61IbFlGhNgwoDRJtvDgkmRBGkycYc2wlo8+ZcLTh1Nd
o+f+GLs66roEErw7kY9C9eDtY87rQvQjXqZtTOy5rtG9XgJgrnH13Wgb79yTC3qntHZsW0elPyoQO3UZXJJRIIcNpX3IOYAxAa
/eYy6hfRZruxSm6IBVEMxaukDyaS68Id2CuTCe17oZREWRIcxdZwH3UDaXUIl6GDwssHQxiHkocF/F5siHpfFX3Aa6MqR9XDtg
abjT8gbwsMGXVhBYBi8y4V99HXT18UjMBI1mrGqxkO7VCAzi04hgH1Lps4LXVIDq/8/xFKpu9jxf/VAMqg5R37Vgxx34za40YU
fdPNfnlvgYBoYcm0dIpZz7qjQrUFdljDckgOdTsNRULqELCMj1A3aivxpylihAKEK7Zk1S6TxBsU7y076angHmB5TVS9zzLrm4
edaCcIBfs75OgIfzpKrfXvccSaFA==
  ip proxy-address 192.168.66.64 255.255.255.0 vlan 1 processor 1.2

;========================== management-access ==========================
  management access console
    authentication primary local
    exit

  management access telnet
    authentication primary active-directory
    authentication secondary radius
    authentication tertiary local
    permit vlan
    exit

  management access ssh
    authentication primary active-directory
    authentication secondary local
    permit vlan
    exit

  management access https
    authentication primary active-directory
    authentication secondary local
    permit vlan
    exit
```

```
  management access snmp
    permit vlan
    exit

  management access http-api
    authentication primary local
    permit vlan
    exit

  management access https-api
    authentication primary local
    permit vlan
    exit

;========================== interface gigabit ==========================
  interface gigabit 1/1
    no shutdown
    exit

;============================ event-severity ============================
  email-severity auto-reboot level critical
  email-severity server-offline level critical

;============================== ip route ==============================
  ip route 0.0.0.0 0.0.0.0 10.1.38.1
  ip route 192.168.78.0 255.255.255.0 10.1.38.1
  ip route 0.0.0.0 0.0.0.0 10.45.110.218 255

;============================== cfg-smtp ==============================
  smtp
    from CAGIVA@wwmed.com
    mail-server email1.wwmed.com
    maximum age 30
    retry interval 10
    to juser@wwmed.com
    exit

;============================ email-event ============================
  email-event chassis-monitor
    description monitoring-chassis
    mail-to alerts@lime.wwmed.com
    enable
    exit

  email-event noc3
    mail-to juser@wwmed.com
    group chassis
    group metadata threshold-counter 5
    group metadata event online threshold-counter 2
    group storage event share-online threshold-counter 3
    group storage event share-remove-complete
    group policy event shadowmetadatasharefreespaceerrorraise threshold-counter 5
    enable
    exit

  email-event tech-support
    description Built-In
    mail-to e-support@wwmed.com
    group chassis
    no enable
    exit
```

```
;========================== ntp server config ==========================
  ntp server 192.168.25.201 version 4
  ntp server 192.168.66.9 version 4

;========================== logging destination ==========================
  logging destination 172.16.202.8

;======================== at/scheduler commands ========================
  at date 03/02/2012 03:41:00 every 5 minutes  do "show sessions" report adminSessions
  at date 03/03/2012 00:00:00 every 1 days  do "copy startup-config
ftp://root:rootpw@172.16.100.183//tmp/acocfg.conf"
  at date 03/02/2012 05:00:00 every 1 days  do "active-directory update forest WELLS.ME.ORG
proxy-user ny_admin"
  at date 03/02/2012 03:45:00 every 1 days  do "active-directory update forest MEDARCH.ORG
proxy-user acoProxy2"
  at date 03/02/2012 03:30:00 every 89 days  do "cifs rekey all"
  at date 03/02/2012 03:39:00 do "show cifs-service user-sessions all" report cifsUsersGgh
;============================ SSL Certificate ============================
  ssl
    cipher ssl-dhe-dss-export-with-des40-cbc-sha
    cipher ssl-dhe-dss-with-3des-ede-cbc-sha
    cipher ssl-dhe-dss-with-des-cbc-sha
    cipher ssl-dhe-rsa-export-with-des40-cbc-sha
    cipher ssl-dhe-rsa-with-3des-ede-cbc-sha
    cipher ssl-dhe-rsa-with-des-cbc-sha
    cipher ssl-rsa-export-with-des40-cbc-sha
    cipher ssl-rsa-export-with-rc4-40-md5
    cipher ssl-rsa-with-3des-ede-cbc-sha
    cipher ssl-rsa-with-des-cbc-sha
    cipher ssl-rsa-with-rc4-128-md5
    cipher ssl-rsa-with-rc4-128-sha
    cipher tls-dhe-dss-with-aes-128-cbc-sha
    cipher tls-dhe-rsa-with-aes-128-cbc-sha
    cipher tls-empty-renegotiation-info-scsv
    cipher tls-rsa-with-aes-128-cbc-sha
    ssl-key-store /u3+7QAAAAIAAAABAAAAAQAGdG9tY2F0AAABNXhrlnQAAAUBMIIE/TAOBgorBgEE
    ssl-key-store ASoCEQEBBQAEggTpxvuf/B+6yAe7YflkQHonJ566sH8s7ua0IluJWuravUrzNkGO
    ssl-key-store qcTXQwRh+nLCgwhPpKh0zSUC3RKkOITоTGwF9CzMof0EiSGbBzTaWXJ/yOvbOq5j
    ssl-key-store LcbIdiQxrXKFskmE771o8L5QT58TDeFxEjyxoj7/NAUlxwTEF+wmbSaJW389RVx1
    ssl-key-store cGElFyBkh5OB2PQ8pNPwijJamXx9rYOpbBjpb50Xg5kE0x1cJ0/2tX+q+EuIhG6v
    ssl-key-store VHjxEnrrYv1zR8pzGsvYPYF1VBZs94W5Tgnte6Ln7cx3OyZF9DMQnwBVRNwfo/pH
    ssl-key-store 3ej6Hfm9PE/+hglgn2ao/0VcSKuiHkDnErMqdbtuMGITOgmOxARFZuQSSoZ//UuZ
    ssl-key-store OnidnjPv5YgWzSe0uOqGMCye5XaycWNJ9Cc7KBJqYKtMPRM6jwEUr8QLPQLVu4Kn
    ssl-key-store SFs0PNHlRD5UVrG621VI3xvnjJqFNyseOtbQ2iZpARBXPsmHfs4t6oh2G6aI1Rxw
    ssl-key-store xVGs7F4R1GhticBDkzEqxJ5QjYqt6VkuIyow672woem80JqZ/WBtABZPZR/Zo8tG
    ssl-key-store eOMgNSrx2oF/207poU7E1bw+C8TNFvdYxxmZmSu81leATvNVDB23prP0CR8G+G0w
    ssl-key-store P7TegyoCy7uQTpPxGBQjnyYJjtuGcfiGWKw2FxsyEUkF5ZYAre+6a9LmGrRNrGyV
    ssl-key-store pD51/w02XBAKpBMKCGT+qLzuXNtdM8gsqpLaCJSCR1P6SQ43rSPoMxYcCf68doxx
    ssl-key-store Pz/WNE96kdLhij21Nw/X4U7Fdm+4XWpXZRTUoIUXYl+U/H+hh8/gGJi/xMZcnJ3n
    ssl-key-store /gvRp6luOLunyIjFbzeZLVBuydNbCTn8HpP5Lj9TCblV/P5kf7yRcOZUFn+MlzrC
    ssl-key-store AQjXrXljWddt79qPfXlopZIIrImjVg7qqGrikD7c5ZLkx4F2N/F/SlDzgbCqxaOQ
    ssl-key-store F0evnQosEGYMfgregAhnPmLhCyIdxd+2OCJXZMyg1+GztkrBbCdo7+2fzlwltu6+
    ssl-key-store lVq3h2lHmMZv0qPcRmzdwBa7mdw5+1dVgN02T+4vgww08gxLvxR+bmSQLfW/OSmr
    ssl-key-store o3dTND1nAqvWmfAKnJNyINME4mUhq0/VVD0RhCKOVUd8T6LbTyWgyabysrawx4E1
    ssl-key-store 8qDHLhx49C+CvGb+7k24g83bwmmvIxvX3l7vMxNvQ4onL3RmWFwO+NDb/MdT4WtD
    ssl-key-store 8xSK08Tem1CLccjDbQugPXnJr1YSDPfAzjAgbj1hmVfIsWw4yQXGzIB70H40hV4n
    ssl-key-store CN7Aaga9Tuy4k4yGJ0Une9/7BZamcTx1dlB8YQ7nTGYR2BSX7+Syhg50c9q4Dkyn
    ssl-key-store dCscpIgmu6pYKHreF58+ezoUmWFt+tihsIXpjrJEAE7pBD+npKKQRozsGjpWh5VN
    ssl-key-store pE6boAvpc0v3Gb7uj4TyyikK8n8gqYCn0He3ybVWMXW718mO+QP6qVAIU0n1Auig
    ssl-key-store D0RGJqWM0UiEF6Utez1Fmeqd39OT9vSWyuSmGmIdi6MEySqpGCR73Xq6+yYogF+v
    ssl-key-store SVZszPe6Jw411Ot/aTRUhJIeZ0EEa6PZ+rTd0JOnDgxoTqCNqi1OLhdogTOhuHPM
```

```
  ssl-key-store SCKrp3pOQ8ITsjwF3lWXtD5LVEfhagTQSxsTxoCd+gCTAwdw6RzW2rpkn8S4XcJj
  ssl-key-store FWZcSSARsrQrXNjPgMAx3krHrb1XAAAAAQAFWC41MDkAAAPWMIID0jCCArqgAwIB
  ssl-key-store AgIETzlzmzANBgkqhkiG9w0BAQUFADCBqjELMAkGA1UEBhMCVVMxFjAUBgNVBAgT
  ssl-key-store DU1hc3NhY2h1c2V0dHMxDzANBgNVBAcTBkxvd2VsbDEUMBIGA1UEChMLRjUgTmV0
  ssl-key-store d29ya3MxFzAVBgNVBAsTDkRhdGEgU29sdXRpb25zMSYwJAYDVQQDEx1jYWdpdmEu
  ssl-key-store ZW5nc21va2UuYWNvcGlhbmV0LmNvbTEbMBkGCSqGSIb3DQEJARYMYWRtaW5AZjUu
  ssl-key-store Y29tMB4XDTEyMDIxMzIwMzMzMVoXDTEzMDIxMjIwMzMzMVowgaoxCzAJBgNVBAYT
  ssl-key-store AlVTMRYwFAYDVQQIEw1NYXNzYWNodXNldHRzMQ8wDQYDVQQHEwZMb3dlbGwxFDAS
  ssl-key-store BgNVBAoTC0Y1IE5ldHdvcmtzMRcwFQYDVQQLEw5EYXRhIFNvbHV0aW9uczEmMCQG
  ssl-key-store A1UEAxMdY2FnaXZhLmVuZ3Ntb2tlLmFjb3BpYW5ldC5jb20xGzAZBgkqhkiG9w0B
  ssl-key-store CQEWDGFkbWluQGY1LmNvbTCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEB
  ssl-key-store AKIlIJc55xXDda1s5UZcdJSwRU8ZbrN1uqvJFTeAvdopTvaJGK05gvFhTIJRevVe
  ssl-key-store Ui+KfuhrINQ3zLeflM9UUvjXwgLQdUJY553WLW2nJbsWTFbdKGG07BLs4KxBFJRa
  ssl-key-store gYxdMPkMk3TIKaOsgzcUsO2WG+6Fta/+s2PNPpbMYZs7As4kF7X1NIpA33COYzLc
  ssl-key-store ZmKo4xR9c+3c0aNPBgpfAa9UR2RKKb8m6yGsMvemrLpUCLTKAltfk/X3msoP8KrY
  ssl-key-store y8r04gOdgZDSEXRNhQUlm9Wer4m/iL4KSMUCNCGHerdu25sA+NT1doc9axH695G+
  ssl-key-store jJJOBeY6w3L50V+7wRMdvMUCAwEAATANBgkqhkiG9w0BAQUFAAOCAQEAV1/q5eaN
  ssl-key-store xOSG0F0/V28MG/F5Zvech1sXtf6usiKtGniNwAA34Pj+JgH4Fk6NQWzXFZ4MALFs
  ssl-key-store 8ciuN/jBMbo/zRTIH99zjzVyRLQ5koznF8ODiEbdD+MZAEyQDkVlBzxZ16Y0vmCo
  ssl-key-store g2+xR/ucCt+9gSEUMHS/QOcRjczFbKRuKjMb3R9KJ5IyUxmzKLud9Ys7CxP3mAV/
  ssl-key-store XhGyBy3IjvVqfrEW95AHVdUJ3/GKa/22KsXapzAngV3uzP6D1V0daVa01OTMsPB2
  ssl-key-store l6P9UWTbg2fcu3BsLdE9eJfOXmQ6jNpaqwgHbQcPKG47NGpdp22Hm41/eFvB39pG
  ssl-key-store /qgNkGKyijbgkfwJTa3GqhGMcvZvXGjjpb6xVcSe
  exit

exit
```

34

ARX API and Notification Rules

clear statistics api

Purpose The ARX supports a SOAP-based API for monitoring general configuration, usage statistics, and specific file and directory changes in its managed volumes. You can use the show statistics api command to show the usage statistics for this API. This command, clear statistics api, resets all of the statistical counters to 0 (zero).

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer or crypto-officer

Syntax **`clear statistics api`**

Default(s) None

Guidelines The show statistics api command shows cumulative statistics for API usage. Use this command to clear those statistical counters.

The CLI prompts for confirmation before clearing any statistics; type **yes** to proceed.

Sample

```
provA# clear statistics api
Clear API statistics? [yes/no] yes
```

Related Commands show statistics api

clear statistics notification

Purpose A *notification rule* is required to support the SOAP-based API for its managed volume. This rule takes regular snapshots, or point-in-time copies, of a volume's files and metadata. These are called *notification snapshots*, and managed volumes keep statistics on them. Use this command to clear the cumulative statistics for a particular notification rule.

Mode priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`clear statistics notification`** ***`namespace volume notification-rule`***

namespace (1-30 characters) identifies a namespace.

volume (1-1024 characters) is a volume that supports notification snapshots.

notification-rule (1-1024 characters) is the rule whose statistics should be cleared.

Default(s) None

Guidelines Thc CLI prompts for confirmation before clearing any notification-snapshot statistics. Enter **yes** to proceed.

The show policy ... details command shows cumulative statistics for notification-snapshot operations. Use this command to clear the statistical counters for one notification rule.

Sample

```
provA# clear statistics notification provMed /mds mdChgs

Confirming this command causes all historic statistical data for the
snapshot rule to be reset.
Proceed? [yes/no] yes
```

clears all notification statistics for the "mdChgs" rule.

Related Commands show policy

enable (gbl-ns-vol-ntfy)

Purpose

A *notification rule* is required to support a SOAP-based API for the current volume. This rule takes regular snapshots, or point-in-time copies, of the volume's files and metadata. The volume can use these snapshots to answer procedure calls from the API's clients.

Use the enable command to enable the current notification rule. The rule must be enabled to create any ARX snapshots, either manually or by a schedule. It also must be enabled to track file changes in real time.

Use no enable to disable the current notification rule.

Mode

gbl-ns-vol-ntfy

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
enable
no enable
```

Default(s)

```
no enable
```

Guidelines

You must enable a rule for the policy engine to use it, or to invoke it manually with snapshot create.

Samples

```
provA(gbl-ns-vol-ntfy[provMed~/mds~mdChgs])# enable
```

enables the current notification rule.

```
provA(gbl-ns-vol-ntfy[provMed~/rns~testNtfy])# no enable
```

disables the "testNtfy" rule in the same namespace.

Related Commands

namespace -> volume -> notification rule
snapshot create

notification rule

Purpose

The ARX supports a SOAP-based API for monitoring its configuration and usage. A client application can use this API to query a managed volume about changes in its files and directories over time. To respond to these queries, the managed volume must keep regular snapshots of its metadata as well as its actual files. The managed volume can compare these snapshots to discover creates, deletes, renames, and changes in any given file or directory. You use the notification rule command to begin the process of taking regularly-scheduled snapshots for this purpose.

Use the no form of the command to remove the notification rule without removing any of the snapshots behind it.

Modes

gbl-ns-vol

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
notification rule name
no notification rule name
```

name (1-1024 characters) is a name you choose for the notification rule.

Default(s)

None

Guidelines: Filer Preparation

Before you begin, you must prepare all the filers behind the volume. The volume must be backed by NetApp filers, EMC Celerra servers, EMC Data Domain systems, and/or Windows servers that support snapshots. In the case of Windows servers, WinRM must also be installed so that the ARX can invoke snapshots through its management API.

The filer that holds the volume's metadata share must also support snapshots.

The ARX volume creates a notification snapshot by issuing CLI commands to each of its back-end filers, including the filer used for the metadata share. The ARX therefore needs information and credentials for accessing each filer's CLI. From gbl-filer mode, use the filer-type command to identify the filer vendor, use proxy-user (gbl-filer) to identify a proxy user with proper management-login credentials, and use manage snapshots to declare that the filer supports snapshots. You can use the ip address ... management command to designate the management-IP address at that station (by default, the ARX logs into the CLI through an external filer's primary-IP address, set with the simplest syntax for the ip address command).

Note

All of the filers behind the managed volume, including the one that contains the volume's metadata share, must support snapshots. The notification rule requires a full set of snapshots so that it can support specific API queries about any file in the volume. This rule does not support "sparse" snapshots, where some back-end shares are excluded. Configure all of the volume's external filers as described above.

Guidelines

An enabled notification rule is the basis for supporting file-change queries in the ARX API. You apply a schedule to the rule so that it takes regular snapshots, and you can also invoke the rule manually with the snapshot create command.

When you create a new notification rule, the CLI prompts for confirmation. Enter **yes** to create the rule. (You can use terminal expert to eliminate confirmation prompts for creating new policy objects.)

This command places you in gbl-ns-vol-ntfy mode, where you enable the rule and where you have some options that you can apply to it. By default, a notification rule retains three snapshots; whenever it successfully creates a new snapshot, it deletes the oldest snapshot so that there are never more than three. You can use the optional retain (gbl-ns-vol-ntfy) command to change the number of retained snapshots. A notification rule requires a regular schedule, which you set with the schedule (gbl-ns-vol-ntfy) command. To enable report-generation for each snapshot, use the report (gbl-ns-vol-ntfy) command. You must use the enable (gbl-ns-vol-ntfy) command enable this rule to take any snapshots at all, even manually. You cannot exclude any storage shares in the managed volume; a notification rule requires that all storage shares in the volume support snapshots. The snapshots automatically exclude any shares with the special replica-snap setting.

You cannot configure a notification rule in a direct volume. A notification rule requires snapshots of the volume's metadata, and direct volumes do not have any metadata.

Guidelines: API

To activate the ARX API, use the management access [http-api | https-api] command followed by the permit command. This opens a port where a SOAP client can send queries to this volume. The HTTP-API port is 83, and the HTTPS-API port is 843.

Use the following URL syntax to access the API documentation through HTTP:

`http://`*`arx-management-ip`*`:83/arx-api/`

where *arx-management-ip* is either the out-of-band management-IP address (interface mgmt) or an in-band (VLAN) management-ip address (interface vlan). The permit command determines which of these address types are available for access. Use show interface mgmt and/or show interface vlan to find the IP addresses for each interface.

This URL accesses the same documentation through HTTPS:

`https://`*`arx-management-ip`*`:843/arx-api/`

where *arx-management-ip* is a valid management IP, as described above.

You can use the show statistics api command to find usage statistics for the API.

Guidelines: Managing Notification Snapshots

You can use snapshot manage to incorporate existing filer snapshots into the notification rule, but you cannot use it for any metadata snapshots. The command is therefore not recommended for notification rules.

Each ARX snapshot has one or more component snapshots on its back-end filers. You can use the snapshot verify command to verify that all of the component snapshots still exist behind a notification rule. This does not check the snapshot on the metadata share.

To see the current state of notification-related snapshots, use the show notification command.

Guidelines: VIP Fencing

During a snapshot-create operation, clients can access the volume and possibly make changes, rendering the filer snapshots inconsistent from one another. This inconsistency creates challenges for API clients: it could, in rare cases, affect queries about file content or file attributes. This has no effect on queries about creates, deletes, or renames, which are all derived from the metadata snapshot.

For sites where these queries are particularly important, you can use the snapshot consistency command to put up a *VIP fence* for snapshots. This fence prevents client access to any VIP that supports this volume. This may affect multiple volumes. The fence stays up until the last filer completes its snapshot or checkpoint, or until a timeout expires.

Guidelines: Snapshot Presentation to Clients

By default, CIFS clients can access their snapshots with Windows Explorer. They select a file or directory, pull up its Properties, and find a list of snapshots for the file or directory in the Previous Versions tab. CIFS clients can use this interface to find and restore previous versions of their files and directories. Microsoft calls this the Volume Shadowing Service, or *VSS*, for Shared Folders.

NFS clients typically access their snapshots through a ".snapshot" directory that appears in every directory in the ARX volume.

Refer to the Guidelines for the snapshot rule command to learn about CLI commands that can affect the client view of snapshots.

Guidelines: Filer-Snapshot Names

Each EMC or NetApp snapshot has a unique name to distinguish it as a backing snapshot for a particular ARX. The snapshot names use the following format

`acopia_`*`id_time-stamp_uuid_filer-volume`*

where

- *id* is an integer that the snapshot software uses internally.
- *time-stamp* is in *yyyymmddHHMM* format. This is the time that the snapshot was created.
- *uuid* is the Universally-Unique ID that identifies the ARX. In a redundant pair, this is the UUID for the peer that was originally configured as senior/active, no matter which peer is currently active. You can find an ARX's UUID in the output of show running-config.
- *filer-volume* is the name of the filer's volume.

Windows servers use a UUID in curly braces ({}) to identify each snapshot. This does not change for ARX-generated snapshots.

Guidelines: Configuration Changes in a Scheduled Rule

If you make a configuration change in a notification rule that is running with a schedule, the configuration change is ineffective until the next time the rule runs.

Guidelines: Removing the Notification Rule

The no form of the command removes the rule without removing any snapshots from the back-end filers. To remove the snapshots from the filers, use the snapshot remove command *before* you remove the rule. This is an efficient method for cleaning all of the supporting snapshots behind the rule. If supporting snapshots remain when you invoke no notification rule, the CLI lists all remaining snapshots when it prompts for confirmation.

The snapshot remove command removes snapshots from all the storage-share filers; you must manually remove snapshots from your metadata-share filer.

Guidelines: Snapshot Reconstitution

Snapshot Reconstitution is not supported for notification rules. Refer to the Guidelines for the snapshot rule command to learn about snapshot reconstitution.

Samples

```
provA(gbl-ns-vol[provMed~/mds])# notification rule mdChgs
This will create a new policy object.

Create object 'mdChgs'? [yes/no] yes
provA(gbl-ns-vol-ntfy[provMed~/mds~mdChgs])# ...
```

instantiates a new notification rule, “mdChgs,” in the current namespace volume.

```
bstnA(gbl-ns-vol[medarcv~/lab_equipment])# no notification rule apiTst

The following snapshots are being managed by the rule 'apiTst':

apiTst_0(acopia_27_200707170500_501f705c-8735-11d8-8936-a58a9e4556df_d
atavol2)

apiTst_2(acopia_25_200707170300_501f705c-8735-11d8-8936-a58a9e4556df_d
atavol2)

apiTst_1(acopia_26_200707170400_501f705c-8735-11d8-8936-a58a9e4556df_d
atavol2)

If this command is confirmed, the rule is deleted and all associated
data related to the aforementioned snapshots will be removed from the
switch.  The snapshots on the filers will not be removed.

Proceed? [yes/no] yes
bstnA(gbl-ns-vol-snap[medarcv~/lab_equipment])# ...
```

removes a notification rule with some existing snapshots remaining on one of its back-end filers.

Related Commands

proxy-user
filer-type
proxy-user (gbl-filer)
ip address `... management`
manage snapshots

retain (gbl-ns-vol-ntfy)
schedule (gbl-ns-vol-ntfy)
report (gbl-ns-vol-ntfy)
enable (gbl-ns-vol-ntfy)
show notification

snapshot consistency

snapshot vss-mode
snapshot directory display
snapshot directory cifs-name
snapshot directory nfs-name

snapshot create
snapshot verify
snapshot remove

snapshot clear

report (gbl-ns-vol-ntfy)

Purpose

A *notification rule* is required to support a SOAP-based API for the current volume. Use this command to enable progress reports for the current notification rule.

Use no report to prevent progress reports.

Mode

gbl-ns-vol-ntfy

Security Role(s)

storage-engineer or crypto-officer

Syntax

`report` ***`file-prefix`***
`no report`

file-prefix (1-1024 characters) sets a prefix for all notification reports. Each report has a unique name in the following format:

*file-prefix*_0_create_*YearMonthDayHourMinuteSecondsMilliseconds*.rpt

For example, md_chgs_0_create_20110823031645354.rpt could be the name for one report with the prefix, "md_chgs."

Default(s)

`no report`

Guidelines

Every time a notification rule fires and creates a coordinated snapshot, the rule can generate a report to show the details of the snapshot-create operation. This includes details about each snapshot at each of the volume's back-end filers. The reporting feature is disabled by default; we recommend that you enable reporting for all notification rules, to diagnose any problems that occur with its snapshots.

Use show reports for a list of reports, or show reports *file-name* to show the contents of one report.

See Figure 34.1 on page 34-11 for a sample snapshot-create report.

Sample

```
provA(gbl-ns-vol-ntfy[provMed~/mds~mdChgs])# report md_chgs
```

enables reports for the snapshot rule, "mdChgs."

Related Commands

namespace -> volume -> notification rule
show reports

Figure 34.1 *Sample Report: md_chgs_0_create_....rpt*

```
provA# show reports md_chgs_0_create_20120307023703404.rpt

  Snapshot Rule Summary
  ---------------------
   Namespace Name:               provMed
   Volume Name:                  /mds
   Snapshot Rule Name:           mdChgs


  Snapshot Properties
  -------------------
   Snapshots Enabled:            Yes
   Guarantee Consistency:        Disabled
   Retain Count:                 31
```

```
  Schedule:                      oncePerDay
  CIFS Directory Name:           ~snapshot
  Directory Display:             None
  Hidden File Attribute:         Not Set
  Restricted Access Configured:  No
  VSS Mode:                      Windows XP
  Contents:
    Metadata:                    No
    Volume Configuration:        No
    User Snapshots:              Yes
  Archive:
    Total Archive Operations:    0
    Total Successful Operations: 0
    Total Failed Operations:     0
    Total Saved Metadata:        0 B
    Total Saved Volume Config:   0 B
    Average Copy Rate:           0 b/s

Snapshot Summary - mdChgs_0
------------------
  Snapshot Name:                 mdChgs_0
  Snapshot Operation:            Create
  Result:                        Success
  Time Requested:                03/07/2012 02:37:03 -0500
  Time Created:                  03/07/2012 02:37:03 -0500
  Last Time Verified:
  Request:                       Create
  Snapshot State:                Complete
  Snapshot Origin:               Manual
  Report Name:                   md_chgs_0_create_20120307023703404.rpt


Included Shares
----------------
      Share Name:                mds (user data)
      Filer:
        Name:                    ntap-pr-1
        CIFS Share:              MDS
        Volume:                  vol1
        Filer Snapshot:          acopia_1_201203070737_db922942-876f-11d8-9110-8dtu78fc8329_vol1

      Share Name:                surgeons (user data)
      Filer:
        Name:                    ntap-prov
        CIFS Share:              SURGEONS
        Volume:                  vol1
        Filer Snapshot:          acopia_1_201203070737_db922942-876f-11d8-9110-8dtu78fc8329_vol1

      Share Name:                N/A (metadata)
      Filer:
        Name:                    ntap-pr-1
        NFS Export:              /vol/vol3/arx_meta
        Volume:                  vol3
        Filer Snapshot:          acopia_1_201203070737_db922942-876f-11d8-9110-8dtu78fc8329_vol3
```

retain (gbl-ns-vol-ntfy)

Purpose

A *notification rule* is required to support a SOAP-based API for the current volume. Each notification rule retains some maximum number of *snapshots*, or point-in-time copies, of its volume's files and metadata. Use this command to set the number of retained snapshots for the current rule.

Use no retain to revert to the default retention count.

Modes

gbl-ns-vol-ntfy

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
retain snap-count
no retain
```

snap-count (1-1024) is the maximum number of snapshots for this rule to retain.

Default(s)

`3`

Guidelines

Whenever the notification rule runs, it creates a new snapshot "0" (zero). The former snapshot "0" becomes snapshot "1," the former snapshot "1" becomes snapshot "2," and so on. The rule deletes the oldest snapshot if keeping it would exceed the retain count. For example, if the rule retains 5 snapshots, it only keeps snapshots 0-4. It removes the older snapshots from client presentations as well as the back-end filers. Additionally, the API clients cannot use these removed snapshots. This prevents wasted snapshot space on the filers.

If you increase the retention count in a rule that already has a full count of snapshots, the rule does not delete any snapshots until it exceeds the count again. For example, if the retain count increases from three to four, the next time the rule runs it moves snapshot "3" to "4" and does not delete any snapshots. On the following run of the rule, it deletes snapshot "4" to avoid exceeding the retain count.

If you decrease the retain count in a rule that has already reached it, the volume removes all excess snapshots the next time the rule runs.

Sample

```
provA(gbl-ns-vol-ntfy[provMed~/mds~mdChgs])# retain 31
```

retains 31 snapshots for the rule named "mdChgs."

Related Commands

namespace -> volume -> notification rule

schedule (gbl-ns-vol-ntfy)

Purpose

A *notification rule* is required to support a SOAP-based API for the current volume. Use this schedule command to assign a schedule to the current notification rule.

Use no schedule to remove the rule's schedule.

Mode

gbl-ns-vol-ntfy

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
schedule name
no schedule
```

name (1-64 characters) identifies the schedule. Use show schedule for a list of configured schedules.

Default(s)

None.

Guidelines

A notification rule takes *snapshots* (point-in-time copies) of an ARX volume on a regular schedule. This command determines which schedule to use.

You can also manually invoke the rule's snapshots with the snapshot create command.

To create a schedule, use the gbl-mode schedule command.

Guidelines: Snapshot Grouping

If more than one snapshot-producing rule (snapshot rule, snapshot replica-snap-rule, and/or notification rule) uses the same schedule, the ARX aggregates all of the filer snapshots at once. This is called *snapshot grouping*. By grouping the filer snapshots, the ARX avoids any duplicate snapshots on a given back-end volume.

Snapshot grouping may create a problem for volumes backed by EMC Celerra servers, which serialize their snapshots. Consider a site with 10 volumes backed by 10 different back-end volumes on a single EMC server. If all 10 ARX volumes use the same snapshot schedule, each ARX snapshot must wait for all 10 EMC checkpoints to finish, one after the other.

Each ARX volume must use the same schedule (by name) for snapshot grouping to occur. If you use schedules with different names but the same dates and times, the ARX volumes independently perform their snapshot operations. This could result in duplicate snapshots on the same back-end volume, but it avoids the EMC-serialization issue.

Sample

```
provA(gbl-ns-vol-ntfy[provMed~/mds~mdChgs])# schedule oncePerDay
```

uses a schedule named "oncePerDay" for the "mdChgs" notification rule.

Related Commands

namespace -> volume -> notification rule
snapshot create
schedule
show schedule

show notification

Purpose

A *notification rule* is required to support a SOAP-based API for the current volume. Use the show notification command to see the current status of one or more notification rules.

Mode

(any)

Security Role(s)

backup-operator or crypto-officer

Syntax

```
show notification
show notification namespace ns
show notification namespace ns volume path
show notification namespace ns volume path rule rule-name
```

ns (1-30 characters) identifies the namespace with the notification rule(s).

path (1-1024 characters) focuses on one volume in the namespace.

rule-name (1-1024 characters) focuses on one notification rule. This changes the output to a detailed view of the rule.

Guidelines: Summary View

The simplest output shows a summary of all selected *snapshots*, or point-in-time copies of an ARX volume. This is a table with one line per snapshot. Each line contains the following fields:

Namespace,

Volume, and

Rule all identify the notification rule.

Name is the name of the particular notification snapshot. This is typically the rule name with an integer ID that indicates the order of these snapshots ("0" is the newest, "1" is the next-newest, and so on).

Status indicates the result of the coordinated snapshot:

- Requested means that someone issued a snapshot create command, or that the rule's schedule just invoked a new snapshot.
- In Progress indicates that a snapshot is currently underway at the back-end filers.
- Created means that all filer snapshots were created successfully.
- Created (Sparse) means that one or more filer snapshots failed. This indicates that the notification snapshot is invalid. All filers behind the volume, including the filer used for metadata, must have the manage snapshots setting for notification snapshots to succeed.
- Incomplete means that the fencing timeout expired before the snapshot finished on all of its included back-end filers. This implies that snapshot consistency is enabled, causing the volume's VIP(s) to put up a *fence* against client access until the snapshots finish.
- Fence failed also implies that snapshot consistency is enabled, and means that the VIP failed to raise its fence against client access. This indicates an internal-software error; contact F5 Support if you see this.

Guidelines: Summary View (Cont.)

– Policy pause failed indicates another internal-software error, and should not appear during normal operation. As above, contact F5 Support if you see this.

– Failed indicates that the snapshot operation failed at one or more back-end filers or servers. Check the snapshot report (available in the detailed output) for more details.

Time Created indicates the time when the final filer finished its snapshot.

Source is "Schedule" or "Manual."

– Schedule indicates that the snapshot was invoked by the notification rule's schedule (gbl-ns-vol-ntfy), and

– Manual means that an administrator manually invoked a snapshot with the snapshot create command.

Guidelines: "Snapshot Summary" in Detailed View

The detailed view of a notification rule appears when you specify the rule in the command. This shows two tables: "Snapshot Summary" and "Snapshot Properties."

The "Snapshot Summary" table summarizes the rule configuration:

Namespace Name,

Volume Name, and

Snapshot Rule Name are defined by the notification rule command.

Guidelines: "Snapshot Properties"

The "Snapshot Properties" Table provides details about the snapshot configuration:

Snapshots Enabled is "Yes" or "No," depending on the setting for enable (gbl-ns-vol-ntfy).

Guarantee Consistency is "Enabled" or "Disabled," depending on whether or not the volume uses VIP fencing for its snapshots. The VIP fence, if enabled, blocks all client access to the volume while the filers take their coordinated snapshots. Use the snapshot consistency command to allow or disallow this fence.

Retain Count is the number of snapshots to retain for this rule. If the rule has this many snapshots when it takes a new one, it moves all of the snapshots to the next slot down and then deletes the oldest snapshot. For example, it creates a new "rnChgs_0" snapshot, moves the old "rnChgs_0" to "rnChgs_1," and so on, until it reaches the retain count; then it deletes any remaining snapshots. You can control this with the retain (gbl-ns-vol-ntfy) command.

Schedule is the name of the schedule for the notification rule, if any. Use the schedule (gbl-ns-vol-ntfy) command to assign a schedule to the snapshot rule.

CIFS Directory Name only appears if the volume supports CIFS. This is the pseudo directory that well-informed CIFS clients (administrators) can use to access their snapshots. You can use the snapshot directory cifs-name command to change this name.

NFS Directory Name only appears if the volume supports NFS. This is the pseudo directory that well-informed NFS clients (administrators) can use to access their snapshots. You can use the snapshot directory nfs-name command to change this name.

Directory Display is "All Exports" (clients see the ~snapshot/.snapshot directory in any front-end share), "Volume Root Only" (clients only see the directory only in a front-end share of the volume's root directory) or "None." You can use the snapshot directory display command to change this.

Guidelines: "Snapshot Properties" (Cont.)

Hidden File Attribute only appears if the rule's volume supports CIFS. This is "Set" if the special ~snapshot directory has its "hidden" DOS attribute raised, or "Not Set" otherwise. Use an optional argument in the snapshot directory display command to control this setting. This option is irrelevant to NFS clients.

Restricted Access Configured also only appears if the volume supports CIFS. This is "Yes" or "No" depending on whether or not someone used the snapshot privileged-access command. If this is "Yes," a small set of privileged CIFS clients can access the volume's snapshots. These clients are members of a Windows-Management-Authorization (*WMA*) group with permission to monitor snapshots. These privileged clients are typically administrators. This option has no impact on NFS clients.

VSS Mode only appears if the volume supports CIFS. This field indicates the client-machine versions for which the volume supports the Volume Shadow-copy Service (*VSS*). VSS is an intuitive interface that Windows clients can use to access their snapshots. This is "Windows XP" (the volume supports VSS for Windows XP client machines, as well as newer machines), "Pre-Windows XP" (the volume also supports VSS for Windows-2000 clients), or "None." Use the snapshot vss-mode command to change this setting. As above, this option has no impact on NFS clients.

Guidelines: "Snapshot Summary - *snapshot-name*"

After the notification rule has run at least once, additional tables appear for each of its snapshots. The first table, "Snapshot Summary - *snapshot-name*" summarizes the notification snapshot. This table contains the following fields:

Snapshot Name is the name of the notification snapshot.

Time Requested is the last date and time that someone issued an ARX command (such as snapshot create or snapshot verify) for this notification snapshot.

Time Created shows the date and time that the last filer behind the volume completed its snapshot or checkpoint operation. "In progress" appears if a snapshot is currently underway.

Last Time Verified is the last time someone ran snapshot verify against this notification snapshot.

Guidelines: "Snapshot Summary - *snapshot-name*" (Cont.)

Request shows the currently-active request for this ARX snapshot. This is one of the following:

- Idle indicates that no command is currently running on this snapshot.
- Create means the rule is currently creating a notification snapshot. This indicates either a scheduled snapshot or a snapshot invoked by the snapshot create command.
- Manage shows that the rule is incorporating an existing snapshot to an ARX snapshot. This means that someone invoked the snapshot manage command.
- Delete indicates that filer snapshots are being removed from behind this rule. This request comes from the snapshot remove command.
- Remove indicates that someone invoked no notification rule to delete the current notification rule.
- Verify shows that someone invoked snapshot verify to verify the integrity of all filer snapshots behind this notification snapshot.
- Cleanout indicates that a notification snapshot failed and the snapshot software is removing filer snapshots that are inconsistent with one another. This should be rare. If you see this, you can use the show cores command to check for core-memory files from the snapshot software. The clean-out operation generates a snapshot report that you can see with the show reports command.

Snapshot State shows the results of the most-recent snapshot operation. This is either "Complete" or "Incomplete." "Complete" indicates a successful snapshot on all of the volume's shares. An "Incomplete" state only applies to a snapshot verify operation: this indicates that one of the included shares is now missing its back-end snapshot or checkpoint.

Snapshot Origin is "Schedule" or "Manual." This has the same possible values as the Source field in the summary view:

- Schedule indicates that the snapshot was invoked by the notification rule's schedule (gbl-ns-vol-ntfy), and
- Manual means that an administrator manually invoked a snapshot with the snapshot create command.

Report Name identifies the report for the rule's most-recent action. You can use show reports *report-name* to view this report.

Guidelines: "Included Shares"

A separate table appears under the "Included Shares" heading for each of the shares in this snapshot. These describes the component snapshots for the notification snapshot described above. Each table contains the following fields to describe the snapshot on a particular filer:

Share Name is the name of the share in the ARX volume, defined by the share command.

Filer is a header for the remaining fields.

Name is the external-filer name, defined with the external-filer command.

NFS Share is the name of the NFS export at the filer.

CIFS Share is the name of the share at the filer.

Volume is the name of the filer volume behind the above filer share.

Volume Snapshot is the name of the filer snapshot.

Guidelines: "Excluded Shares" and "Offline Shares"

If any of the volume's shares were administratively excluded from the snapshot, an "Excluded Shares" section appears to describe them. You can use no manage snapshots to exclude all shares on a filer. A notification rule requires that all storage shares are in every snapshot, so this causes the notification snapshot to fail. These tables contain the same fields that are in the "Included Shares" tables, except the fields to describe the back-end snapshot.

If the snapshot operation failed at any of the volume's back-end share, an "Offline Shares" section appears to describe them. These share tables contain the same fields as those in the "Excluded Shares" section.

Samples

`provA# show notification`

shows a summary of all notification snapshots on the ARX. See Figure 34.2, below, for sample output.

`provA# show notification namespace provMed volume /mds rule mdChgs`

shows details for the "mdChgs" rule. See Figure 34.3 on page 34-19 for sample output.

Related Commands

notification rule

Figure 34.2 *Sample Output: show notification*

```
provA# show notification

Namespace    Volume  Rule         Name           Status       Time Created               Source
------------ ------- ------------ -------------- ------------ -------------------------- --------
provMed      /mds    mdChgs       mdChgs_0       Created      03/07/2012 02:37:03 -0500 Manual
```

Figure 34.3 *Sample Output: show notification namespace provMed volume /mds rule mdChgs*

```
provA# show notification namespace provMed volume /mds rule mdChgs

  Snapshot Summary
  ----------------
   Namespace Name:               provMed
   Volume Name:                  /mds
   Snapshot Rule Name:           mdChgs


  Snapshot Properties
  -------------------
   Snapshots Enabled:            Yes
   Guarantee Consistency:        Disabled
   Retain Count:                 31
   Schedule:                     oncePerDay
   CIFS Directory Name:          ~snapshot
   Directory Display:            None
   Hidden File Attribute:        Not Set
   Restricted Access Configured: No
   VSS Mode:                     Windows XP
   Contents:
     Metadata:                   No
     Volume Configuration:       No
     User Snapshots:             Yes
   Archive:
```

```
    Total Archive Operations:    0
    Total Successful Operations: 0
    Total Failed Operations:     0
    Total Saved Metadata:        0 B
    Total Saved Volume Config:   0 B
    Average Copy Rate:           0 b/s

Snapshot Summary - mdChgs_0
-----------------
 Snapshot Name:                mdChgs_0
 Time Requested:               03/07/2012 02:37:03 -0500
 Time Created:                 03/07/2012 02:37:03 -0500
 Last Time Verified:
 Request:                      Idle
 Snapshot State:               Complete
 Snapshot Origin:              Manual
 Report Name:                  md_chgs_0_create_20120307023703404.rpt


Included Shares
---------------
     Share Name:               mds (user data)
     Filer:
       Name:                   ntap-pr-1
       CIFS Share:             MDS
       Volume:                 vol1
       Filer Snapshot:         acopia_1_201203070737_db922942-876f-11d8-9110-8dtu78fc8329_vol1

     Share Name:               surgeons (user data)
     Filer:
       Name:                   ntap-prov
       CIFS Share:             SURGEONS
       Volume:                 vol1
       Filer Snapshot:         acopia_1_201203070737_db922942-876f-11d8-9110-8dtu78fc8329_vol1

     Share Name:               N/A (metadata)
     Filer:
       Name:                   ntap-pr-1
       NFS Export:             /vol/vol3/arx_meta
       Volume:                 vol3
       Filer Snapshot:         acopia_1_201203070737_db922942-876f-11d8-9110-8dtu78fc8329_vol3
```

show statistics api

Purpose The ARX supports a SOAP-based API for monitoring configuration and usage. A client application can also use this API to query a managed volume about changes in its files and directories over time. Use the show statistics api command to see statistics for the usage of this API.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show statistics api
```

Guidelines The output is a two-column table showing the number of calls to each of the API's supported functions. The functions are separated into the following major categories:

- API is a generalized category with meta information about the API itself.
- Chassis contains the procedure calls for information about the hardware chassis, and possibly its redundant peer.
- Export is a group of procedures related to front-end-CIFS shares (see export (gbl-cifs)) and front-end-NFS exports (export (gbl-nfs)) on the ARX.
- FileChangeNotification is a set of procedures for finding out about file changes in a managed volume. These procedures access the snapshots produced by a notification rule (see the *Guidelines* below).
- FileServer is a group of procedures related to external-filers behind the ARX.
- ManualMigrateRule is a group of procedures for migrating files from one back-end filer to another behind a managed volume.
- Namespace is the heading for procedures related to namespaces on the ARX.
- Network is a group of procedures related to network configuration and statistics.
- Policy is a group of procedures related to policy configuration, including schedules.
- Report is the heading for report-access procedures. These allow an API client to access the same reports you can see with show reports.
- Schedules is a group of procedures for creating, editing, and viewing ARX schedules.
- Share contains procedure calls related to shares on the ARX.
- VirtualService is a group of procedures related to global servers, virtual servers, cifs services, and nfs services on the ARX.
- Volume contains procedure calls related to volumes on the ARX.

The statistics clear whenever the chassis reboots. You can use the clear statistics api command to clear them manually.

Guidelines: Setting Up File-Change Notifications for the API The procedure calls under FileChangeNotifications require some preparation. For each managed volume where you want to track file changes, create a notification rule. This rule creates snapshots of the volume's metadata and all of its shares. The API monitors the volume over time, and it probes these snapshots as needed to fulfill its procedure calls.

Guidelines: API

To activate the ARX API, use the management access command followed by the permit command. This opens a port where a SOAP client can send queries to this volume. The HTTP-API port is 83, and the HTTPS-API port is 843.

Use the following URL syntax to access the API documentation through HTTP:

http://*arx-management-ip*:83/arx-api/

where *arx-management-ip* is either the out-of-band management-IP address (interface mgmt) or an in-band (VLAN) management-ip address (interface vlan). The permit command determines which of these address types are available for access. Use show interface mgmt and/or show interface vlan to find the IP addresses for each interface.

This URL accesses the same documentation through HTTPS:

https://*arx-management-ip*:843/arx-api/

where *arx-management-ip* is a valid management IP, as described above.

You can use this command, show statistics api, to find usage statistics for the API.

Sample

`provA# show statistics api`

shows all API statistics for the "provA" system. See Figure 34.4, below, for sample output.

Related Commands

clear statistics api
notification rule
management access
permit

Figure 34.4 *Sample Output: show statistics api*

```
provA# show statistics api

API Calls                                Total
-----------------------------------------------
Api:
    get_version                            0
Chassis:
    get_ha_peer                            0
    get_ha_status                          0
    get_health                             0
    get_hostname                           0
    get_hwversion                          0
    get_model                              0
    get_processor_stats                    0
    get_serial                             0
    get_storage_status                     0
    get_swversion                          0
    get_type                               0
Export:
    get_configuration                      0
    get_definition                         0
    get_definition2                        0
    get_list                               0
    get_list2                              0
    get_statistics                         0
    get_status                             0
FileChangeNotification:
    finish_notifications                   0
    get_notifications                      0
```

```
    get_snapshots                   0
    get_snapshot_definition         0
    start_notifications             0
FileServer:
    get_configuration               0
    get_definition                  0
    get_list                        0
    get_status                      0
ManualMigrateRule:
    create                          0
    create_and_configure            0
    get_configuration               0
    get_list                        0
    get_statistics                  0
    get_status                      0
    migrate_files                   0
    remove                          0
    set_enable                      0
    set_migrate_close_file          0
    set_report                      0
Namespace:
    get_configuration               0
    get_definition                  0
    get_list                        0
Network:
    get_definition                  0
    get_list                        0
    get_stats                       0
Policy:
    get_definition                  0
    get_list                        0
    get_schedule                    0
Report:
    get_list                        0
    get_report                      0
    get_status                      0
Schedule:
    create                          0
    create_and_configure            0
    get_configuration               0
    get_list                        0
    get_status                      0
    remove                          0
    set_description                 0
    set_duration                    0
    set_interval                    0
    set_start                       0
    set_stop                        0
Share:
    get_configuration               0
    get_definition                  0
    get_definition2                 0
    get_list                        0
    get_statistics                  0
    get_status                      0
VirtualService:
    get_cifs_open_files             0
    get_cifs_user_sessions          0
    get_configuration               0
    get_definition                  0
    get_definition2                 0
    get_list                        0
```

```
    get_statistics                      0
    get_stats                           0
    get_status                          0
Volume:
    get_configuration                   0
    get_definition                      0
    get_definition2                     0
    get_list                            0
    get_status                          0
---------------------------------------------
Total                                   0
```

35

Disaster Recovery Between ARX Clusters

activate configs

Purpose Use the activate configs command to enable all of the volumes and services in a pre-loaded global-config file. If this site is a backup for another site, this is the final step in performing a disaster-recovery operation.

Mode priv-exec

Security Role(s) crypto-officer

Syntax

```
activate configs file-name [partial][take-ownership][tentative]
activate replicated-configs file-name ...
```

file-name (1-255 characters) is the name of the global-config file. You can activate a replicated global-config from a remote cluster or a locally-defined global-config file. Use show replicated-configs for a list of replicated-config files from other clusters, or use show configs for a list of locally-defined configuration files.

partial (optional) activates a subset of the full configuration. This can be any of these keywords, or any combination of them:

global-server *fqdn* (1-128 characters) activates a particular global-server configuration, along with all of its component objects (such as its namespace and all the shares behind it). This is the fully-qualified domain name (for example, "myserver.myorg.org") for one global server in the global-config. Use show replicated-configs *file-name* to view the global config and see all of its global servers.

shares activates all shares.

volumes activates all volumes.

global-servers activates all global servers.

policies activates all rules and share farms.

If you specify any of the above, the command activates only the specified option(s). If you omit all of them, the command enables all global servers, volumes, and shares but does *not* activate any policy rules.

take-ownership (optional, but recommended) causes the command to use the take-ownership option for every managed-volume share that it enables (see the documentation for enable (gbl-ns-vol-shr)). The filer-replication process typically copies hidden files from the back-end shares at the source site, including files that mark the back-end share as "already imported" by the source site's ARX. The managed volumes will not import these shares unless the take-ownership flag is raised.

tentative (optional) creates a report to describe the how the command runs, but does not actually enable any configuration objects. The CLI shows the name of the report after you issue the command. You can use show reports *report-name* to review the report and confirm that the command implements the configuration as desired.

activate replicated-configs *file-name* ... is an alternative syntax, with the same options as described above for activate configs.

Default(s) ***partial*** - **shares volumes global-servers**

Guidelines: Preparation for Disaster

This command is used in a disaster-recovery operation, which requires preparation before the disaster occurs.

Preparing Back-End Filers

You begin by using file-replication facilities (such as NetApp's SnapMirror, EMC's SRDF, or RoboCopy) to duplicate the storage from an active site to a backup site. Continuously replicate all of the back-end shares behind the ARX. Also create back-end shares for managed-volume metadata at the backup site, but you do not need to replicate the source site's metadata. The ARX software at the backup site will recreate the metadata there.

Then you begin the process of building a global configuration to be regularly copied from the active site to the backup site. The global configuration includes all storage-service parameters; see show global-config. The peers at the active site run the global configuration during normal service, and the peers at the backup site run it after a disaster.

Sharing Master Keys and Removing Global-Config from Backup Site

The ARX peers at both sites must share the same master key in order to share the same global configuration. The *master key* encrypts and decrypts all of the passwords and other Critical Security Parameters (CSPs) in the global-config file. Use the show master-key command to show an encrypted copy of the master key on the current ARX. The master keys should be synchronized during installation time, as described in the Hardware Installation manuals. If the backup peers have the wrong master key, reset the peers to factory defaults and reset the master key. They should only have their network parameters configured (show running-config), along with the correct master key. The documentation for show master-key provides details for this process.

Connecting Sites with a Full Mesh of RON Tunnels

At the ARX CLI, use the ron tunnel command to connect both ARX peers at the active site to both peers at the backup site. That is, each ARX must have two tunnels, one to each peer at the other site.

Removing All Storage Services from the Backup Cluster

We recommend removing all namespaces, global servers, and front-end services from the backup cluster before you begin. This prevents any significant configuration conflicts between the active site and the backup site, which could possibly prevent a successful disaster recovery. Ideally, the backup cluster should only have network parameters configured (which you can see with show running-config), along with administrative accounts (show group users) and any AD forests that have been discovered in the network (show active-directory).

If there are any storage services on the backup cluster, we recommend removing them. From the backup cluster's CLI, use remove service *namespace* for each namespace in the backup cluster's configuration. The remove service command removes the namespace configuration along with all associated external-filer definitions, and it also removes ARX metadata from the filers behind the namespace. Then use no global server, no cifs, and/or no nfs to remove all of the global servers and front-end services at the backup cluster.

Guidelines: Preparation for Disaster (Cont.)

Building a Shared Global-Config

Once all ARX peers have the same master key and are connected by RON tunnels, use the cluster-name command to establish a cluster name for each site. The cluster-name command also assigns each peer to a cluster. You then expand the active cluster's global configuration so that it can also include the backup cluster's file servers and Virtual-IP addresses:

- Add a new external-filer object for each of the backup site's filers. After you define them, the active cluster has definitions for all the active site's filers and all the backup site's filers. With both sets of filers defined, you can use the following commands to associate each filer to its role in each cluster:
  - metadata share *ext-filer* ... cluster *n* chooses *ext-filer* to host a metadata share in cluster *n*. You run the command again to choose a different filer for this metadata share at the other cluster.
  - filer *ext-filer* ... cluster *n* chooses a filer to host a storage share in a particular cluster. As above, you run the command twice for each share; once to choose the filer at the active cluster, and another time to choose the filer at the backup cluster.
  - sam-reference *ext-filer* ... cluster *n* chooses *ext-filer* to play the role of a SAM-reference filer in cluster *n*. You run the command again to choose a different filer for SAM at the other cluster.
- Use the virtual server ... *vip* ... cluster *n* command to set the *vip* (Virtual-IP) addresses for use at cluster *n*. This brings you into a submode, with several commands for setting CIFS-service names. The submode, and those service names, are the same for both clusters; if the submode's commands were already set for the active cluster, you can exit the mode without repeating them for the backup cluster.

Regularly Copying the Global-Config to the Backup Cluster

This prepares the active cluster's global config so that it is valid at either the active cluster or the backup cluster. Then create a config-replication rule to regularly copy that global config from the active cluster to the backup cluster.

Preparing for Snapshot Reconstitution

If you use any snapshot rules to back up client data, and if the filer-replication application copies the snapshots to the backup site, we advise that you also prepare for *snapshot reconstitution*. You prepare for snapshot reconstitution by regularly copying snapshot reports off of the active cluster, along with a Perl script in the configs directory. The guidelines for snapshot rule describe this preparation process in detail: see Guidelines: Preparing for Snapshot Reconstitution, on page 30-50.

Guidelines: Preparation for Disaster (Cont.)

Testing the "Load Configs" Operation at the Backup Cluster

After the global-config is successfully copied to the backup cluster, we recommend an occasional test of the load configs command at that cluster. This loads all of the active cluster's storage services into the database without enabling any shares or volumes. It also loads the Active Directory configuration, which may not be correct for the backup cluster; after the load operation, run active-directory update seed-domain there to confirm that the AD configuration is correct.

If you are running CIFS services with constrained delegation, you require some additional configuration at the Domain Controller (DC) to support your CIFS clients. The back-end CIFS servers at the backup cluster must be on the "delegate-to" list for each front-end CIFS service; for each cifs service, run the probe delegate-to command to find any back-end servers that need to be added to the list, and add them at the DC. When you finish, the CIFS service's "delegate-to" list includes the back-end servers behind *both* clusters. This configuration occurs at the DC, and persists after this load-configs test.

Check the remaining configuration to confirm that it is correct. Then use remove service to remove each of the namespaces that you added with the load configs command. Finally, use no global server, no cifs, and/or no nfs to remove all of the global servers and front-end services.

Guidelines: Disaster Recovery

To recover from a disaster, go to the backup cluster and run load configs on the global-config file that was copied earlier. This loads the global-config into the database without enabling any shares, volumes, global servers, or policy rules. You can use this command, activate configs, to enable any or all of those objects.

Unless you use the tentative option, the CLI prompts for confirmation before activating the global-config file. Enter **yes** to proceed.

The best practice is to activate the configuration in stages and verify the success of each stage. Use the following command order, as time allows:

- activate configs ... shares take-ownership, then
- activate configs ... volumes. Periodically use show namespace status to verify that the volumes come online, and that there are no issues blocking your imports. You can use show namespace for detailed status.
- activate configs ... global-servers. You can use show cifs-service and/or show nfs-service to confirm that client services are running. Also verify that you can connect to the global server(s) as a client.
- activate configs ... policy. (The policy keyword is required for the activate configs command to enable any rules.) Use show policy to confirm that all rules are migrating or otherwise functioning properly.

Alternatively, you can bring up the shares, volumes, and global servers all at once by omitting any specific keywords. After confirming that the services are running properly, you can use activate configs ... policy to activate all rules.

If you also prepared for snapshot reconstitution, play back the snapshot reconstitution script as described in Guidelines: Reconstituting ARX Snapshots, on page 30-50.

If the cluster at the failed site is still on the network after the disaster, log into the cluster and use no enable (gbl-cfg-repl) to stop copying the global configuration from the now-backup site.

Guidelines: Report

This command always creates a report to show its actions. The CLI shows the name of the report after you issue this command. You can use tail *report-name* follow to follow the report as the command runs, or show reports *report-name* to view the report after the command has finished.

The configuration file is an ordered list of CLI commands that recreate the global configuration, including all volumes and storage services. The report shows these CLI commands as it runs them. If it reaches a configuration object that already exists in the local database, it shows an INFO message to explain that the object is being ignored. An X appears at the beginning of each ignored line. This includes the enable command for each activated configuration object (such as a share or volume); the load configs command does not enable any configuration objects.

Samples

```
newptA# activate configs provSvcs.rcfg tentative

  Generating activate configuration report:
'act_provSvcs.rcfg_201007060305'.
```

performs a tentative activation of the global-config file, "provSvcs.rcfg."

```
newptA# activate configs provSvcs.rcfg shares take-ownership

This enables the global configuration.

Are you sure? [yes/no] yes

  Generating activate configuration report:
'aact_provSvcs.rcfg_20100709113618'.
```

activates all shares from the above configuration, and takes ownership of them. See Figure 35.1 for a sample report.

Related Commands

load configs
show global-config

Figure 35.1 *Sample Report: Activate Configs Report*

```
newptA# show reports act_provSvcs.rcfg_20100709113618.rpt

  Generating activate configuration report: 'act_provSvcs.rcfg_20100709113618'.

X newptA- terminal character-set unicode-utf-8
X newptA- global
X newptA(gbl)-   cluster-name providence member provA
X newptA(gbl)-   cluster-name newport member newptA
X newptA(gbl)-   user adm1 encrypted-password YyB4UgM33t1HemjAuV2A70ZEt3Rx0eG5LLPAhIYIOEo=

% INFO: Configuration object type not able to be activated.

X newptA(gbl-user[adm1])-     exit
X newptA(gbl)-   user adm12 encrypted-password
OIm2BgM33t1HemjAuV2A70ZEt3Rx0eG5bt0PfjQJhaN+DVVUSwXi5w==

% INFO: Configuration object type not able to be activated.

X newptA(gbl-user[adm12])-     exit
X newptA(gbl)-   user admin encrypted-password Sd/LKQM33t1HemjAuV2A70ZEt3Rx0eG5O91ib1WbGcs=
```

```
% INFO: Configuration object type not able to be activated.

X newptA(gbl-user[admin])-     exit
X newptA(gbl)-   user newadmin encrypted-password
HmTKsQM33t1HemjAuV2A70ZEt3Rx0eG5zV2nUYRa+pGp6Sx7MRbxrg==

% INFO: Configuration object type not able to be activated.

X newptA(gbl-user[newadmin])-     exit
X newptA(gbl)-   group Administrators
...
X newptA(gbl)-   external-filer ntap-npt2

% INFO: Configuration object type not able to be activated.

X newptA(gbl-filer[ntap-npt2])-     description "New filer, rack 4 (Newport)"
X newptA(gbl-filer[ntap-npt2])-     ip address 10.51.100.173
X newptA(gbl-filer[ntap-npt2])-     exit
X newptA(gbl)-   external-filer ntap-nwpt

% INFO: Configuration object type not able to be activated.

X newptA(gbl-filer[ntap-nwpt])-     description "big filer in Newport's front lab"
X newptA(gbl-filer[ntap-nwpt])-     ip address 10.51.100.172
X newptA(gbl-filer[ntap-nwpt])-     exit
X newptA(gbl)-   external-filer ntap-pr-1

% INFO: Configuration object type not able to be activated.

X newptA(gbl-filer[ntap-pr-1])-     cifs-port 445
X newptA(gbl-filer[ntap-pr-1])-     description "New NAS filer, rack 5"
X newptA(gbl-filer[ntap-pr-1])-     ip address 192.168.103.52
X newptA(gbl-filer[ntap-pr-1])-     exit
X newptA(gbl)-   external-filer ntap-prov

% INFO: Configuration object type not able to be activated.

X newptA(gbl-filer[ntap-prov])-     cifs-port 445
X newptA(gbl-filer[ntap-prov])-     description "NAS filer in computer lab, rack 2"
X newptA(gbl-filer[ntap-prov])-     ip address 192.168.103.7
X newptA(gbl-filer[ntap-prov])-     exit
X newptA(gbl)-   schedule oncePerDay

% INFO: Configuration object type not able to be activated.

X newptA(gbl-schedule[oncePerDay])-     every 1 days
X newptA(gbl-schedule[oncePerDay])-     start 07/09/2010:11:22:00
...
X newptA(gbl)-   namespace provMed
X newptA(gbl-ns[provMed])-     protocol cifs
X newptA(gbl-ns[provMed])-     cifs anonymous-access
X newptA(gbl-ns[provMed])-     cifs authentication kerberos
X newptA(gbl-ns[provMed])-     cifs authentication ntlm
X newptA(gbl-ns[provMed])-     cifs authentication ntlmv2
X newptA(gbl-ns[provMed])-     cifs filer-signatures
X newptA(gbl-ns[provMed])-     proxy-user prov_admin
X newptA(gbl-ns[provMed])-     windows-mgmt-auth mmc-admins
X newptA(gbl-ns[provMed])-     sam-reference ntap-pr-1
X newptA(gbl-ns[provMed])-     volume /mds
X newptA(gbl-ns-vol[provMed~/mds])-       filer-subshares native-names-only
X newptA(gbl-ns-vol[provMed~/mds])-       modify
X newptA(gbl-ns-vol[provMed~/mds])-       reimport-modify
```

```
X newptA(gbl-ns-vol[provMed~/mds])-       reserve files 4000000
X newptA(gbl-ns-vol[provMed~/mds])-       auto sync files
X newptA(gbl-ns-vol[provMed~/mds])-      metadata share ntap-pr-1 nfs3 /vol/vol3/arx_meta cluster
providence
X newptA(gbl-ns-vol[provMed~/mds])-       metadata share ntap-npt2 nfs3 /vol/vol3/arxMeta cluster
newport
X newptA(gbl-ns-vol[provMed~/mds])-       no compressed-files
X newptA(gbl-ns-vol[provMed~/mds])-       named-streams
X newptA(gbl-ns-vol[provMed~/mds])-       no persistent-acls
X newptA(gbl-ns-vol[provMed~/mds])-       no sparse-files
X newptA(gbl-ns-vol[provMed~/mds])-       unicode-on-disk
X newptA(gbl-ns-vol[provMed~/mds])-       share mds
X newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         description "home dirs for new doctors"
X newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         import priority 1
X newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         import skip-managed-check
X newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         import sync-attributes
X newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         filer ntap-pr-1 cifs MDS cluster providence
X newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         filer ntap-npt2 cifs MDS cluster newport
X newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         enable
X newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         exit
X newptA(gbl-ns-vol[provMed~/mds])-       share surgeons
X newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         description "home dirs for veteran
surgeons"
X newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         import skip-managed-check
X newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         import sync-attributes
X newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         filer ntap-prov cifs SURGEONS cluster
providence
X newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         filer ntap-nwpt cifs SURGEONS cluster
newport
X newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         enable
X newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         exit
X newptA(gbl-ns-vol[provMed~/mds])-       place-rule pop-md-files-2-new
X newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         schedule oncePerDay
X newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         report tier1 error-only
X newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         inline report daily tier1
error-only
X newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         from fileset popular-files
X newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         target share mds
X newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         enable
X newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         exit
X newptA(gbl-ns-vol[provMed~/mds])-       place-rule unpop-md-files-2-old
X newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-          schedule oncePerDay
X newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-          from fileset unpopular-files
X newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-          target share surgeons
X newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-          no inline notify
X newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-          enable
X newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-          exit
X newptA(gbl-ns-vol[provMed~/mds])-       place-rule masters-2-new
X newptA(gbl-ns-vol-plc[provMed~/mds~masters-2-new])-         from fileset all-files match
directories promote-directories
X newptA(gbl-ns-vol-plc[provMed~/mds~masters-2-new])-         target share mds
X newptA(gbl-ns-vol-plc[provMed~/mds~masters-2-new])-         enable
X newptA(gbl-ns-vol-plc[provMed~/mds~masters-2-new])-         exit
X newptA(gbl-ns-vol[provMed~/mds])-       volume-group 1
X newptA(gbl-ns-vol[provMed~/mds])-       enable
X newptA(gbl-ns-vol[provMed~/mds])-       exit
X newptA(gbl-ns[provMed])-     volume /rns
X newptA(gbl-ns-vol[provMed~/rns])-       filer-subshares native-names-only
X newptA(gbl-ns-vol[provMed~/rns])-       modify
X newptA(gbl-ns-vol[provMed~/rns])-       reimport-modify
X newptA(gbl-ns-vol[provMed~/rns])-       reserve files 4000000
X newptA(gbl-ns-vol[provMed~/rns])-       auto sync files
```

```
X newptA(gbl-ns-vol[provMed~/rns])-      metadata share ntap-pr-1 nfs3 /vol/vol4/arx_meta cluster
providence
X newptA(gbl-ns-vol[provMed~/rns])-       metadata share ntap-npt2 nfs3 /vol/vol4/arxMeta cluster
newport
X newptA(gbl-ns-vol[provMed~/rns])-       no compressed-files
X newptA(gbl-ns-vol[provMed~/rns])-       named-streams
X newptA(gbl-ns-vol[provMed~/rns])-       no persistent-acls
X newptA(gbl-ns-vol[provMed~/rns])-       no sparse-files
X newptA(gbl-ns-vol[provMed~/rns])-       unicode-on-disk
X newptA(gbl-ns-vol[provMed~/rns])-       share nurses
X newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         description "home dirs for veteran
surgeons"
X newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         import skip-managed-check
X newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         import sync-attributes
X newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         filer ntap-prov cifs NURSES cluster
providence
X newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         filer ntap-nwpt cifs NURSES cluster newport
X newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         enable
X newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         exit
X newptA(gbl-ns-vol[provMed~/rns])-       share rns
X newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         description "home dirs for new nurses"
X newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         import priority 1
X newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         import skip-managed-check
X newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         import sync-attributes
X newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         filer ntap-pr-1 cifs RNS cluster providence
X newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         filer ntap-npt2 cifs RNS cluster newport
X newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         enable
X newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         exit
X newptA(gbl-ns-vol[provMed~/rns])-       place-rule pop-rn-files-2-new
X newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         schedule oncePerDay
X newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         report tier1 error-only
X newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         inline report daily tier1
error-only
X newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         from fileset popular-files
X newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         target share rns
X newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         enable
X newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         exit
X newptA(gbl-ns-vol[provMed~/rns])-       place-rule unpop-rn-files-2-old
X newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-          schedule oncePerDay
X newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-          from fileset unpopular-files
X newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-          target share nurses
X newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-          no inline notify
X newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-          enable
X newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-          exit
X newptA(gbl-ns-vol[provMed~/rns])-       place-rule masters-2-new
X newptA(gbl-ns-vol-plc[provMed~/rns~masters-2-new])-         from fileset all-files match
directories promote-directories
X newptA(gbl-ns-vol-plc[provMed~/rns~masters-2-new])-         target share rns
X newptA(gbl-ns-vol-plc[provMed~/rns~masters-2-new])-         enable
X newptA(gbl-ns-vol-plc[provMed~/rns~masters-2-new])-         exit
X newptA(gbl-ns-vol[provMed~/rns])-       volume-group 1
X newptA(gbl-ns-vol[provMed~/rns])-       enable
X newptA(gbl-ns-vol[provMed~/rns])-       exit
X newptA(gbl-ns[provMed])-     exit
X newptA(gbl)-   namespace provMed
X newptA(gbl-ns[provMed])-     protocol cifs
X newptA(gbl-ns[provMed])-     cifs anonymous-access
X newptA(gbl-ns[provMed])-     cifs authentication kerberos
X newptA(gbl-ns[provMed])-     cifs authentication ntlm
X newptA(gbl-ns[provMed])-     cifs authentication ntlmv2
X newptA(gbl-ns[provMed])-     cifs filer-signatures
X newptA(gbl-ns[provMed])-     proxy-user prov_admin
```

```
X newptA(gbl-ns[provMed])-     windows-mgmt-auth mmc-admins
X newptA(gbl-ns[provMed])-     sam-reference ntap-pr-1
X newptA(gbl-ns[provMed])-     exit
X newptA(gbl)-   global server provmed.MEDARCH.ORG
X newptA(gbl-gs[provmed.MEDARCH.ORG])-                description "CIFS service - home dirs"
X newptA(gbl-gs[provmed.MEDARCH.ORG])-                windows-domain MEDARCH.ORG
X newptA(gbl-gs[provmed.MEDARCH.ORG])-                active-directory proxy-user prov_admin
X newptA(gbl-gs[provmed.MEDARCH.ORG])-                virtual server newptA 192.168.74.61
255.255.0.0 vlan 103 cluster newport
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                 wins 192.168.25.102
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                 active-directory alias
pubs_test_provmed
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                 active-directory alias
pubs_test_ntap-prov
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                 active-directory alias
pubs_test_ntap-pr-1
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                 active-directory alias
pubs_test_provmed.MEDARCH.ORG
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                 active-directory alias
pubs_test_ntap-prov.MEDARCH.ORG
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                 active-directory alias
pubs_test_ntap-pr-1.MEDARCH.ORG
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                 enable
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                 exit
X newptA(gbl-gs[provmed.MEDARCH.ORG])-                virtual server provA 192.168.74.91
255.255.0.0 vlan 103 cluster providence

% INFO: Cluster 'providence' is not the local cluster. This configuration option will not take
effect on this switch.

X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 wins 192.168.25.102
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_provmed
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_ntap-prov
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_ntap-pr-1
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_provmed.MEDARCH.ORG
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_ntap-prov.MEDARCH.ORG
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_ntap-pr-1.MEDARCH.ORG
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_provmed
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_ntap-prov
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_ntap-pr-1
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_provmed.MEDARCH.ORG
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_ntap-prov.MEDARCH.ORG
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_ntap-pr-1.MEDARCH.ORG
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 enable
X newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 exit
X newptA(gbl-gs[provmed.MEDARCH.ORG])-                enable
X newptA(gbl-gs[provmed.MEDARCH.ORG])-                exit
X newptA(gbl)-   cifs provmed.MEDARCH.ORG
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                browsing provMed
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                description "user home directories"
```

```
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-              kerberos-creds provmed.MEDARCH.ORG
MEDARCH.ORG            ECqJ5gJ0zkU9hskVeO7a2w== provmed$ HOST/provmed.MEDARCH.ORG            2
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                signatures
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                dynamic-dns pubs_test_provmed
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                dynamic-dns pubs_test_ntap-prov
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                dynamic-dns pubs_test_ntap-pr-1
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                export provMed /mds as MDS description
"home dirs for the docs"
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                export provMed /rns as RNS description
"home dirs for the nurses"
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                export provMed /mds/blake as BLAKE
filer-subshare
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                export provMed /mds/hogtroff as HOGTROFF
filer-subshare
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                export provMed /mds/mccoy as MCCOY
filer-subshare
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                export provMed /mds/mhoward as MHOWARD
filer-subshare
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                export provMed /mds/pierce as PIERCE
filer-subshare
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                enable
X newptA(gbl-cifs[provmed.MEDARCH.ORG])-                exit
X newptA(gbl)-   config-replication prov2newport cluster providence

% INFO: Cluster 'providence' is not the local cluster. This configuration option will not take
effect on this switch.

% INFO: Configuration object type not able to be activated.

X newptA(gbl-cfg-repl[prov2newport])-     description "send service config to Newport site"
X newptA(gbl-cfg-repl[prov2newport])-     report gblRepl2newport
X newptA(gbl-cfg-repl[prov2newport])-     schedule oncePerDay
X newptA(gbl-cfg-repl[prov2newport])-     target-cluster newport
X newptA(gbl-cfg-repl[prov2newport])-     target-file provSvcs.rcfg
X newptA(gbl-cfg-repl[prov2newport])-     user admin encrypted-password
Sd/LKQM33t1HemjAuV2A70ZEt3Rx0eG5O91ib1WbGcs=
X newptA(gbl-cfg-repl[prov2newport])-     enable
X newptA(gbl-cfg-repl[prov2newport])-     exit
X newptA(gbl)-   exit
```

cluster-name

Purpose

A *cluster* is single ARX site, comprised of a redundant pair of ARX systems or a standalone ARX. Use this command to declare a cluster name for any ARX (or redundant-ARX pair) in the current Resilient Overlay Network (RON). This name is required for configuring disaster recovery between two ARX sites.

Use no cluster-name to remove a cluster name from the ARX configuration.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
cluster-name cluster-name member peer-a [member peer-b]
no cluster-name cluster-name member peer-a [member peer-b]
```

cluster-name (1-64 characters) is the name you choose for this cluster.

peer-a (1-64 characters) is an ARX that you assign to the above cluster. Each ARX can be assigned to a single cluster only. The ARX must already be on the same RON as the current switch; you can use the show ron command for a full list of all switches on the RON.

peer-b (1-64 characters) is the redundant peer to the above ARX, if there is one. You can use the show ron *peer-a* command to find if *peer-a* has a redundant peer.

Default(s)

None.

Guidelines

This command establishes an ARX site as a *cluster*, capable of failing its global-configuration over to another cluster. You can create two clusters on your RON and set them up in a disaster-recovery configuration, so that one can take over service for the other in the event of a catastrophic failure. The ron tunnel command creates a RON tunnel between any two ARX systems on the same WAN; use that command to create a full mesh of RON tunnels between all four peers. Additionally, ensure that all four peers have the same master key, so that they can share their global configurations (see show master-key for details).

To set up disaster recovery between two clusters, you designate one cluster as "active" and the other as the "backup" site. Use filer-replication facilities (for example, NetApp's SnapMirror or EMC's SRDF) to copy all files and directories from the active site's back-end filers to their counterparts at the backup site. Once the files and directories are synchronized, add all of the backup cluster's external-filers to the active cluster's configuration. Use the filer, metadata share, and sam-reference commands to establish how those filers should be used in the backup cluster. Then use the virtual server command to establish a new Virtual-IP address and configuration for each global server in the backup cluster. Finally, create a config-replication rule to regularly copy this configuration from the active cluster to the backup cluster. This prepares the backup site to take over service from the active site.

If the active site fails, you can then resume your storage services at the backup cluster. To accomplish this, go to the backup cluster and use load configs to load the active site's configuration into the backup cluster's database, then use activate configs to enable all volumes and services at the backup cluster.

Guidelines: Only One Metadata Share Per Managed Volume

A cluster allows the configuration of a single metadata share per managed volume. If any managed volume has more than one metadata share, this command fails. A single metadata share per volume is required to ensure that the disaster-recovery software can reliably re-create the volume configuration on a backup cluster.

Guidelines: Removing Cluster Configuration

To remove a remote cluster along with all its associated filers, metadata shares, and virtual servers, use the remove cluster-config command. You can use the same command on the local cluster to remove the cluster object and all references to it, without actually removing the above objects. The no cluster-name command removes only the cluster definition; the CLI rejects this operation if the cluster name is referenced by any of the above objects.

Sample

```
provA(gbl)# cluster-name providence member provA provB
provA(gbl)# cluster-name newport member newptA newptB
```

creates two new clusters, “providence” and “newport.” You can now set up one as a backup site for the other.

Related Commands

config-replication
load configs
activate configs
remove cluster-config

config-replication

Purpose

Use the config-replication command to create a rule for replicating the global configuration to an ARX cluster at another site. A *cluster* is a redundant pair of ARX systems. This prepares for a disaster-recovery scenario. If your back-end filers mirror their data between the sites, this rule mirrors the configuration data for managing those filers at each site.

Use the no form of the command to delete the rule.

Modes

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
config-replication name cluster cluster-name
no config-replication name cluster cluster-name
```

name (1-32 characters) is the name you choose for the rule.

cluster-name (1-64 characters) is the name of the source cluster. This cluster must be defined in the database; you can use the show cluster command for a complete list of pre-defined clusters.

Default(s)

None

Guidelines: Preparing to Share Configs Between Clusters

Before you copy the global configuration file from one cluster to another, the cluster members require require matching master keys and a full mesh of RON tunnels.

The *master key* for an ARX encrypts and decrypts the passwords in its global-config file, so it must be the same for all switches that share the global-config. The backup switches should have the same master key as their active counterparts; the master keys should be set accordingly when the backup switches are installed. Refer to the documentation for the show master-key command for information on showing the key and possibly changing it on the backup switches.

Every switch in the active cluster must have a ron tunnel to every switch in the backup cluster. This provides multiple pathways for copying the global configuration between clusters.

The backup cluster should not have any storage services configured, so that it accepts the latest storage-services configuration from the active cluster without any conflicts. If the backup cluster has any namespaces, use remove service to cleanly remove each of them. Also use no global server, no cifs, and/or no nfs to remove all of the global servers and front-end services at the backup site.

Guidelines

The config-replication command creates a rule to regularly copy the active cluster's global configuration and send it to the backup cluster. The command puts you into gbl-cfg-repl mode, where you have various configuration commands for configuring the config-replication rule. Use the target-cluster command to choose the ARX cluster to receive this configuration. Then use the user (gbl-cfg-repl) command to enter your administrative credentials at the remote cluster; this gives you permission to write the configuration file there. The target-file command provides a name for the file copy. Use the schedule (gbl-cfg-repl) command to choose a regular schedule for duplicating the configuration file. We recommend using the report (gbl-cfg-repl) command to create a report about each config-replication event. You can also use the description (gbl-cfg-repl) command to describe this rule in the output of show policy. Finally, use the enable (gbl-cfg-repl) command to start the rule.

Guidelines: Schedule

To create a schedule to be applied to the config-replication rule, use the gbl schedule command at the source cluster.

The rule fires as dictated by its schedule: if the schedule has a start time that is earlier than now, the first config replication begins as soon as you enable the rule.

Guidelines: Reports

As mentioned above, we recommend that you use the config-replication report feature to keep a detailed log of all config-replication events. The reports appear on the source cluster: the show reports command shows all reports on the switch, including config-replication reports. You can use the standard file-management commands with these reports: delete, rename, show reports file-name, tail, and/or grep.

Guidelines: Preparing for Disaster Recovery

The global-config file replicated by this rule must support both clusters before you can use it to recover from a disaster. Expand the active cluster's global configuration so that it can also include the backup cluster's file servers and Virtual-IP addresses:

- Add a new external-filer object for each filer at the backup site. With both sets of filers defined on the active cluster, you can use the following commands to associate each filer to its role in each cluster:
  - metadata share *ext-filer* ... cluster *n* chooses *ext-filer* to host a metadata share in cluster *n*. You run the command again to choose a different filer for this metadata share at the other cluster.
  - filer *ext-filer* ... cluster *n* chooses a filer to host a storage share in a particular cluster. As above, you run the command twice for each share; once to choose the filer at the active cluster, and another time to choose the filer at the backup cluster.
  - sam-reference *ext-filer* ... cluster *n* chooses *ext-filer* to play the role of a SAM-reference filer in cluster *n*. You run the command again to choose a different filer for SAM at the other cluster.
- Use the virtual server ... *vip* ... cluster *n* command to set the *vip* (Virtual-IP) addresses for use at cluster *n*. This brings you into a submode, with several commands for setting CIFS-service names. The submode, and those service names, are the same for both clusters.

This prepares the active cluster's global config so that it is valid at either the active cluster or the backup cluster.

Sample

```
provA(gbl)# config-replication prov2newport cluster providence
```

instantiates a new config-replication rule, "prov2newport."

Related Commands

target-cluster
user (gbl-cfg-repl)
target-file
schedule (gbl-cfg-repl)
report (gbl-cfg-repl)
description (gbl-cfg-repl)
enable (gbl-cfg-repl)
schedule
show reports

description (gbl-cfg-repl)

Purpose Use the optional description command to set a descriptive string for the current config-replication rule. This appears in show commands.

Use the no form of the command to delete the description.

Mode gbl-cfg-repl

Security Role(s) storage-engineer or crypto-officer

Syntax

```
description text
no description
```

text (1-48 characters) is your description. Surround the text with quotation marks (“”) if it contains any spaces.

Default(s)

```
no description
```

Guidelines The description appears in the output for show config-replication.

Sample

```
provA(gbl-cfg-repl[prov2newport])# description “send service config to
Newport site”
```

specifies a description for the current config-replication rule.

Related Commands config-replication
show config-replication

enable (gbl-cfg-repl)

Purpose Use the enable command to enable the current config-replication rule.

Use no enable to disable the rule.

Modes gbl-cfg-repl

Security Role(s) storage-engineer or crypto-officer

Syntax

```
enable
no enable
```

Default(s)

```
no enable
```

Guidelines You must enable the config-replication rule for the policy engine to use it. Once enabled, the config-replication rule follows its assigned schedule (set with the schedule (gbl-cfg-repl) command) and copies its global-config file to the target ARX cluster (target-cluster).

The no form of the command is useful in a formerly-active cluster that is still on the network. Use the no enable command to stop the cluster at the failed site from copying its configuration to the newly-active site.

Samples

```
provA(gbl-cfg-repl[prov2newport])# enable
```

enables the “prov2newport” rule.

```
bstnA(gbl-cfg-repl[dr-test])# no enable
```

disables the “dr-test” rule.

Related Commands config-replication

load configs

Purpose

The *global-config* parameters are shared among both ARXes in a redundant pair. (The traditional running-config applies only to the local switch.) You can copy the global-config from one redundant pair to another, so that the second redundant pair (called an ARX *cluster*) can act as a backup for the first. Use this command, load configs, to load a global-config into the local database without enabling any configuration objects.

This is the first step in disaster recovery; this cluster can use the loaded configuration to take over services from the failed cluster.

Mode

priv-exec

Security Role(s)

crypto-officer

Syntax

`load configs` ***`file-name`*** **`[global-server`** ***`fqdn`*****`] [tentative]`**
`load replicated-configs` ***`file-name`*** **`...`**

file-name (1-255 characters) is the name of the global-config file. You can load a replicated global-config from a remote cluster or a locally-defined global-config file. Use show replicated-configs for a list of replicated-config files from other clusters, or use show configs for a list of locally-defined configuration files.

fqdn (optional, 1-128 characters) selects a particular global-server configuration, along with all of its component objects (such as a namespace and all the external-filers behind it). This is the fully-qualified domain name (for example, "myserver.myorg.org") for one global server in the global-config. Use show replicated-configs *file-name* to view the global config and see all of its global servers.

tentative (optional) creates a report to describe the how the command runs, but does not actually commit any configuration changes to the database. The CLI shows the name of the report after you issue the command. You can use show reports *report-name* to review the report and confirm that the command implements the configuration as desired.

load replicated-configs *file-name* ... is an alternative syntax, with the same options as described above for load configs.

Default(s)

None

Guidelines

After proper preparations have been made for disaster recovery (see the documentation for the config-replication command), you can use this command as a first step toward recovering from a disaster at a remote site. After you run this command to load the configuration, use activate configs to enable it and start client services at the current cluster.

Guidelines: Reports

This command always creates a report to show its actions. The CLI shows the name of the report after you issue this command. You can use tail *report-name* follow to follow the report as the command runs, or show reports *report-name* to view the report after the command has finished.

The configuration file is an ordered list of CLI commands that recreate the global configuration, including all volumes and storage services. The report shows these CLI commands as it runs them. For every configuration object that the load operation ignores, an INFO message appears before the object to explain why it was skipped. An X appears at the beginning of each ignored line.

Guidelines: Rules for the Load Operation

The load operation skips configuration objects that meet any of the following criteria:

- A configuration object of the same name already exists on this ARX. The load operation does not overwrite existing configuration objects.
- The exception to the above rule is the namespace configuration. The load operation merges namespaces by creating any volumes, shares, and rules that are not already present in the local configuration.

Guidelines: Loading Configurations from Higher-Rev Systems

If you use a configuration from a newer release of ARX software, edit the script before you load it. Command syntax may have changed as well as the best practices for command order. Consult this manual (and other documentation from the current release) for the correct syntax, and change the configuration file's commands as needed.

The load operation prompts for confirmation if it finds that the configuration is for a newer release. After you edit the configuration, enter **yes** to proceed.

Guidelines: Recommendations

We recommend running this command on a cluster without any storage services in its global configuration. If any object in the global-config file has the same name as an object that already exists, the load operation typically does *not* change the object that already exists. Therefore, if you load a configuration onto a backup cluster and leave it there, change the configuration on the active cluster, and copy the changed configuration to the backup cluster to load it again, many (possibly all) of the configuration changes are ignored at the backup cluster.

If you run a test load before any disaster, we recommend using remove service for each loaded service after you finish. Also use no global server, no cifs, and/or no nfs to remove all of the global servers and front-end services.

Contact F5 Support for advice if you wish to retain different global-config parameters at the backup site.

Samples

```
newptA# load configs provSvcs.rcfg tentative

  Generating load configuration report:
'load_provSvcs.rcfg_201007010258'.
```

performs a tentative load of the global-config file, "provSvcs.rcfg."

```
newptA# load config provSvcs.rcfg

  Generating load configuration report:
'load_provSvcs.rcfg_201007091136'.
```

loads the same configuration file into the local database. A sample report appears in Figure 35.2.

Related Commands

config-replication
activate configs
show replicated-configs

Figure 35.2 *Sample Report: Load Configs Report*

```
newptA# show reports load_provSvcs.rcfg_20100709113148.rpt

  Generating load configuration report: 'load_provSvcs.rcfg_20100709113148'.

X newptA- terminal character-set unicode-utf-8
X newptA- global
newptA(gbl)-   cluster-name providence member provA
newptA(gbl)-   cluster-name newport member newptA
newptA(gbl)-   user adm1 encrypted-password YyB4UgM33t1HemjAuV2A70ZEt3Rx0eG5LLPAhIYIOEo=

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-user[adm1])-     exit
newptA(gbl)-   user adm12 encrypted-password
OIm2BgM33t1HemjAuV2A70ZEt3Rx0eG5bt0PfjQJhaN+DVVUSwXi5w==

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-user[adm12])-     exit
newptA(gbl)-   user admin encrypted-password Sd/LKQM33t1HemjAuV2A70ZEt3Rx0eG5O91ib1WbGcs=

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-user[admin])-     exit
newptA(gbl)-   user newadmin encrypted-password
HmTKsQM33t1HemjAuV2A70ZEt3Rx0eG5zV2nUYRa+pGp6Sx7MRbxrg==

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-user[newadmin])-     exit
newptA(gbl)-   group Administrators

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-group[Administrators])-     role backup-operator
```

```
X newptA(gbl-group[Administrators])-     role crypto-officer
X newptA(gbl-group[Administrators])-     role network-engineer
X newptA(gbl-group[Administrators])-     role operator
X newptA(gbl-group[Administrators])-     role storage-engineer
X newptA(gbl-group[Administrators])-     windows-domain MEDARCH.ORG
X newptA(gbl-group[Administrators])-     exit
newptA(gbl)-   group "Backup Operators"

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-group[Backup Operators])-     role backup-operator
X newptA(gbl-group[Backup Operators])-     role operator
X newptA(gbl-group[Backup Operators])-     windows-domain MEDARCH.ORG
X newptA(gbl-group[Backup Operators])-     exit
newptA(gbl)-   group "Domain Admins"

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-group[Domain Admins])-     role backup-operator
X newptA(gbl-group[Domain Admins])-     role crypto-officer
X newptA(gbl-group[Domain Admins])-     role network-engineer
X newptA(gbl-group[Domain Admins])-     role operator
X newptA(gbl-group[Domain Admins])-     role storage-engineer
X newptA(gbl-group[Domain Admins])-     windows-domain MEDARCH.ORG
X newptA(gbl-group[Domain Admins])-     exit
newptA(gbl)-   group "Domain Users"

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-group[Domain Users])-     role operator
X newptA(gbl-group[Domain Users])-     windows-domain MEDARCH.ORG
X newptA(gbl-group[Domain Users])-     exit
newptA(gbl)-   group "Enterprise Admins"

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-group[Enterprise Admins])-     role backup-operator
X newptA(gbl-group[Enterprise Admins])-     role crypto-officer
X newptA(gbl-group[Enterprise Admins])-     role network-engineer
X newptA(gbl-group[Enterprise Admins])-     role operator
X newptA(gbl-group[Enterprise Admins])-     role storage-engineer
X newptA(gbl-group[Enterprise Admins])-     windows-domain MEDARCH.ORG
X newptA(gbl-group[Enterprise Admins])-     exit
newptA(gbl)-   group admins

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-group[admins])-     role storage-engineer
X newptA(gbl-group[admins])-     user adm1
X newptA(gbl-group[admins])-     user adm12
X newptA(gbl-group[admins])-     exit
newptA(gbl)-   group crypto-officer

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-group[crypto-officer])-     user admin
X newptA(gbl-group[crypto-officer])-     user newadmin
X newptA(gbl-group[crypto-officer])-     exit
newptA(gbl)-   group operator

% INFO: Configuration object already exists so ignoring new CLI commands.
```

```
X newptA(gbl-group[operator])-     user admin
X newptA(gbl-group[operator])-     user newadmin
X newptA(gbl-group[operator])-     user adm1
X newptA(gbl-group[operator])-     user adm12
X newptA(gbl-group[operator])-     exit
newptA(gbl)-   radius-server 192.168.25.201

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-radius[192.168.25.201])-     exit
newptA(gbl)-   radius-server 192.168.25.207

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-radius[192.168.25.207])-     auth-port 5555
X newptA(gbl-radius[192.168.25.207])-     retries 4
X newptA(gbl-radius[192.168.25.207])-     timeout 10
X newptA(gbl-radius[192.168.25.207])-     exit
newptA(gbl)-   ntlm-auth-db ntlmMap2
newptA(gbl-ntlm-auth-db[ntlmMap2])-     user lab encrypted-password
Azfe3Ud6aMC5XYDvRkS3dHHR4bkHAs2mghfmP1T0QWUDU3ax
newptA(gbl-ntlm-auth-db[ntlmMap2])-     exit
newptA(gbl)-   proxy-user acoProxy1

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-proxy-user[acoProxy1])-     description "jq's admin account"
X newptA(gbl-proxy-user[acoProxy1])-     user jqprivate encrypted-password
CHhf0AM33t1HemjAuV2A70ZEt3Rx0eG5Zo4yXOFoRo/qiA6xIQOgyw==
X newptA(gbl-proxy-user[acoProxy1])-     windows-domain WWMEDNET.COM pre-win2k-name WWMEDNET
X newptA(gbl-proxy-user[acoProxy1])-     exit
newptA(gbl)-   proxy-user acoProxy3

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-proxy-user[acoProxy3])-     user jqtesterencrypted-password
gsglrwM33t1HemjAuV2A70ZEt3Rx0eG5/383ySYcxkkm97vJ6rS19Q==
X newptA(gbl-proxy-user[acoProxy3])-     windows-domain FDTESTNET.COM pre-win2k-name BOSTONCIFS
X newptA(gbl-proxy-user[acoProxy3])-     exit
newptA(gbl)-   proxy-user ny_admin

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-proxy-user[ny_admin])-     user jqpublic  encrypted-password
CCNAZgM33t1HemjAuV2A70ZEt3Rx0eG5csH6+tzyZik=
X newptA(gbl-proxy-user[ny_admin])-     windows-domain WELLS.ME.ORG pre-win2k-name WELLS
X newptA(gbl-proxy-user[ny_admin])-     exit
newptA(gbl)-   proxy-user prov_admin
newptA(gbl-proxy-user[prov_admin])-     user administrator encrypted-password
Sd/LKQM33t1HemjAuV2A70ZEt3Rx0eG5O91ib1WbGcs=
newptA(gbl-proxy-user[prov_admin])-     windows-domain MEDARCH.ORG            pre-win2k-name
MEDARCH
newptA(gbl-proxy-user[prov_admin])-     exit
newptA(gbl)-   proxy-user acoProxy2

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-proxy-user[acoProxy2])-     description "user with backup and admin creds on our
servers"
X newptA(gbl-proxy-user[acoProxy2])-     user jqpublic  encrypted-password
CCNAZgM33t1HemjAuV2A70ZEt3Rx0eG5csH6+tzyZik=
```

```
X newptA(gbl-proxy-user[acoProxy2])-     windows-domain MEDARCH.ORG            pre-win2k-name
MEDARCH
X newptA(gbl-proxy-user[acoProxy2])-     exit
newptA(gbl)-   active-directory-forest MEDARCH.ORG

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-forest[MEDARCH.ORG])-               forest-root MEDARCH.ORG           pre-win2k-name
MEDARCH  192.168.25.102 preferred
X newptA(gbl-forest[MEDARCH.ORG])-               forest-root MEDARCH.ORG           pre-win2k-name
MEDARCH  192.168.25.104 preferred
X newptA(gbl-forest[MEDARCH.ORG])-               forest-root MEDARCH.ORG           pre-win2k-name
MEDARCH  192.168.25.111 preferred
X newptA(gbl-forest[MEDARCH.ORG])-               forest-root MEDARCH.ORG           pre-win2k-name
MEDARCH  192.168.25.109 preferred
X newptA(gbl-forest[MEDARCH.ORG])-               forest-root MEDARCH.ORG           pre-win2k-name
MEDARCH  192.168.25.110 preferred
X newptA(gbl-forest[MEDARCH.ORG])-                tree-domain bostonmed.org      pre-win2k-name
BOSTONMED 172.16.74.89 preferred
X newptA(gbl-forest[MEDARCH.ORG])-                tree-domain bostonmed.org      pre-win2k-name
BOSTONMED 172.16.74.88 preferred
X newptA(gbl-forest[MEDARCH.ORG])-                tree-domain fdtestnet.net pre-win2k-name
FDTESTNET 172.16.168.21 preferred
X newptA(gbl-forest[MEDARCH.ORG])-                tree-domain fdtestnet.net pre-win2k-name
FDTESTNET 172.16.168.22preferred
X newptA(gbl-forest[MEDARCH.ORG])-                child-domain bostoncifs.fdtestnet.net
pre-win2k-name BOSTONCIFS 10.19.230.94 preferred
X newptA(gbl-forest[MEDARCH.ORG])-                child-domain bostoncifs.fdtestnet.net
pre-win2k-name BOSTONCIFS 10.19.230.88 preferred
X newptA(gbl-forest[MEDARCH.ORG])-                child-domain westcoast.medarch.org
pre-win2k-name WESTCOAST 192.168.202.9 preferred
X newptA(gbl-forest[MEDARCH.ORG])-                child-domain westcoast.medarch.org
pre-win2k-name WESTCOAST 192.168.202.10 preferred
X newptA(gbl-forest[MEDARCH.ORG])-                child-domain westcoast.medarch.org
pre-win2k-name WESTCOAST 192.168.202.11 preferred
X newptA(gbl-forest[MEDARCH.ORG])-               child-domain ma.ne.medarch.org
pre-win2k-name MA      192.168.25.103 preferred
X newptA(gbl-forest[MEDARCH.ORG])-               child-domain ma.ne.medarch.org
pre-win2k-name MA      192.168.25.105 preferred
X newptA(gbl-forest[MEDARCH.ORG])-                child-domain ne.medarch.org
pre-win2k-name NE    172.16.124.73 preferred
X newptA(gbl-forest[MEDARCH.ORG])-                child-domain ne.medarch.org
pre-win2k-name NE    172.16.124.19preferred
X newptA(gbl-forest[MEDARCH.ORG])-                exit
newptA(gbl)-   active-directory-forest WELLS.ME.ORG

% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-forest[WELLS.ME.ORG])-     forest-root wells.me.org pre-win2k-name WELLS
172.16.108.139 preferred
X newptA(gbl-forest[WELLS.ME.ORG])-     forest-root wells.me.org pre-win2k-name WELLS
172.16.108.136
X newptA(gbl-forest[WELLS.ME.ORG])-     child-domain adk.wells.me.org pre-win2k-name ADK
172.16.110.8 preferred
X newptA(gbl-forest[WELLS.ME.ORG])-     child-domain adk.wells.me.org pre-win2k-name ADK
172.16.110.5
X newptA(gbl-forest[WELLS.ME.ORG])-     child-domain york.wells.me.orgpre-win2k-name YORK
172.16.120.22preferred
X newptA(gbl-forest[WELLS.ME.ORG])-     child-domain york.wells.me.orgpre-win2k-name YORK
172.16.120.5preferred
X newptA(gbl-forest[WELLS.ME.ORG])-     exit
newptA(gbl)-   active-directory-forest VT.COM
```

```
% INFO: Configuration object already exists so ignoring new CLI commands.

X newptA(gbl-forest[VT.COM])-     forest-root VT.COM pre-win2k-name VT 172.16.213.8 preferred
X newptA(gbl-forest[VT.COM])-     forest-root VT.COM pre-win2k-name VT 172.16.213.9
X newptA(gbl-forest[VT.COM])-     tree-domain ATLANTIC.ME.ORG pre-win2k-name ATLANTIC
172.16.210.14preferred
X newptA(gbl-forest[VT.COM])-     tree-domain ATLANTIC.ME.ORG pre-win2k-name ATLANTIC
172.16.210.7preferred
X newptA(gbl-forest[VT.COM])-     child-domain MCNIELS.VT.COM pre-win2k-name MCNIELS
172.16.240.70preferred
X newptA(gbl-forest[VT.COM])-    child-domain MCNIELS.VT.COM pre-win2k-name MCNIELS 172.16.240.88
preferred
X newptA(gbl-forest[VT.COM])-     child-domain BSH.ATLANTIC.ME.ORG pre-win2k-name BASSHARBOR
10.51.100.6 preferred
X newptA(gbl-forest[VT.COM])-     child-domain BSH.ATLANTIC.ME.ORG pre-win2k-name BASSHARBOR
10.52.160.1 preferred
X newptA(gbl-forest[VT.COM])-     exit
newptA(gbl)-   external-filer ntap-npt2
newptA(gbl-filer[ntap-npt2])-     description "New filer, rack 4 (Newport)"
newptA(gbl-filer[ntap-npt2])-     ip address 10.51.100.173
newptA(gbl-filer[ntap-npt2])-     exit
newptA(gbl)-   external-filer ntap-nwpt
newptA(gbl-filer[ntap-nwpt])-     description "big filer in Newport's front lab"
newptA(gbl-filer[ntap-nwpt])-     ip address 10.51.100.172
newptA(gbl-filer[ntap-nwpt])-     exit
newptA(gbl)-   external-filer ntap-pr-1
newptA(gbl-filer[ntap-pr-1])-     cifs-port 445
newptA(gbl-filer[ntap-pr-1])-     description "New NAS filer, rack 5"
newptA(gbl-filer[ntap-pr-1])-     ip address 192.168.103.52
newptA(gbl-filer[ntap-pr-1])-     exit
newptA(gbl)-   external-filer ntap-prov
newptA(gbl-filer[ntap-prov])-     cifs-port 445
newptA(gbl-filer[ntap-prov])-     description "NAS filer in computer lab, rack 2"
newptA(gbl-filer[ntap-prov])-     ip address 192.168.103.7
newptA(gbl-filer[ntap-prov])-     exit
newptA(gbl)-   schedule oncePerDay
newptA(gbl-schedule[oncePerDay])-     every 1 days
newptA(gbl-schedule[oncePerDay])-     start 07/09/2010:11:22:00
newptA(gbl-schedule[oncePerDay])-     exit
newptA(gbl)-   schedule twicePerMonth
newptA(gbl-schedule[twicePerMonth])-     every day-list 1,15
newptA(gbl-schedule[twicePerMonth])-     start 07/09/2010:11:22:00
newptA(gbl-schedule[twicePerMonth])-     exit
newptA(gbl)-   policy-age-fileset popular-files
newptA(gbl-fs-age[popular-files])-     select-files newer-than 2 weeks
newptA(gbl-fs-age[popular-files])-     exit
newptA(gbl)-   policy-age-fileset unpopular-files
newptA(gbl-fs-age[unpopular-files])-     select-files older-than 2 weeks
newptA(gbl-fs-age[unpopular-files])-     exit
newptA(gbl)-   policy-filename-fileset all-files
newptA(gbl-fs-name[all-files])-     recurse
newptA(gbl-fs-name[all-files])-     exit
newptA(gbl)-   windows-mgmt-auth mmc-admins
newptA(gbl-mgmt-auth[mmc-admins])-     permit all any
newptA(gbl-mgmt-auth[mmc-admins])-     user lab99 windows-domain MEDARCH.ORG
newptA(gbl-mgmt-auth[mmc-admins])-     user mhoward_md windows-domain MEDARCH.ORG
newptA(gbl-mgmt-auth[mmc-admins])-     user meg windows-domain MEDARCH.ORG
newptA(gbl-mgmt-auth[mmc-admins])-     exit
newptA(gbl)-   namespace provMed
newptA(gbl-ns[provMed])-     protocol cifs
newptA(gbl-ns[provMed])-     cifs anonymous-access
```

```
newptA(gbl-ns[provMed])-     cifs authentication kerberos
newptA(gbl-ns[provMed])-     cifs authentication ntlm
newptA(gbl-ns[provMed])-     cifs authentication ntlmv2
newptA(gbl-ns[provMed])-     cifs filer-signatures
newptA(gbl-ns[provMed])-     proxy-user prov_admin
newptA(gbl-ns[provMed])-     windows-mgmt-auth mmc-admins
newptA(gbl-ns[provMed])-     sam-reference ntap-pr-1
newptA(gbl-ns[provMed])-     volume /mds
newptA(gbl-ns-vol[provMed~/mds])-       filer-subshares native-names-only
newptA(gbl-ns-vol[provMed~/mds])-       modify
newptA(gbl-ns-vol[provMed~/mds])-       reimport-modify
newptA(gbl-ns-vol[provMed~/mds])-       reserve files 4000000
newptA(gbl-ns-vol[provMed~/mds])-       auto sync files
newptA(gbl-ns-vol[provMed~/mds])-       metadata share ntap-pr-1 nfs3 /vol/vol3/arx_meta cluster
providence
newptA(gbl-ns-vol[provMed~/mds])-       metadata share ntap-npt2 nfs3 /vol/vol3/arxMeta cluster
newport
newptA(gbl-ns-vol[provMed~/mds])-       no compressed-files
newptA(gbl-ns-vol[provMed~/mds])-       named-streams
newptA(gbl-ns-vol[provMed~/mds])-       no persistent-acls
newptA(gbl-ns-vol[provMed~/mds])-       no sparse-files
newptA(gbl-ns-vol[provMed~/mds])-       unicode-on-disk
newptA(gbl-ns-vol[provMed~/mds])-       share mds
newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         description "home dirs for new doctors"
newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         import priority 1
newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         import skip-managed-check
newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         import sync-attributes
newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         filer ntap-pr-1 cifs MDS cluster providence
newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         filer ntap-npt2 cifs MDS cluster newport
newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         enable
newptA(gbl-ns-vol-shr[provMed~/mds~mds])-         exit
newptA(gbl-ns-vol[provMed~/mds])-       share surgeons
newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         description "home dirs for veteran
surgeons"
newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         import skip-managed-check
newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         import sync-attributes
newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         filer ntap-prov cifs SURGEONS cluster
providence
newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         filer ntap-nwpt cifs SURGEONS cluster
newport
newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         enable
newptA(gbl-ns-vol-shr[provMed~/mds~surgeons])-         exit
newptA(gbl-ns-vol[provMed~/mds])-       place-rule pop-md-files-2-new
newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         schedule oncePerDay
newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         report tier1 error-only
newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         inline report daily tier1
error-only
newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         from fileset popular-files
newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         target share mds
newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         enable
newptA(gbl-ns-vol-plc[provMed~/mds~pop-md-files-2-new])-         exit
newptA(gbl-ns-vol[provMed~/mds])-       place-rule unpop-md-files-2-old
newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-         schedule oncePerDay
newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-         from fileset unpopular-files
newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-         target share surgeons
newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-         no inline notify
newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-         enable
newptA(gbl-ns-vol-plc[provMed~/mds~unpop-md-files-2-old])-         exit
newptA(gbl-ns-vol[provMed~/mds])-       place-rule masters-2-new
newptA(gbl-ns-vol-plc[provMed~/mds~masters-2-new])-         from fileset all-files match
directories promote-directories
newptA(gbl-ns-vol-plc[provMed~/mds~masters-2-new])-         target share mds
```

```
newptA(gbl-ns-vol-plc[provMed~/mds~masters-2-new])-         enable
newptA(gbl-ns-vol-plc[provMed~/mds~masters-2-new])-         exit
newptA(gbl-ns-vol[provMed~/mds])-       volume-group 1
newptA(gbl-ns-vol[provMed~/mds])-       enable
newptA(gbl-ns-vol[provMed~/mds])-       exit
newptA(gbl-ns[provMed])-     volume /rns
newptA(gbl-ns-vol[provMed~/rns])-       filer-subshares native-names-only
newptA(gbl-ns-vol[provMed~/rns])-       modify
newptA(gbl-ns-vol[provMed~/rns])-       reimport-modify
newptA(gbl-ns-vol[provMed~/rns])-       reserve files 4000000
newptA(gbl-ns-vol[provMed~/rns])-       auto sync files
newptA(gbl-ns-vol[provMed~/rns])-       metadata share ntap-pr-1 nfs3 /vol/vol4/arx_meta cluster
providence
newptA(gbl-ns-vol[provMed~/rns])-       metadata share ntap-npt2 nfs3 /vol/vol4/arxMeta cluster
newport
newptA(gbl-ns-vol[provMed~/rns])-       no compressed-files
newptA(gbl-ns-vol[provMed~/rns])-       named-streams
newptA(gbl-ns-vol[provMed~/rns])-       no persistent-acls
newptA(gbl-ns-vol[provMed~/rns])-       no sparse-files
newptA(gbl-ns-vol[provMed~/rns])-       unicode-on-disk
newptA(gbl-ns-vol[provMed~/rns])-       share nurses
newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         description "home dirs for veteran surgeons"
newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         import skip-managed-check
newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         import sync-attributes
newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         filer ntap-prov cifs NURSES cluster
providence
newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         filer ntap-nwpt cifs NURSES cluster newport
newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         enable
newptA(gbl-ns-vol-shr[provMed~/rns~nurses])-         exit
newptA(gbl-ns-vol[provMed~/rns])-       share rns
newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         description "home dirs for new nurses"
newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         import priority 1
newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         import skip-managed-check
newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         import sync-attributes
newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         filer ntap-pr-1 cifs RNS cluster providence
newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         filer ntap-npt2 cifs RNS cluster newport
newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         enable
newptA(gbl-ns-vol-shr[provMed~/rns~rns])-         exit
newptA(gbl-ns-vol[provMed~/rns])-       place-rule pop-rn-files-2-new
newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         schedule oncePerDay
newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         report tier1 error-only
newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         inline report daily tier1
error-only
newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         from fileset popular-files
newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         target share rns
newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         enable
newptA(gbl-ns-vol-plc[provMed~/rns~pop-rn-files-2-new])-         exit
newptA(gbl-ns-vol[provMed~/rns])-       place-rule unpop-rn-files-2-old
newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-          schedule oncePerDay
newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-         from fileset unpopular-files
newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-         target share nurses
newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-         no inline notify
newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-         enable
newptA(gbl-ns-vol-plc[provMed~/rns~unpop-rn-files-2-old])-         exit
newptA(gbl-ns-vol[provMed~/rns])-       place-rule masters-2-new
newptA(gbl-ns-vol-plc[provMed~/rns~masters-2-new])-         from fileset all-files match
directories promote-directories
newptA(gbl-ns-vol-plc[provMed~/rns~masters-2-new])-         target share rns
newptA(gbl-ns-vol-plc[provMed~/rns~masters-2-new])-         enable
newptA(gbl-ns-vol-plc[provMed~/rns~masters-2-new])-         exit
newptA(gbl-ns-vol[provMed~/rns])-       volume-group 1
newptA(gbl-ns-vol[provMed~/rns])-       enable
```

```
newptA(gbl-ns-vol[provMed~/rns])-       exit
newptA(gbl-ns[provMed])-     exit
newptA(gbl)-   namespace provMed

% INFO: Configuration object already exists. Merging child objects.

X newptA(gbl-ns[provMed])-     protocol cifs
X newptA(gbl-ns[provMed])-     cifs anonymous-access
X newptA(gbl-ns[provMed])-     cifs authentication kerberos
X newptA(gbl-ns[provMed])-     cifs authentication ntlm
X newptA(gbl-ns[provMed])-     cifs authentication ntlmv2
X newptA(gbl-ns[provMed])-     cifs filer-signatures
X newptA(gbl-ns[provMed])-     proxy-user prov_admin
X newptA(gbl-ns[provMed])-     windows-mgmt-auth mmc-admins
X newptA(gbl-ns[provMed])-     sam-reference ntap-pr-1
X newptA(gbl-ns[provMed])-     exit
newptA(gbl)-   global server provmed.MEDARCH.ORG
newptA(gbl-gs[provmed.MEDARCH.ORG])-                description "CIFS service - home dirs"
newptA(gbl-gs[provmed.MEDARCH.ORG])-                windows-domain MEDARCH.ORG
newptA(gbl-gs[provmed.MEDARCH.ORG])-                active-directory proxy-user prov_admin
newptA(gbl-gs[provmed.MEDARCH.ORG])-               virtual server newptA 192.168.74.61 255.255.0.0
vlan 103 cluster newport
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                  wins 192.168.25.102
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                  active-directory alias
pubs_test_provmed
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                  active-directory alias
pubs_test_ntap-prov
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                  active-directory alias
pubs_test_ntap-pr-1
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                  active-directory alias
pubs_test_provmed.MEDARCH.ORG
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                  active-directory alias
pubs_test_ntap-prov.MEDARCH.ORG
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                  active-directory alias
pubs_test_ntap-pr-1.MEDARCH.ORG
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                  enable
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.61])-                  exit
newptA(gbl-gs[provmed.MEDARCH.ORG])-               virtual server provA 192.168.74.91 255.255.0.0
vlan 103 cluster providence
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  wins 192.168.25.102
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_provmed
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_ntap-prov
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_ntap-pr-1
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_provmed.MEDARCH.ORG
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_ntap-prov.MEDARCH.ORG
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_ntap-pr-1.MEDARCH.ORG
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_provmed
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_ntap-prov
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_ntap-pr-1
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_provmed.MEDARCH.ORG
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                  active-directory alias
pubs_test_ntap-prov.MEDARCH.ORG
```

```
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 active-directory alias
pubs_test_ntap-pr-1.MEDARCH.ORG
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 enable
newptA(gbl-gs-vs[provmed.MEDARCH.ORG~192.168.74.91])-                 exit
newptA(gbl-gs[provmed.MEDARCH.ORG])-               enable
newptA(gbl-gs[provmed.MEDARCH.ORG])-               exit
newptA(gbl)-   cifs provmed.MEDARCH.ORG
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               browsing provMed
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               description "user home directories"
newptA(gbl-cifs[provmed.MEDARCH.ORG])-              kerberos-creds provmed.MEDARCH.ORG
MEDARCH.ORG            ECqJ5gJ0zkU9hskVeO7a2w== provmed$ HOST/provmed.MEDARCH.ORG           2
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               signatures
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               dynamic-dns pubs_test_provmed
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               dynamic-dns pubs_test_ntap-prov
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               dynamic-dns pubs_test_ntap-pr-1
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               export provMed /mds as MDS description "home
dirs for the docs"
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               export provMed /rns as RNS description "home
dirs for the nurses"
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               export provMed /mds/blake as BLAKE
filer-subshare
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               export provMed /mds/hogtroff as HOGTROFF
filer-subshare
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               export provMed /mds/mccoy as MCCOY
filer-subshare
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               export provMed /mds/mhoward as MHOWARD
filer-subshare
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               export provMed /mds/pierce as PIERCE
filer-subshare
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               enable
newptA(gbl-cifs[provmed.MEDARCH.ORG])-               exit
newptA(gbl)-   config-replication prov2newport cluster providence
newptA(gbl-cfg-repl[prov2newport])-     description "send service config to Newport site"
newptA(gbl-cfg-repl[prov2newport])-     report gblRepl2newport
newptA(gbl-cfg-repl[prov2newport])-     schedule oncePerDay
newptA(gbl-cfg-repl[prov2newport])-     target-cluster newport
newptA(gbl-cfg-repl[prov2newport])-     target-file provSvcs.rcfg
newptA(gbl-cfg-repl[prov2newport])-     user admin encrypted-password
Sd/LKQM33t1HemjAuV2A70ZEt3Rx0eG5O91ib1WbGcs=
newptA(gbl-cfg-repl[prov2newport])-     enable
newptA(gbl-cfg-repl[prov2newport])-     exit
newptA(gbl)-   exit
```

remove cluster-config

Purpose Use the remove cluster-config command to delete a cluster-name from the database along with any references to that cluster name. If you use this with the remote cluster name, this removes all filers, metadata shares, and/or virtual servers associated with that cluster. If you use this with the current cluster name, this only removes the cluster name from those configuration objects, leaving them otherwise intact.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax `remove cluster-config` ***`cluster-name`***

cluster-name (1-64 characters) identifies the cluster. You can use the show cluster command for a complete list of configured clusters.

Default(s) None.

Guidelines The CLI prompts for confirmation before removing or changing any configuration objects; enter **yes** to proceed.

For a remote cluster, this removes the following objects from the system:

- any filer,
- any metadata share, or
- any virtual server that resides in the cluster.

For the local cluster, this only removes the cluster name from the above objects. They remain in the configuration, but without the cluster name as part of their definitions.

If any of the above object types refer to a cluster name, you cannot use no cluster-name to remove the cluster name. This command exists as a convenience; it is a single command for removing the cluster name and all of its dependent configuration.

To change the name of the local cluster, you can use this command to remove the current name and then use the cluster-name command to add the new cluster name to all of the above objects at once.

Sample

```
bstnA# remove cluster-config portland

This removes all configuration for this cluster.

Are you sure? [yes/no] yes
```

removes all configuration objects associated with the “portland” cluster.

Related Commands cluster-name

report (gbl-cfg-repl)

Purpose Use this command to enable progress reports for the current config-replication rule.

Use no report to prevent progress reports.

Mode gbl-cfg-repl

Security Role(s) storage-engineer or crypto-officer

Syntax

```
report file-prefix [verbose]
no report
```

file-prefix (1-1024 characters) sets a prefix for all config-replication reports from this rule. Each report has a unique name in the following format:

prefixYearMonthDayHourMinute.rpt

For example, gblRepl2newport201003031200.rpt could be the name for one report with the "gblRepl2newport" prefix.

verbose (optional) enables verbose data in the reports.

Default(s) `no report`

Guidelines Every time the config-replication rule fires and duplicates the global-config, the rule can generate a report to show the details of the copy session. The reporting feature is disabled by default; we recommend that you enable reporting for all config-replication rules, to keep a log of all these copy sessions.

Use show reports for a list of reports, or show reports *file-name* to show the contents of one report.

Samples

```
provA(gbl-cfg-repl[prov2newport])# report gblRepl2newport
```

enables reports for the config-replication rule, "prov2newport." For a sample report, see Figure 35.3.

```
bstnA(gbl-cfg-repl[dr-test])# report drTest verbose
```

enables verbose reports for the config-replication rule, "dr-test."

```
newptA(gbl-cfg-repl[newpt2prov])# no report
```

disables reports for the "newpt2prov" rule.

Related Commands config-replication

Figure 35.3 *Sample Report: Config-Replication Report*

```
provA# show reports gblRepl2newport.rpt
**** Configuration Replication Report: Started at Tue Jul 13 11:42:22 2010 ****
**** Software Version: 5.02.000.12592 (Jul  9 2010 20:13:13) [nbuilds]
**** Hardware Platform: ARX-500
**** Report Destination:

% NOTICE: CONFIG_REPL_UPDATE (42991634):
  Config-replication for cluster 'newport' member 'newptA' was successful.
```

```
% NOTICE: CONFIG_REPL_MISSING_MEMBER (42991635):
  Config-replication for cluster 'newport' is missing member '2'.

**** Total processed:               0
**** Elapsed time:           00:00:02
**** Configuration Replication Report: DONE at Tue Jul 13 11:42:24 2010 ****
```

schedule (gbl-cfg-repl)

Purpose

Use this schedule command to assign a schedule to the current config-replication rule.

Use no schedule to remove the rule's schedule.

Mode

gbl-cfg-repl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
schedule name
no schedule
```

name (1-64 characters) identifies the schedule. Use show schedule for a list of configured schedules.

Default(s)

None.

Guidelines

A config-replication rule copies the global-configuration from the current ARX to another ARX at a backup site. This copy operation occurs on the schedule you choose with this command.

To create a schedule, use the gbl-mode schedule command.

You cannot use a schedule with a fixed duration; a config-replication rule must always run to completion in order to succeed.

Sample

```
provA(gbl-cfg-repl[prov2newport])# schedule oncePerDay
```

assigns the "oncePerDay" schedule to the config-replication rule, "prov2newport."

Related Commands

config-replication
schedule
show schedule

show cluster

Purpose Use the show cluster command to get a list of all ARX *clusters*, or redundant pairs, known to the current ARX. You can set up ARX clusters so that one cluster can act as a backup for another.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show cluster [`*`cluster-name`*`]`**

cluster-name (optional, 1-64 characters) shows only the chosen cluster.

Guidelines The summary output shows a table of clusters, one row per cluster. Each row has the following columns:

Cluster Name identifies the cluster. You set this when you use the cluster-name command to create the cluster.

Member is the hostname of one ARX in the cluster (typically a redundant pair).

Member in the next column is for the above switch's redundant peer, if there is one.

Samples `provA# show cluster`

shows all clusters. See Figure 35.4 for sample output.

`provA# show cluster newport`

shows one cluster. See Figure 35.5 for sample output.

Related Commands cluster-name

Figure 35.4 *Sample Output: show cluster*

```
provA# show cluster

Cluster Name               Member                       Member
-------------------------------------------------------------------------------
providence                 provA
newport                    newptA
```

Figure 35.5 *Sample Output: show cluster newport*

```
provA# show cluster newport

Cluster Name               Member                       Member
-------------------------------------------------------------------------------
newport                    newptA
```

show config-replication

Purpose Use the show config-replication command to see the current status of one or more config-replication rules. A config-replication rule copies the global config (see show global-config) from the current site to an ARX at a disaster-recovery site.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show config-replication [rule-name]`

rule-name (optional, 1-32 characters) provides configuration details for a particular config-replication rule.

Guidelines: Summary Output The summary output shows a table of config-replication rules, one row per rule. Each row has the following columns:

Name identifies the rule. You set this when you use the config-replication command to create the rule.

Cluster is the name of the target ARX cluster (or redundant pair) where the rule sends its global configuration file. You can set this with the target-cluster command.

Filename is the name of the configuration file that the rule creates at the remote cluster. You can change this with the target-file command.

Schedule is the name of the schedule for the config-replication rule. Use the schedule (gbl-cfg-repl) command to assign a schedule to the rule.

Admin State is "Enabled" or "Disabled," depending on the setting of the enable (gbl-cfg-repl) command.

Guidelines: Detailed Output Detailed output appears when you specify a config-replication rule in the command. The detailed output contains these fields:

Name identifies the rule, as described above.

Description describes the rule, as specified with the description (gbl-cfg-repl) command.

Cluster is the name of target cluster, as described for the summary output above.

Target Filename is the name of the configuration file that the rule creates. This is the same as the Filename in the summary output.

Report Prefix shows the prefix used in all of the rule's reports. The reports appear on the source cluster. You can use the report (gbl-cfg-repl) command to change this prefix.

Username is the administrative user name used for authentication at the remote cluster. These credentials confirm that the rule is authorized to write the configuration file. This is a valid administrative user at the ARX cluster that receives the copy.

Admin State is "Enabled" or "Disabled," as described above for the summary output.

Guidelines: Detailed Output (Cont.)

Status shows the results of the last attempt to copy the global-config:

- "Copy Successful" indicates a success.
- "Not Copied Yet" means that the rule's schedule (gbl-cfg-repl) has not started yet.
- "The config-replication target-cluster is not defined for *rule*." shows that the copy failed because the config-replication rule has no target. You can use the target-cluster command to define a target cluster for the copy operation.
- "The config-replication user is not defined for *rule*." shows that the copy failed because the config-replication rule has no user defined. The rule needs an administrative user name that is valid at the target cluster, so that it has permission to copy a file onto the remote ARX. You can use the user (gbl-cfg-repl) command to identify an administrative account for the copy operation.
- "File *file-name* does not exist." means that an internal process failed to open *file-name*, a temporary version of the global-config file. This indicates an internal error; contact F5 Support if you see this.
- Any other error indicates a low-level error with the SCP process. Contact F5 Support if you see any such error in this field.

Schedule is the schedule's configured name, chosen with the schedule (gbl-cfg-repl) command. The fields in this table show the configuration of the schedule.

- Start Time can be reset by the gbl-schedule start command. This is the date and time of day used for all time-based calculations: for example, if a schedule runs every Saturday, this determines the time of day on Saturdays when the schedule fires.
- Interval is set by the gbl-schedule every command. This is the time between scheduled runs.
- Status is "Paused" (the schedule is between run times), "Running," or "Waiting for start time" (the schedule has never been invoked).
- Previous is a sub-table showing the Run time (when the run started) and End time of the most-recent run.
- Current appears if the rule is currently running. This has the same fields as the Previous table, Run time and End time.
- Next shows when the schedule's next run will begin. It also has the Run time and End time fields.

Samples

```
provA# show config-replication
```

shows the status of all config-replication rules. See Figure 35.6 for sample output.

```
provA# show config-replication prov2newport
```

shows the detailed status for a particular rule. See Figure 35.7 for sample output.

Related Commands config-replication

Figure 35.6 *Sample Output: show config-replication*

```
provA# show config-replication

Name               Cluster          Filename        Schedule        Admin State
-------------------------------------------------------------------------------
prov2newport       newport          provSvcs.rcfg   oncePerDay      Enabled
```

Figure 35.7 *Sample Output: show config-replication prov2newport*

```
provA# show config-replication prov2newport
  Name:             prov2newport
  Description:      send service config to Newport site
  Cluster:          newport
  Target Filename:  provSvcs.rcfg
  Report Prefix:    gblRepl2newport
  Username:         admin
  Admin State:      Enabled
  Status:           Copy Successful

  Schedule:         oncePerDay
    Start Time:       Mon Sep 20 03:36:00 2010
    Interval:         1 days
    Status:           Running (runs in 23:56:30)

    Previous:
      Run Time:         Mon Sep 20 03:36:00 2010 (Schedule Start)
      End Time:         N/A

    Next:
      Run Time:         Tue Sep 21 03:36:00 2010
      End Time:         N/A
```

target-cluster

Purpose Use the target-cluster command to choose a remote ARX cluster as a target for copies of the local cluster's configuration. The current config-replication rule copies the global-configuration from the local cluster on a schedule. This prepares the remote cluster as a backup for the local cluster in the event of a disaster at the local site.

Mode gbl-cfg-repl

Security Role(s) storage-engineer or crypto-officer

Syntax **`target-cluster`** ***`cluster-name`***

cluster-name (1-64 characters) is the target cluster for the copy of the global-config file. This cluster must be defined on the current ARX; use the show cluster command for a list of defined clusters, or the cluster-name command to define one.

Default(s) None.

Guidelines This command sets a cluster target for the current config-replication rule. A *cluster* is a redundant pair of ARX peers (see redundancy), or possibly a standalone ARX.

Sample

```
provA(gbl-cfg-repl[prov2newport])# target-cluster newport
```

uses the "newport" cluster as the target for the "prov2newport" rule.

Related Commands

config-replication
cluster-name
show cluster

target-file

Purpose The current config-replication rule copies the global-configuration from the local cluster to a remote cluster on a schedule. Use the target-file command to choose the name of the remote copy. The copy operation prepares the remote cluster as a backup for the local cluster in the event of a disaster.

Mode gbl-cfg-repl

Security Role(s) storage-engineer or crypto-officer

Syntax **`target-file`** ***`file-name`***

file-name (1-255 characters) is the file name for the destination file.

Default(s) None.

Guidelines This command sets a file-name target for the current config-replication rule. A *cluster* is a redundant pair of ARX peers (see redundancy), or possibly a standalone ARX. A config-replication rule copies the local cluster's global configuration (see show global-config) to a remote cluster. This prepares the remote cluster as a backup for the local cluster, in case of a disaster at the local site.

Sample `provA(gbl-cfg-repl[prov2newport])# target-file provSvcs.rcfg`
uses the name, "provSvcs.rcfg," as the target file name for the "prov2newport" rule.

Related Commands config-replication

user (gbl-cfg-repl)

Purpose

When you use a config-replication rule to copy the global configuration to a remote ARX cluster, you require an administrative username and password for the remote peer. Use the user command to set this username and password.

Use the no form of this command to remove the administrative credentials, effectively disabling the config-replication rule.

Mode

gbl-cfg-repl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
user name
no user
```

name (1-32 characters) is a valid administrative user at a remote ARX.

Default(s)

None.

Guidelines

The CLI prompts twice for a password. See the sample below.

Sample

```
provA(gbl-cfg-repl[prov2newport])# user admin
Password: s3cr3tpasswd
Validate Password: s3cr3tpasswd
provA(gbl-cfg-repl[prov2newport])# ...
```

provides a username and password, admin and s3cr3tpasswd, for sending global-config copies to a remote ARX.

Related Commands

config-replication

36

Software Upgrades

boot system

Purpose To upgrade (or otherwise change) switch software, you must arm the switch with a new release file and then restart all of the switch modules. Use the boot system command to arm the ARX with a release file.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `boot system release-file`

release-file (1-255 characters) identifies the release file. Use the show directories releases command to see all available release files. Use copy to copy a release file to the switch.

Default(s) None

Guidelines To upgrade switch software, you must

- obtain a new release file from https://downloads.f5.com/esd/productlines.jsp,
- use copy to copy the new release file to the switch,
- use license activate or license activate file if you have never done so, or if you renewed your license to be eligible for this release,
- use boot system to arm the switch with the new release file, and
- use reload to restart all modules and put the new release file into service.

Use show boot to verify that the switch is armed with the new release file.

To revert to an older release after an upgrade, you can downgrade the switch to the old software. You should only do this under the guidance of F5 Support. Downgrading deletes all configuration parameters, both the running-config and the global-config; save both configuration files off of the box before you use this command to downgrade (use copy running-config and copy global-config). Then edit the config files so that they only use CLI commands from the downgrade release. Use boot system to select the old release, then use reload to install it. When you enter the boot system command, a CLI prompt asks for confirmation before deleting the configuration. Enter **yes** to continue. After the reload, copy the edited config files onto the system, run the running-config, use license activate to enable the license, and then run the global-config.

Each release file may contain module firmware bundled with the software. The release installation (after reboot) puts all such firmware into flash memory on each module, ready to be installed. Use the show firmware upgrade command to enumerate any modules that have newer versions of firmware waiting for them. To install the new firmware onto the module(s), use firmware upgrade.

In a redundant pair of switches, you upgrade the backup peer before you upgrade the active peer. This command only runs on the backup peer in any redundant pair.

Guidelines: Licensing for New Releases

The following error indicates that you are not licensed for the new software release:

```
The current license is not eligible to upgrade to x.y.z. The service
contract is not active as of service-check-date. Please re-activate the
license to get current service contract information before performing
the upgrade.
```

Contact F5 Support to renew or extend your service contract, then run the license activate or license activate file command to reactivate your license.

Guidelines: ARX-1500 and ARX-2500 Redundancy

The ARX-1500 and ARX-2500 store their metalog data on their internal disks, along with logs, software-release files, and other management data. Managed volumes write their metalog data as clients change the volume state; the metalog is used to restore the volume configuration in the event of a failure. The metalog is also copied to the redundant peer, if there is one (see redundancy). The speed of many volume operations depends on fast metalog writes.

Some other system operations create a large number of writes to the internal disk, potentially slowing metalog writes. This can slow volume performance, even if it occurs on the backup peer. The process of upgrading the software release is extremely disk intensive, and may cause a noticeable performance degradation. During an upgrade, you use

- the copy command (copy ftp, copy {nfs|cifs}, copy scp, or copy tftp) to copy a full release file to the disk, and
- this command (boot system) to unpack the release file on the disk.

You should perform such disk-intensive operations during off-peak hours on the ARX-1500 and ARX-2500. This is true whether you run the operations on the active peer or the backup.

Samples

```
bstnA# copy nfs wwmed /rcrds/.admin/ 12345.rel releases new.rel
bstnA# boot system new.rel

% INFO: The boot system command may take up to 5 minutes to complete.
```

arms the switch with a release file, “new.rel,” to be loaded onto all modules at next reboot.

```
bstnA# boot system previous.rel

WARNING: Arming this release will cause your configuration to be reset
to factory default because it is a lower version than the running
release.

Proceed? [yes/no] yes

% INFO: The boot system command may take up to 5 minutes to complete.
```

reverts the switch to an older release. This removes *all* configuration parameters from the system.

Related Commands

show directories `releases`
copy {nfslcifs}
copy ftp
copy scp
copy tftp
license activate
license activate file
reload
show boot

firmware upgrade

Purpose A software installation may include firmware upgrades for the chassis modules, but does not install the new firmware. Use the firmware upgrade command to install firmware from a new software release.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `firmware upgrade {slot-id | all}`

slot-id (1-2) loads the firmware into the given slot only. The firmware is installed and activated the next time the switch reboots. This option is only valid on an ARX-4000.

all loads and installs the latest firmware on all modules in the chassis. This causes the ARX to reboot.

Default(s) None.

Platforms ARX-500, ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines The CLI prompts for confirmation before upgrading any firmware; enter **yes** to proceed.

Use this command only as directed by F5 personnel.

Firmware is low-level software, such as BIOS, the RAID controller, software for the FPGAs, and the software components that manage the system's boot process. Each module type has a different set of FPGAs; these include the System-Status Bus (SSB) FPGA, the Non-Volatile RAM (NVR) FPGA, the MetaLog (MTL) FPGA, the franklog FPGA, and the macau FPGA. Module startup is managed by a set of three additional firmware components: the *bootstrap* software starts the hardware, then the *diagnostics* run to verify the hardware functions, and finally the *boot-loader* loads and starts all of the system software. The show firmware upgrade and show chassis moduleinfo commands show the versions for all of these components and FPGAs.

You can use the show firmware upgrade command to see if there is a difference between running and waiting versions of firmware, or to see the status of a firmware upgrade that you invoked with this command.

The firmware upgrade all command reboots the ARX. If the ARX is in a redundant pair, you can only use this syntax on the backup ARX. Certain RAID-maintenance processes also block the firmware upgrade all command: a raid rebuild (which you can stop with raid offline) or a raid verify (which you can stop with no raid verify).

A full firmware upgrade can take several minutes. If new BIOS is included, the upgrade can take 30 minutes or more.

Use the boot system and reload commands to install a full software release package, including (possibly) new firmware releases.

Sample

```
bstnA# firmware upgrade all

Confirmation of this command commences a firmware upgrade on the entire
chassis. During the upgrade process, the chassis reboots automatically
to complete the upgrade process. If this includes a bios upgrade, this
could take at least 30 minutes.

Proceed? [yes/no] yes
```

upgrades the firmware in all slots.

Related Commands

show firmware upgrade
show chassis **moduleinfo**
boot system

show boot

Purpose The ARX can be *armed* with a waiting release file, meaning that the waiting release is swapped with the running release the next time the modules are rebooted. Use the show boot command to see both the running release and the armed release.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show boot`**

Guidelines The armed release is labeled "System Release" and the running release is labeled "System Configuration." The following format appears for each build:

`Version` ***`version`*** **`(`*****`build-date build-time`*****`) [`*****`user`*****`]`**

Use boot system to arm the switch with the new release file. Use reload to restart all modules and put the new release file into service.

Sample

```
bstnA(cfg)# show boot

System Release
 Version 5.00.000.12345 (Feb 26 2009 20:38:56) [nbuilds] %X0A
System Configuration
 Version 500000.7
```

Related Commands boot system
reload

show chassis software

Purpose Use the show chassis software command to see the software version(s) running, armed, and otherwise available on the chassis.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show chassis software`**

Guidelines System Software shows two software-release files:

- Armed is the software-release file to be loaded on the next reboot, if any. Use boot system to arm the switch with a new release file. Use reload to restart all modules and put the new release file into service.
- Running is the software-release file that is currently running.

The next table describes all release files on the switch:

- Version
- Date is the date the release-file was installed on the ARX.
- Size is in bytes.
- Filename

Related Commands boot system
reload

Figure 36.1 *Sample Output: show chassis software*

```
bstnA(cfg)# show chassis software

System Software
---------------
Armed   :    test1.rel
Running :    test1.rel


Version        Date             Size      Filename
-------------------------------------------------
5.00.000.11569 12/01/2008 06:12  781981419 test2.rel
4.01.000.11275 12/02/2008 00:09      16384 firmware1.rel
5.00.000.11569 12/01/2008 06:12      16384 firmware2.rel
4.01.000.11275 12/02/2008 00:09  769506787 test1.rel
4.01.000.11275 12/01/2008 01:13  769506787 test3.rel
```

show firmware upgrade

Purpose Module firmware comes bundled with new software releases, but it is not installed on the modules along with the software. Use the show firmware upgrade command to determine whether or not the latest software release has new firmware for your modules.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `show firmware upgrade [verbose]`

verbose (optional) expands the output with detailed version numbers for each firmware package and module.

Valid Platforms ARX-500, ARX-1500, ARX-2000, ARX-2500, and ARX-4000

Guidelines You can use the firmware upgrade command to manually install the firmware after the software is installed. Use the show firmware upgrade command to view the running versions of firmware along with the versions that may be waiting.

The summary version shows a table with two columns: Slot and Status Summary. The Status Summary for each slot compares the running firmware with the firmware that is available from the installed software release: “Up to date,” “Pending,” “Upgrade available,” and “Upgrade in progress” are possible values.

The verbose output has a larger table with the following columns:

Slot is the slot number for the module in question.

Firmware shows the type of firmware. They appear in the following order:

- The “BIOS Firmware” runs on the processor cores. This only appears for the platforms that support BIOS-firmware upgrades.
- The “Raid Ctrl,” “Intel(R) RAID Controller ...,” or “MegaRAID SCSI ..." is the RAID-controller firmware for the internal hard drives.
- The remaining rows contain slot-specific firmware. “Bootstrap,” “Diag,” and “BootLoader” are the three components that manage the module’s boot process. The NSM versions of these are “NSM Boot,” “NSM Diag,” “NSM Exp ROM,” and “NSM BootLdr.” “SSB,” “NVR,” “MTL,” “LBA,” “Firetruck,” “Hunchback,” “Franklog,” and “Macau” are examples of FPGAs.

Running is the running version of each firmware package.

Current is the version that is bundled with the running-software release.

Status is the same as shown in the summary output.

Samples

```
bstnA(cfg)# show firmware upgrade
Show Firmware Update
--------------------

Slot Status Summary
-------------------

1    Up to date
2    Up to date
```

shows the high-level status for an ARX-4000 chassis.

```
bstnA(cfg)# show firmware upgrade verbose
```

shows detailed firmware status for the same ARX. See Figure 36.2 for sample output.

```
prtlndA# show firmware upgrade verbose
```

shows the same level of detail for an ARX-2000. See Figure 36.3 for sample output.

Related Commands firmware upgrade

Figure 36.2 *Sample Output: show firmware upgrade verbose (bstnA)*

```
bstnA> show firmware upgrade verbose
Show Firmware Update Verbose
----------------------------

Slot  Firmware      Running              Current              Status
-------------------------------------------------------------------------------

1    Bios Firmware  094.101320081858     094.101320081858     Up to date
1    Intel(R) RAID Controller SROMBSAS18E  7.0.1-0075           7.0.1-0075           Up to date
2    NSM Diag       5.03.000.13012       5.03.000.13012       Up to date
2    NSM BootLdr    5.03.000.13012       5.03.000.13012       Up to date
2    NVR            71                   71                   Up to date
2    LBA            65                   65                   Up to date
```

Figure 36.3 *Sample Output: show firmware upgrade verbose (prtlndA)*

```
prtlndA> show firmware upgrade verbose
Show Firmware Update Verbose
----------------------------

Slot  Firmware      Running              Current              Status
-------------------------------------------------------------------------------

1    Bios Firmware  046.121720091524     046.121720091524     Up to date
1    Intel(R) RAID Controller SRCSASLS4I  11.0.1-0017          11.0.1-0017          Up to date
1    NSM Diag       5.03.000.13012       5.03.000.13012       Up to date
1    NSM BootLdr    5.03.000.13012       5.03.000.13012       Up to date
1    NVR            71                   71                   Up to date
1    LBA            66                   66                   Up to date
```

37

Namespace Check (nsck) and Sync

cancel migrate-metadata

Purpose You can use the nsck ... migrate-metadata command to move a managed volume's metadata from one back-end share to another. Use this command to cancel a metadata migration that is currently in progress.

Modes priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, crypto-officer, or backup-operator

Syntax **`cancel migrate-metadata`** ***`namespace`*** **`volume`** ***`volume`***

namespace (1-30 characters) and

volume (1-1024 characters) specify the volume whose metadata is migrating between back-end shares. The nsck utility performs metadata migrations; use the show nsck command for a list of volumes where the nsck utility is running.

Default(s) None

Guidelines The CLI prompts for confirmation before canceling the metadata migration; enter **yes** to continue. This stops the volume's metadata from migrating and returns the volume to service. The managed volume uses its original metadata share after the cancellation.

If the migration is complete and the volume is in the process of restarting with the new metadata share, this command fails.

Use the nsck ... migrate-metadata command to start a metadata migration.

Sample

```
bstnA# cancel migrate-metadata insur volume /claims
Are you sure you want to cancel this migration? [yes/no] yes
Metadata migration cancelled successfully.
bstnA#
```

cancels the migration in progress for the "insur~/claims" volume.

Related Commands nsck ... migrate-metadata
show nsck

cancel migrate-volume

Purpose You can use the nsck ... migrate-volume command to move an ARX volume from one volume-group to another. Use this command to cancel a volume migration that is currently in progress.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`cancel migrate-volume`** ***`namespace`*** **`volume`** ***`volume`***

namespace (1-30 characters) and

volume (1-1024 characters) specify the volume that is migrating between volume groups.

Default(s) None

Guidelines The CLI prompts for confirmation before canceling the volume migration; enter **yes** to continue. This stops the volume from migrating and returns the volume to service in its original volume group. The cancellation may take a long time in a large volume.

If the migration has already progressed to the point where the volume is restarting in its new volume group, this command fails. If necessary, you can re-run the nsck ... migrate-volume command to send the volume back to the original group.

You can use show volume-group to show the current volume group for every volume. Use the show nsck command to monitor the progress of a volume migration.

Sample

```
bstnA# cancel migrate-volume medarcv volume /rcrds
Are you sure you want to cancel this migration? [yes/no] yes
...
bstnA#
```

cancels the migration in progress for the “medarcv~/rcrds” volume.

Related Commands

nsck ... migrate-volume
show volume-group
show nsck

cancel nsck report

Purpose To stop an nsck-report job (nsck ... report metadata-only or nsck ... report inconsistencies), use the cancel nsck report command.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`cancel nsck report`** ***`nsck-id`***

nsck-id (1-2,147,483,647) identifies one nsck-report job that is currently running. The job ID appears if you run show nsck without any options.

Guidelines This command prompts for confirmation before stopping the report; enter **yes** to continue.

This command has no effect on nsck jobs that do more than produce a report, such as those invoked by nsck ... destage or nsck ... rebuild.

Samples

```
bstnA# cancel nsck report 22
Cancel the nsck metadata-only report operation 'metadata_only.21.rpt'?
[yes/no]  yes
bstnA#
```

cancels a metadata-only report.

Related Commands

show nsck
nsck ... report inconsistencies
nsck ... report metadata-only

cancel sync

Purpose If a file or directory changes on a back-end filer without the managed volume's awareness, the volume metadata is *inconsistent* for that file. A *sync* operation discovers which files or directories are inconsistent and re-synchronizes them with the metadata; the cancel sync command stops a sync operation.

Modes priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`cancel sync {files | directories}`** ***`namespace`*** **`volume`** ***`volume`*** **`path`** ***`dir-path`***

files | directories chooses the type of sync operation to cancel.

namespace (1-30 characters) is the namespace where the sync is occurring. Use the show sync command for a list of namespaces that are being synchronized.

volume (1-1024 characters) specifies the managed volume where the sync is occurring.

dir-path (1-1024 characters) specifies the directory that is being synchronized within the volume, if any. This path is relative to the ***volume*** root, above. Use forward slashes (/) for NFS or CIFS paths.

Default(s) None

Guidelines The CLI prompts for confirmation before canceling the sync operation; enter **yes** to continue.

Use the sync files command to find files with metadata inconsistencies and repair them, or the sync directories command to find new directories on the filer that are not recorded in the metadata. Sync operations happen in the background while you continue to use the CLI; you can use the wait-for sync command to block the CLI until the sync is complete. To show all sync operations, use show sync.

Sample

```
bstnA# cancel sync files wwmed volume /acct path /jsmith
Cancel specified sync operations? [yes/no] yes
bstnA#
```

cancels the sync-files operation in progress for the /jsmith directory in the "wwmed~/acct" volume.

Related Commands

sync files
sync directories
wait-for sync
show sync

clear nsck

Purpose Use this command to clear a pending or complete nsck job, or to clear all such nsck jobs.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear nsck [`*`nsck-id`*`]`**

nsck-id (optional, 1-2,147,483,647) identifies the nsck job to clear. The job ID appears in the report name for the nsck job. If you do not include an ID, the command clears all nsck jobs.

Default(s) None

Guidelines The CLI prompts for confirmation before executing the command.

Use the show nsck command to view all nsck jobs that have run on this switch. After you issue this command, show nsck does not show any pending or complete nsck jobs.

This command does not delete any nsck report files, which you can view with the show reports command.

Sample(s)

```
bstnA# clear nsck 4
Clear NSCK ID 4? [yes/no] yes
```

clears the nsck job for nsck ID #4.

Related Commands show nsck
show reports

clear sync

Purpose The ARX holds its sync-operation records indefinitely, so that you can see the results of every sync operation through show sync. Use the clear sync command to remove sync operations from this running history.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
clear sync {files | directories}
clear sync {files | directories} namespace
clear sync {files | directories} namespace volume vol-path
clear sync {files | directories} namespace volume vol-path path
path
```

files | directories is a required choice. This selects the type of sync-operation records to clear.

The ***namespace***, ***vol-path***, and ***path*** arguments mimic the ones used in the sync files and sync directories commands. These identify one or more sync operations; if you omit all of them, the command clears all sync-files or sync-directories operations from the switch.

namespace (1-30 characters) specifies a namespace where one or more sync operations have occurred.

vol-path (optional, 1-1024 characters) specifies a volume. Use forward slashes (/) for all volumes, including CIFS.

path (optional, 1-1024 characters) identifies a particular path that has been synchronized. As with volumes, use forward slashes in all paths.

Default(s) None

Guidelines The CLI prompts for confirmation before deleting any sync records. Enter **yes** to continue.

Use show sync to show all recorded sync operations. This command clears them; after you run it, the show sync command no longer displays the cleared operation(s).

The sync files and sync directories commands each start a sync operation (which runs in the background while you continue to use the CLI), cancel sync cancels a running sync, and wait-for sync blocks the CLI until a running sync completes.

Samples

```
bstnA# clear sync files wwmed
Clear sync records? [yes/no] yes
bstnA#
```

clears all sync-files operations from the “wwmed” namespace.

```
bstnA# clear sync directories
Clear sync records? [yes/no] yes
bstnA#
```

clears all sync-directories operations on the switch. The show sync directories command displays nothing after you do this.

Related Commands

sync files
sync directories
show sync
cancel sync
wait-for sync

nsck ... destage

Purpose Use the nsck destage command to release all back-end shares from a managed volume. Once destaged, you can access the volume's filer(s) directly without risking any metadata inconsistencies.

Mode priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax

```
nsck name destage [force]
nsck name destage volume-group id [force]
nsck name destage volume volname[force]
```

name (1-30 characters) is the name of the namespace where the destage occurs. Use the show namespace command to see a list of configured namespaces.

destage directs the nsck utility to destage the namespace metadata and prepare it for re-import.

volume-group *id* (optional) chooses all of the volumes in the given volume group. A *volume group* shares CPU and memory resources, and can fail over from one chassis to another in a redundant configuration. It is also a failure domain for a group of volumes in the same namespace. Use the show volume-group command to show all volume groups and their volumes. Before a volume is enabled, you can use volume-group to assign it to a group. This command destages all of the volumes assigned to the group that you choose.

volume *volname* (optional, 1-256 characters) is the volume to destage (for example, '/var').

force (optional) indicates that you want to force the destage even if another nsck job is currently under way. You can force any nsck job with this option except nsck ... migrate-metadata or nsck ... migrate-volume; you must cancel those jobs explicitly (with cancel migrate-metadata or cancel migrate-volume) before you can destage the volume.

Default(s) None.

Guidelines This applies to a managed volume only, not to a direct volume.

The CLI asks for confirmation before releasing any volume's shares; enter **yes** to continue. This stops all client access to any of the destaged volumes.

Once the filer or share(s) have been removed and/or restored to health, you can use the enable shares command to re-import all shares into a managed volume (see enable (gbl-ns, gbl-ns-vol)). After the re-import, all NFS clients must unmount and remount any of the volume's front-end NFS exports (export (gbl-nfs)).

◆ Note

Nsck turns off a volume's modify *flag before it destages the volume, making the volume read-only (see* modify*). Before you run the* nsck ... destage *command, you can avoid this by running the gbl-ns-vol* reimport-modify *command.*

Guidelines: Correcting a Problem at an NFS-Metadata Share

If a volume's metadata share supports NFS and that NFS mount is hung, all of the volumes in the same volume group will eventually become unusable. The show volume-group command shows the volume-group assignments for your volumes. You can use the volume-group clause in this command to destage those specific volume groups, then reload the ARX to initiate a failover. The destaged volumes must re-import after this failover, but no other volumes are affected. NFS clients must un-mount and re-mount the NFS front-end exports for these volumes; CIFS clients are unaffected by the re-import. This re-establishes storage service on the peer ARX with minimal disruption to clients.

Sample

```
bstnA# nsck archives destage volume /mktg

Volume /mktg is in use by CIFS global service 'fs10.medarch.org'.

This operation will remove entries for all shares in the volume
from the namespace metadata.  To reimport a destaged share,
enable the share.

% WARNING: Volume /mktg in namespace wwmed is in use by global
services.

Destaging the volume will disrupt all clients using this volume.

Destage the volume anyway? [yes/no] yes
bstnA# ...
```

destages the '/mktg' volume in the 'archives' namespace.

Related Commands

show namespace
show reports
show external-filer
enable (gbl-ns-vol-shr)

nsck ... migrate-metadata

Purpose Use the nsck migrate-metadata command to move a managed-volume's metadata from one share to another.

Mode priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`nsck`** ***`name`*** **`migrate-metadata volume`** ***`volname`*** **`target`** ***`filer-name`*** **`{nfs3 | nfs3tcp | cifs}`** ***`path`***

name (1-30 characters) is the name of the namespace.

migrate-metadata is required.

volume ***volname*** (1-256 characters) is the chosen managed volume.

target ***filer-name*** (1-64 characters) is the new filer or file server for the volume's metadata. Use the show external-filer command for a list of configured filers.

nfs3 | nfs3tcp | cifs is the protocol for accessing the above filer (NFSv3 over UDP, NFSv3 over TCP, or CIFS). This can be different from the protocol(s) for the managed volume, though you cannot choose a CIFS share for a managed volume in an NFS-only namespace.

path (1-1024 characters) is the path to the back-end filer share, or a CIFS-share name (possibly with a sub-path). Use forward slashes (/) for Unix exports or CIFS shares.

Default(s) None.

Guidelines This applies to a managed volume only, not to a direct volume.

This stops client access to the managed volume while the metadata migration proceeds. The CLI asks for confirmation before moving the volume's metadata; enter **yes** to continue.

The migration occurs in stages for maximum protection of the volume's metadata. It copies all metadata to the new share, verifies its integrity, and brings the volume back online before it deletes the metadata from the old share. If the job is canceled or if it fails before the final transition, the managed volume comes back online with its original metadata share.

Very large volumes may take a long time for their metadata to migrate. The command returns immediately, providing you with the name of a detailed report. You can continue using the CLI while the migration occurs in the background. The show nsck command shows the current status of one or more nsck operations.

The migration report is named "migrate_metadata.*nsck-job-id*.rpt." Use show reports to list all reports, including this one. To follow the progress of the migration, you can use tail reports *report-name* follow. For a sample report, see Figure 37.1 on page 37-13.

To cancel the metadata migration (and return the volume to service with its original metadata share), use cancel migrate-metadata.

The metadata share command establishes a metadata share for a managed volume. If you decide later to move the metadata to a new share, you can use this command (nsck ... migrate-metadata) to perform the metadata migration.

Sample

```
bstnA# nsck wwmed migrate-metadata volume /acct target nas3 nfs3
/exports/meta7

Volume /acct is in use by NFS global service
'ac1.engsmoke.acopianet.com'.
Volume /acct in namespace wwmed is in use or being browsed by global
services.

Migrating metadata for this volume will disrupt all clients accessing
this volume through these services.
No other volume will be affected by this procedure.

Proceed with metadata migration? [yes/no] yes
Starting metadata migration for wwmed:/acct. 2162688 bytes to be
transferred.
Detailed status contained in report: migrate_metadata.39.rpt on switch
bstnA
```

migrates metadata for the 'wwmed~/acct' volume to a share on the "nas3" filer. See Figure 37.1 for a sample report.

Related Commands

metadata share
cancel migrate-metadata
show external-filer
show reports

Figure 37.1 *Sample Report: nsck migrate-metadata*

```
bstnA# show reports migrate_metadata.39.rpt
**** Metadata Migration Report: Started at 01/19/2012 01:51:40 -0500 ****
**** Software Version: 6.02.000.14314 (Jan 16 2012 20:04:23) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

Namespace:                  wwmed
Volume:                     /acct

Original Location:          nas1(192.168.25.21): /vol/vol2/meta1 (nfs3)
Target Location:            nas3(192.168.25.47): /exports/meta7 (nfs3)

Metadata Size:              2,162,688 (2.0 MB)

**** Total processed:               0
**** Elapsed time:           00:00:05
**** Metadata Migration Report: DONE at 01/19/2012 01:51:45 -0500 ****
```

nsck ... migrate-volume

Purpose Every ARX volume resides in a *volume group*, where it shares CPU and memory resources with the other volumes in the same group. Each volume group is a single failure domain, where certain catastrophic failures affect the other volumes in the same group but do not affect any volumes outside the group. As your deployment grows, you may want to redistribute your volumes to mitigate possible down time. Use the nsck migrate-volume command to move a volume from one volume-group to another.

Mode priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax `nsck` ***name*** `migrate-volume` ***volname*** `volume-group` ***id*** `[check-limits]`

name (1-30 characters) is the name of the namespace.

migrate-volume ***volname*** (1-256 characters) is the volume to migrate.

volume-group ***id*** chooses the destination volume group. Use the show volume-group command to show all currently-existing volume groups and their volumes. Choose a volume group from the same namespace as the source volume, or choose a new one. This volume group must be able to support the new volume without exceeding any of its licensed limits. To confirm that the migration meets all these criteria, you can use the **check-limits** option (below).

check-limits verifies that the volume migration is possible without actually migrating the volume. You can use this option to run a test before you perform the migration. This checks the licensed limits at the destination volume group and any other possible obstacles for a successful migration.

Default(s) None.

Guidelines This command stops client access to the volume until the migration is complete. The CLI lists all affected front-end services and asks for confirmation; enter **yes** to continue.

A volume cannot migrate while it is importing files from its back-end shares. (You invoke volume import with the enable (gbl-ns, gbl-ns-vol) command.) The import must be complete before you move the volume to another volume group.

If you run multiple volume migrations at the same time, they must all be to different volume groups. A volume group can only accept one migrating volume at a time.

Very large volumes may take a long time to migrate. The command returns immediately, providing you with the name of a detailed report. You can continue using the CLI while the migration occurs in the background. The show nsck command shows the current status of the migration, along with the status of all other nsck operations. After the nsck job is complete, the front-end service(s) require a few additional seconds before they allow client access to the volume.

The migration report is named "migrate_volume_group.*nsck-job-id*.rpt." Use show reports to list all reports, including this one. To follow the progress of the migration, you can use tail reports *report-name* follow. For a sample report, see Figure 37.2 on page 37-15.

To cancel the volume migration (and return the volume to service in its original volume group), use cancel migrate-volume.

Samples

```
bstnA# nsck medarcv migrate-volume /test_results volume-group 3
check-limits
% INFO: Volume-group migration limit checking passed.
```

confirms that it is possible to migrate the 'medarcv~/test_results' volume to volume group 3.

```
bstnA# nsck medarcv migrate-volume /test_results volume-group 3

Volume /test_results is in use by CIFS global service
'ac1.medarch.org'.
Volume /test_results in namespace medarcv is in use or being browsed by
global services.

Migrating the volume to a different Volume-Group will disrupt all
clients accessing this volume through these services.
No other volume will be affected by this procedure.

Proceed with volume migration? [yes/no] yes
Scheduling report: migrate_volume_group.40.rpt on switch bstnA
```

migrates 'medarcv~/test_results' to volume group 3. See Figure 37.2 for a sample report. This is a direct volume, so the report does not include any information about metadata.

Related Commands

volume-group
cancel migrate-volume
show volume-group
show reports

Figure 37.2 *Sample Report: nsck migrate-volume (Direct Volume without Metadata)*

```
bstnA# show reports migrate_volume_group.40.rpt
**** Migrate Volume Group Report: Started at 01/19/2012 01:51:58 -0500 ****
**** Software Version: 6.02.000.14314 (Jan 16 2012 20:04:23) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

Namespace:                  medarcv
Volume:                     /test_results
Migrating from:         2
Migrating to:           3

**** Total processed:                0
**** Elapsed time:           00:00:00
**** Migrate Volume Group Report: DONE at 01/19/2012 01:51:58 -0500 ****
```

nsck ... rebuild

Purpose

The namespace-check (*nsck*) utility finds and repairs inconsistencies between a managed volume's metadata and actual files on back-end storage. An nsck *rebuild* operation deletes all metadata and re-imports all shares. Use the nsck rebuild command to re-initialize all managed volumes in a namespace, or a particular volume.

Mode

priv-exec

Security Role(s)

storage-engineer or crypto-officer

Syntax

`nsck` *`name`* `rebuild [volume` *`volname`*`] [force]`

name (1-30 characters) is the name of the namespace to rebuild. Use the show namespace command to see a list of configured namespaces.

volume *volname* (optional, 1-256 characters) limits the rebuild to the chosen volume (for example, '/home').

force (optional) indicates that you want to force the destage even if another nsck job is currently under way. You can force any nsck job with this option except nsck ... migrate-metadata or nsck ... migrate-volume; you must cancel those jobs explicitly (with cancel migrate-metadata or cancel migrate-volume) before you can rebuild the volume.

Default(s)

None.

Guidelines: Determining if a Rebuild is Necessary

If clients get front-end-service (NFS or CIFS) failures for a particular managed volume, the volume's metadata might be inconsistent the files on the back end. As further indication of inconsistency, search through the syslog file for DB_RUNRECOVERY messages. Use show logs to view the current syslog, or use grep syslog to search the syslog for certain patterns. Use show logs to view all log files, including backups, in the logs directory.

You can use the nsck ... report inconsistencies command to check the namespace, volume, or share for inconsistencies without changing any metadata. To attempt to repair inconsistencies, you can use the sync files command to synchronize the metadata with actual files, or the sync directories command to find any new directories and add them to the metadata. If a sync fails, nsck ... rebuild is a last resort. This operation stops all volume processing, rebuilds all metadata, and re-syncs all filer subshares; this causes a complete service interruption.

Guidelines: Modify Flag

Each managed volume has a "modify" flag that you can set with the gbl-ns-vol modify command. This flag, if raised, allows the volume to rename any imported files or directories with the same pathname as previously-imported files. If lowered, such files and directories are not imported *and* clients cannot write to the volume.

The nsck ... rebuild command lowers this flag for each managed volume that it rebuilds. By default, it keeps the flag lowered after the rebuild is over. This is a safety feature to avoid renamed files when they might not be expected. To allow modifications, you can run the gbl-ns-vol reimport-modify command for each desired volume. This command tells nsck to re-instate the modify flag after the rebuild.

Guidelines: Directory Mastership After the Rebuild

If your volume has tiered shares, where some shares reside on faster and/or more reliable filers than others, we recommend placing all of your *master directories* onto your Tier 1 shares. Whenever two or more directories have the same name and path, one is considered *master* and the others are called *stripes*. The volume's clients cannot access a directory if its master instance is unavailable from the back end, so it is important to keep the masters on reliable filers.

Use a place-rule to select all directories in the volume, place them onto one or more Tier 1 shares, and promote them to master. The rebuild process follows all directory-place rules during the re-import. Additionally, set all of your Tier 1 shares to import priority 1, to force the volume to scan those shares before the lower-priority shares. This ensures that all master directories for the volume are on Tier 1 shares.

Guidelines: Running the Command

The CLI prompts for confirmation before disabling the namespace or volume: you must answer **yes** to proceed with the rebuild operation.

The optional **force** flag is useful for working through a hung nsck job. An nsck job may hang if one of the back-end shares or filers goes offline at the wrong time.

All NFS clients that use the volume(s) must unmount and remount after the rebuild is finished.

Samples

```
bstnA# nsck ns rebuild
Reimport all shares in the namespace ns? [yes/no] yes
bstnA#
```

rebuilds the "ns" namespace.

```
bstnA# nsck archives rebuild force
```

forces a rebuild in the "archives" namespace.

```
bstnA# nsck archives rebuild volume /etc
```

rebuilds the "/etc" volume in the "archives" namespace.

Related Commands

grep
nsck ... report inconsistencies
sync files
show namespace
show logs
place-rule

nsck ... report dir-structure

Purpose A managed volume's *metadata* is its background information about its files and directories, such as their physical locations on back-end filers. You can use the nsck report dir-structure command to get a report of a volume's directory structure from its metadata. This is faster than the full nsck ... report metadata-only command, and the report provides a quick overview of the volume's directory layout.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax `nsck name report dir-structure [path] [norecurse] [summary] [outputfile outputfile]`

name (1-30 characters) is the namespace on which to report. Use the show namespace command for a list of configured namespaces.

path (optional, 1-256 characters) narrows the scope of the report to a specific virtual path in the namespace (for example, '/home/jrandom'). If you omit this, the command generates one report per volume.

norecurse (optional) specifies to not recurse into subdirectories during the report. The directory listing therefore only includes the root of the volume, or the ***path*** if you chose that option.

summary (optional) removes the table of directories from the report.

outputfile *outputfile* (optional; 1-255 characters) specifies a prefix for a customized file name (for example, 'jrandom_dir.rpt'), as opposed to the default. Use the show reports command to display the file in the maintenance directory.

Default(s) ***path*** - the root of the namespace (all volumes).

outputfile *outputfile* - 'dir_structure.*nsck-id*.rpt.' The *nsck-id* is a unique number assigned to each nsck job.

Guidelines The name of the nsck report appears after you issue the command. Use tail to follow the report as it is written. Use show reports *file-name* to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands.

The non-summary report shows

- the configuration options chosen by this command,
- the back-end shares that hold the directories,
- a table of directories, where each directory appears in one row,
- a summary of the directory structure (maximum and average entries, files, links, and so forth amongst these directories), and
- a listing of total files, directories, and other high-level information.

The summary report omits the table of directories.

If you want to cancel the report before it finishes, use the cancel nsck report command.

For a full report of all files and directories in a managed volume, use the nsck ... report metadata-only command.

Samples

```
bstnA# nsck medarv report dir-structure /rcrds
Scheduling report: dir_structure.34.rpt on switch bstnA
bstnA# ...
```

shows the directory structure for the 'medarcv~/rcrds' volume. The report appears in 'dir_structure.34.rpt.' A sample appears below in Figure 37.3 on page 37-19.

```
bstnA# nsck insur report dir-structure
Scheduling report: dir_structure.43.rpt on switch bstnA
bstnA# ...
```

shows the directory structure for all volumes in the 'insur' namespace. This produces a single report, , for the namespace's single volume. A sample appears in Figure 37.4 on page 37-20.

Related Commands

show namespace
show reports
cancel nsck report
tail
grep
copy ftp
copy scp
copy {nfs|cifs}
copy tftp
delete

Figure 37.3 *Sample Report: nsck medarcv report dir-structure /rcrds*

```
bstnA# show reports dir_structure.34.rpt
**** Directory Structure Report: Started at 01/17/2012 01:16:22 -0500 ****
**** Software Version: 6.02.000.14313 (Jan 13 2012 20:08:22) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:
**** Namespace: medarcv
**** Volume: /rcrds
**** Path: /rcrds

Share                Physical Filer
-------------------  ------------------------------------------------------------
[rx                ]  192.168.25.29:prescriptions
[charts            ]  192.168.25.20:histories
[bulk              ]  192.168.25.27:bulkstorage

**** Legend:
****   NL = No Lock: Was unable to lock parent directory during report.

Flag  Dirs      Files     Links    Symlinks  Path
----  --------- --------- -------- --------  ------------------------------------------------
[  ]         15        17        0        0  /rcrds/
[  ]          5         7        0        0  /rcrds/2004/
[  ]          4         8        0        0  /rcrds/2005/
[  ]          0         0        0        0  /rcrds/2008/
[  ]          0         0        0        0  /rcrds/2010/
[  ]          1         0        0        0  /rcrds/2011/
[  ]          0         3        0        0  /rcrds/fractures/
[  ]          3         7        0        0  /rcrds/2000/
[  ]          0         2        0        0  /rcrds/infections/
[  ]          3         7        0        0  /rcrds/2001/
[  ]          0         2        0        0  /rcrds/RECORD~1/
```

```
[  ]          0         9        0        0  /rcrds/sprains/
[  ]          1         0        0        0  /rcrds/VIP_wing/
[  ]          1         0        0        0  /rcrds/2011/mp3downloads/
[  ]          1         1        0        0  /rcrds/2007/
[  ]          0         2        0        0  /rcrds/records_predating_y2k/
[  ]          0         9        0        0  /rcrds/2000/planA/
[  ]          0         9        0        0  /rcrds/2001/planA/
[  ]          0         9        0        0  /rcrds/2004/planA/
[  ]          0        10        0        0  /rcrds/2005/planA/
[  ]          0         1        0        0  /rcrds/2007/late_test_results./
[  ]          0         9        0        0  /rcrds/2000/planB/
[  ]          0         9        0        0  /rcrds/2001/planB/
[  ]          0         9        0        0  /rcrds/2004/planB/
[  ]          0         9        0        0  /rcrds/2005/planB/
[  ]          0         6        0        0  /rcrds/2000/planQ/
[  ]          0         6        0        0  /rcrds/2001/planQ/
[  ]          0         6        0        0  /rcrds/2004/planQ/
[  ]          0         6        0        0  /rcrds/2005/planQ/
[  ]          0         3        0        0  /rcrds/2004/planR./
[  ]          0         6        0        0  /rcrds/2011/mp3downloads/recoveryStats/
[  ]          0         3        0        0  /rcrds/2004/recoveryStats/
[  ]          0         8        0        0  /rcrds/VIP_wing/fractures/
[  ]          0         4        0        0  /rcrds/recoveryStats/
[  ]          0         2        0        0  /rcrds/2005/recoveryStats/

Directory Structure Summary:
  Maximum Entries in a directory:                  32
  Maximum sub-directories in a directory:          15
  Maximum Files in a directory:                    17
  Maximum Links in a directory:                     0
  Maximum Symlinks in a directory:                  0
  Average Entries in a directory:                   6
  Average Files in a directory:                     5

**** Total Files:                          189
**** Total Directories:                     34
**** Total Leaf Directories:                26
**** Total Interior Directories:             9
**** Total Hard Links (nlink>1):             0
**** Total Symlinks:                         0

**** Total items:                 223
**** Elapsed time:           00:00:00
**** Directory Structure Report: DONE at 01/17/2012 01:16:22 -0500 ****
```

Figure 37.4 *Sample Report: nsck insur report dir-structure*

```
bstnA# show reports dir_structure.43.rpt
**** Directory Structure Report: Started at 01/17/2012 01:17:25 -0500 ****
**** Software Version: 6.02.000.14313 (Jan 13 2012 20:08:22) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:
**** Namespace: insur
**** Volume: /claims
**** Path: /claims

Share                Physical Filer
-------------------  ------------------------------------------------------------
[shr1-old          ]  192.168.25.21 NFS:/vol/vol2/insurance, CIFS:insurance
[shr1-next         ]  192.168.25.51NFS:/root_vdm_4/patient_records, CIFS:patient_records

**** Legend:
```

```
****   NL = No Lock: Was unable to lock parent directory during report.

Flag  Dirs      Files     Links    Symlinks  Path
----  --------- --------- -------- --------  -------------------------------------------------
[  ]         14         8        0        0  /claims/
[  ]          0         0        0        0  /claims/Claims:2001/
[  ]          0         0        0        0  /claims/claims:2005/
[  ]          1        24        0        2  /claims/common/
[  ]          0         4        0        0  /claims/draft_proposals./
[  ]          0        50        0        0  /claims/images/
[  ]          0         0        0        0  /claims/POLICI~2/
[  ]          2         7        0        1  /claims/specs/
[  ]          2         9        0        0  /claims/stats/
[  ]          0         8        0        0  /claims/tools/
[  ]          0         0        0        0  /claims/Tools/
[  ]          0         0        0        0  /claims/y2kclaims/
[  ]          0         0        0        0  /claims/dir0134(U+0134)_shr1-old-12/
[  ]          0         3        0        0  /claims/stats/in_home:2005/
[  ]          0        22        0        0  /claims/specs/common/
[  ]          0         5        0        0  /claims/stats/otj_latest/
[  ]          1         1        0        1  /claims/specs/protos/
[  ]          0         0        0        0  /claims/specs/protos/accepted/
[  ]          0         1        0        0  /claims/common/eastLab/
[  ]          0         9        0        0  /claims/Y2KCLA~1/
[  ]          0         0        0        0  /claims/Y2KCLA~4/

Directory Structure Summary:
  Maximum Entries in a directory:                  50
  Maximum sub-directories in a directory:          14
  Maximum Files in a directory:                    50
  Maximum Links in a directory:                     0
  Maximum Symlinks in a directory:                  2
  Average Entries in a directory:                   8
  Average Files in a directory:                     7

**** Total Files:                         151
**** Total Directories:                    20
**** Total Leaf Directories:               16
**** Total Interior Directories:            5
**** Total Hard Links (nlink>1):            0
**** Total Symlinks:                        4

**** Total items:                171
**** Elapsed time:           00:00:00
**** Directory Structure Report: DONE at 01/17/2012 01:17:25 -0500 ****
```

nsck ... report inconsistencies

Purpose

A managed volume's *metadata* is its background information about its files and directories, such as their physical locations on back-end filers. If a client has direct access to one of the volume's back-end shares (that is, access that does not go through the volume), the volume's metadata is likely to be inconsistent. Metadata inconsistencies can also occur in a multi-protocol (NFS and CIFS) volume without any client intervention: on a multi-protocol filer, a file or directory name may differ between NFS and CIFS.

Use the nsck report inconsistencies command to find (but not repair) any inconsistencies in a namespace's metadata, or to prove that there are none.

Modes

priv-exec

Security Role(s)

storage-engineer or crypto-officer

Syntax

`nsck` *`name`* `report inconsistencies [`*`path`*`] [share` *`share-name`*`] [norecurse] [nofilehandles] [multi-protocol] [outputfile` *`outputfile`*`]`

name (1-30 characters) is the namespace on which to report. Use the show namespace command for a list of configured namespaces.

path (optional, 1-256 characters) narrows the scope of the report to a specific virtual path in the namespace (for example, '/home/jrandom').

share *share-name* (optional, 1-64 characters) limits the report to a single share in the volume.

norecurse (optional) specifies to not recurse into subdirectories during the report.

nofilehandles (optional) specifies to not validate file handles during the report.

multi-protocol (optional) focuses the report on inconsistencies between NFS names and CIFS names. This only applies to a volume in a multi-protocol (NFS and CIFS) namespace.

outputfile *outputfile* (optional; 1-255 characters) specifies a prefix for a customized file name (for example, 'jrandom_inconsistencies.rpt'), as opposed to the default. Use the show reports command to display the file in the maintenance directory.

Default(s)

path - the root of the namespace (all volumes).

outputfile *outputfile* - 'inconsistencies.*nsck-id*.rpt.' The *nsck-id* is a unique number assigned to each nsck job.

share-name - all shares in the volume.

Guidelines

The name of the nsck report appears after you issue the command. Use tail to follow the report as it is written. Use show reports *file-name* to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands.

The report shows various types of inconsistency:

- Metadata disagrees with actual files or directories on the back-end filers,
- A file or directory name matches an alternate "8.3" name on a back-end-CIFS share, or
- A file or directory name matches a filer-generated name (FGN) on a back-end share that supports both CIFS and NFS.

The ARX CLI Maintenance Guide discusses all of these inconsistency types in detail, and describes methods for addressing each inconsistency type.

If this operation finds any metadata inconsistencies, you can use sync files to synchronize the metadata with the filer.

If you want to cancel the report before it finishes, use the cancel nsck report command.

Guidelines: Transitional States

This report also describes any network issues, filer issues, or other unusual problems that temporarily prevent a full listing. These are denoted by the following codes:

- MF - The file is being migrated from one share to another, so it is in a transitional state. If there is an additional FF (Found-File) entry for the same file, the FF indicates that it found an extra copy of the file while it existed (for an instant) on both the source share and the destination share. You can probably disregard the FF state for that file. Run the report later for an accurate assessment of its metadata consistency.
- NL - The nsck software could not lock the parent directory, typically due to heavy client activity. Without this lock, numerous changes (such as file additions, edits, and deletions) can be underway, making it possible to mistakenly assume an inconsistency. Run the report later on this directory to get an accurate measure of its inconsistencies. You can use the optional ***path*** in this command to focus on a particular path in the volume.
- FE - The nsck software lost contact with the filer during the report. Check for connectivity issues with this filer. Once the problem is corrected, run the report again.
- FO -The filer is offline. Address the problem at the filer and rerun this report.
- DC - A client issued a "delete-on-last-close" command for this file, but the filer deleted it early. EMC Celerra servers have shown this behavior in lab testing. The client application has already announced its intention to delete the file, so this condition is cleared when the client application closes the file.

Samples

```
bstnA# nsck wwmed report inconsistencies
Scheduling report: inconsistencies.2.rpt on switch bstnA
bstnA# ...
```

shows the metadata for all managed volumes in the 'wwmed' namespace. The report appears in 'inconsistencies.2.rpt.' A successful sample appears below in Figure 37.5 on page 37-24.

```
bstnA# nsck medarcv report inconsistencies outputfile medarcv5-6
Scheduling report: medarcv5-6.3.rpt on switch bstnA
bstnA# ...
```

checks the 'medarcv' namespace and puts the report in a file, 'medarcv5-6.3.rpt.'

```
bstnA# nsck insur report inconsistencies /claims
Scheduling report: inconsistencies.25.rpt on switch bstnA
bstnA# ...
```

shows the metadata for /claims in the 'insur' namespace. The report appears in 'inconsistencies.25.rpt.' An unsuccessful sample appears in Figure 37.6 on page 37-25.

Related Commands

show namespace
sync files
show reports
cancel nsck report
tail
grep
copy ftp
copy scp
copy {nfs|cifs}
copy tftp
delete

Figure 37.5 *Sample Report: nsck report inconsistencies*

```
bstnA# show reports inconsistencies.1.rpt
**** Inconsistencies Report: Started at 01/19/2012 01:27:44 -0500 ****
**** Software Version: 6.02.000.14314 (Jan 16 2012 20:04:23) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:
**** Namespace: wwmed
**** Volume: /acct
**** Path: /acct

Share                Physical Filer
-------------------  ------------------------------------------------------------
[budget           ]  192.168.25.19:/exports/budget
[bills            ]  192.168.25.25:/work1/accting
[bills2           ]  192.168.25.23:/exports/acct2
[it5              ]  192.168.25.24:/lhome/it5

**** Legend:
****   LF = File exists in the metadata, but is missing from the physical filer.
****   LD = Directory exists in the metadata, but is missing from the physical filer.
****   FF = File exists on the physical filer, but is missing from the metadata.
****   FD = Directory exists on the physical filer, but is missing from the metadata.
****   LL = File is a symlink in the metadata, but is a regular file on the filer.
```

```
****   FL = File is a symlink on the filer, but is a regular file in the metadata.
****   IF = Filehandles in the metadata do not match the filehandles on the physical filer.
****   MF = The file is currently being migrated.
****   NL = Unable to lock parent directory during report.
****   FE = Error contacting filer during report.
****   FO = Filer Offline: The filer is offline or disabled.
****   F8 = A file name matches a CIFS alternate "8.3" name on another share.
****   D8 = A directory name matches a CIFS alternate "8.3" name; its contents will be skipped.
****   DC = A client has the file or directory open for delete-on-close, but the filer has already
deleted it.
****   SD = Striped leaf directory found on filer, expected on other shares.
****   SL = File is a symlink.
****   UT = Name contains characters that are invalid UTF-8; must solve issue directly on the filer
****   IS = Inconsistent attributes on one of this directory's stripes (discovered)
****   MI = Attributes are consistent, metadata marked as inconsistent
****   SI = Attributes are inconsistent, metadata not marked as inconsistent


Type                        Share                Path
--------------------------  -------------------  -------------------------------

**** Total Found Items:                   0
**** Total Lost Items:                    0
**** Total Invalid Filehandles:           0
**** Total Migrating Files:               0
**** Total Deleted Before Close:          0
**** Total Locking Errors:                0
**** Total Filer Errors:                  0
**** Total 8.3 Errors:                    0
**** Total Found Stripes:                 0
**** Total Inconsistent Attrs:            0

**** Total processed:          4,432
**** Elapsed time:           00:00:04
**** Inconsistencies Report: DONE at 01/19/2012 01:27:48 -0500 ****
```

Figure 37.6 *Sample Report: nsck report inconsistencies (with Inconsistent Files)*

```
bstnA# show reports inconsistencies.25.rpt
**** Inconsistencies Report: Started at Tue Apr 27 01:27:41 2010 ****
**** Software Version: 5.02.000.12561 (Apr 23 2010 20:15:19) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:
**** Namespace: insur
**** Volume: /claims
**** Path: /claims

Share                Physical Filer
-------------------  ------------------------------------------------------------
[shr1-old          ]  192.168.25.21 NFS:/vol/vol1/NTFS_QTREE/insurance, CIFS:insurance
[shr1-next         ]  192.168.25.51NFS:/root_vdm_4/patient_records, CIFS:patient_records

**** Legend:
****   LF = File exists in the metadata, but is missing from the physical filer.
****   LD = Directory exists in the metadata, but is missing from the physical filer.
****   FF = File exists on the physical filer, but is missing from the metadata.
****   FD = Directory exists on the physical filer, but is missing from the metadata.
****   LL = File is a symlink in the metadata, but is a regular file on the filer.
****   FL = File is a symlink on the filer, but is a regular file in the metadata.
****   IF = Filehandles in the metadata do not match the filehandles on the physical filer.
****   MF = The file is currently being migrated.
```

```
****   NL = Unable to lock parent directory during report.
****   FE = Error contacting filer during report.
****   FO = Filer Offline: The filer is offline or disabled.
****   F8 = A file name matches a CIFS alternate "8.3" name on another share.
****   D8 = A directory name matches a CIFS alternate "8.3" name; its contents will be skipped.
****   DC = A client has the file or directory open for delete-on-close, but the filer has already
deleted it.
****   SD = Striped leaf directory found on filer, expected on other shares.
****   SL = File is a symlink.
****   UT = Name contains characters that are invalid UTF-8; must solve issue directly on the filer
****   IS = Inconsistent attributes on one of this directory's stripes (discovered)
****   MI = Attributes are consistent, metadata marked as inconsistent
****   SI = Attributes are inconsistent, metadata not marked as inconsistent


Type                        Share                Path
--------------------------  -------------------  -------------------------------
[   F8                  ]  [shr1-old          ]  /OVERDU~1.DOC -> overdueclaimsmemo.doc
[   D8                  ]  [shr1-old          ]  /Y2KCLA~1/ -> y2kclaims

**** Total Found Items:                   0
**** Total Lost Items:                    0
**** Total Invalid Filehandles:           0
**** Total Migrating Files:               0
**** Total Deleted Before Close:          0
**** Total Locking Errors:                0
**** Total Filer Errors:                  0
**** Total 8.3 Errors:                    2
**** Total Found Stripes:                 0
**** Total Inconsistent Attrs:            0

**** Total processed:             171
**** Elapsed time:           00:00:00
**** Inconsistencies Report: DONE at Tue Apr 27 01:27:41 2010 ****
```

nsck ... report metadata-only

Purpose

A managed volume's *metadata* is its background information about its files and directories, such as their physical locations on back-end filers. Use the nsck report metadata-only command to show a namespace's or volume's metadata.

Modes

priv-exec

Security Role(s)

storage-engineer or crypto-officer

Syntax

`nsck` *`name`* `report metadata-only [`*`path`*`] [share` *`share-name`*`] [norecurse] [outputfile` *`outputfile`*`]`

name (1-30 characters) specifies a namespace on which to report. Use the show namespace command for a list of configured namespaces.

path (optional, 1-256 characters) narrows the scope of the report to a specific managed volume or path in the namespace (for example, '/eng/share').

share *share-name* (optional, 1-64 characters) limits the report to a single share in the volume.

norecurse (optional) specifies to not recurse into subdirectories during the report.

outputfile *outputfile* (optional, 1-255 characters) specifies a prefix for a customized file name (for example, 'jrandom_inconsistencies.rpt'), as opposed to the default. Use the show reports command to display the file in the maintenance directory.

Default(s)

path - the root of the namespace (all managed volumes).

share-name - all shares in the volume.

outputfile *outputfile* - 'metadata_only'

Guidelines

When a share is first imported into a managed volume, a report similar to this one is generated. You can compare that report to this one to see the effects of policy on your file locations. The import reports are accessible through show reports. Their names have the following formats:

`import.`*`share-id`*`.`*`share-name`*`.`*`job-id`*`.rpt`

Use the find command to show the physical location of a single file.

The name of the nsck report appears after you issue the command. Use tail to follow the report as it is written. Use show reports *file-name* to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands.

If you want to cancel the report before it finishes, use the cancel nsck report command.

You can use the nsck ... report dir-structure command for a faster report that focuses on the volume's directory structure.

Guidelines: Report Content

The report lists all files (labeled FL) and directories (labeled DR) recorded in the volume's metadata. An SL indicates that the file or directory is actually a symbolic link, or symlink. Most of the other labels indicate a naming issue; you can examine these thoroughly with the nsck ... report inconsistencies command.

The NL condition is typically transitory: it indicates that the nsck software is currently unable to get a lock on the parent directory for the given file or directory. This typically indicates high client activity (such as file creates and deletes) in the parent directory, which may lead to temporarily-inaccurate metadata readings in the report. The volume updates its metadata as soon as it is able to get the lock.

Guidelines: File Tracking

This identifies the back-end locations of all the volume's files as of now; for file locations in the past, you can use *file tracking*. This can be useful for finding the correct filer-backup tape for a lost or compromised file.

The file-tracking feature requires some configuration before you can make any file queries. Specifically, a managed volume requires a snapshot rule to regularly copy its configuration and metadata to a file-history archive. After some copies of metadata have been archived, you can use the show file-history virtual-service command to query file locations at different dates.

Samples

```
bstnA# nsck wwmed report metadata-only
Scheduling report: metadata_only.5.rpt on switch bstnA
bstnA# ...
```

shows the metadata for the 'wwmed' namespace. The report appears in 'metadata_only.5.rpt.' A sample report appears in Figure 37.7 on page 37-29.

```
bstnA# nsck archives report metadata-only /arch
```

checks the virtual path, '/arch,' in the archives namespace. Note that this path could span multiple volumes.

```
bstnA# nsck archives report metadata-only /arch share fs6
```

checks only the files from /arch that are stored on the 'fs6' share.

```
bstnA# nsck acopia3 report metadata-only outputfile 6-5mdata
Scheduling report: 6-5mdata.15.rpt on switch bstnA
bstnA# ...
```

checks the 'acopia3' namespace and puts the report in a file, '6-5mdata.15.rpt.'

Related Commands

show namespace
show reports
nsck ... report dir-structure
cancel nsck report
find
tail
grep
copy ftp
copy scp
copy {nfs|cifs}
copy tftp
delete

Figure 37.7 *Sample Report: nsck metadata-only*

```
bstnA# show reports metadata_only.5.rpt
**** Metadata-Only Report: Started at Tue Apr 27 01:26:10 2010 ****
**** Software Version: 5.02.000.12561 (Apr 23 2010 20:15:19) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:
**** Namespace: wwmed
**** Volume: /acct
**** Path: /acct

Share                Physical Filer
-------------------  ------------------------------------------------------------
[budget            ]  192.168.25.19:/exports/budget
[bills             ]  192.168.25.25:/work1/accting
[bills2            ]  192.168.25.23:/exports/acct2
[it5               ]  192.168.25.24:/lhome/it5

**** Legend:
****   FL = File: The reported entry is a file.
****   DR = Directory: The reported entry is a directory.
****   SL = Symlink: The reported entry is a symbolic link.
****   LN = Link: The reported entry has a link count greater than one.
****   NL = No Lock: Was unable to lock parent directory during report.
****   CC = NFS case-blind name collision.
****   IC = Name contains invalid CIFS characters.
****   FN = Name may conflict with a filer-generated name.
****   SP = A persistent split is registered in the metadata, due to a FGN.
****   NF = Name is only accessible to NFS clients.
****   IA = Inconsistent attributes between this directory's master and stripes (recorded).
****   IS = Inconsistent attributes on this specific directory stripe (recorded).


Type                   Share                Path
--------------------  -------------------  -------------------------------------
[FL                 ]  [bills2            ]  /layer3.fm.lck
[   DR              ]  [bills2            ]  /planner
[FL                 ]  [bills2            ]  /layer2.fm.lck
[FL                 ]  [budget            ]  /cliOperator.book
[   DR              ]  [budget            ]  /production
[FL                 ]  [budget            ]  /finance_sites.html
[FL                 ]  [bills             ]  /shellCmds_winXP.fm
[FL                 ]  [budget            ]  /sampleNet.html
[FL                 ]  [budget            ]  /shellCmds_winXP.fm.lck
. . .
[FL                 ]  [bills2            ]  /receipts/giftShop/sep2006.csv
[FL                 ]  [bills             ]  /receipts/giftShop/xmas2009.csv
[FL                 ]  [bills2            ]  /receipts/giftShop/aug2009.csv

**** Total Files:                      3,994
**** Total Directories:                  438
**** Total Hard Links (nlink>1):           0
**** Total Symlinks:                       9
**** Total Locking Errors:                 0

**** Total items:              4,432
**** Elapsed time:           00:00:00
**** Metadata-Only Report: DONE at Tue Apr 27 01:26:10 2010 ****
```

nsck ... report symlinks

Purpose A managed volume's *metadata* is its background information about its files and directories, such as their physical locations on back-end filers. The metadata also contains a flag for all NFS symbolic link (symlink) files. A multi-protocol (CIFS and NFS) volume uses this flag to support symlink de referencing for its CIFS clients. Use the nsck report symlinks command to show all of the NFS symlinks that are identified in the current volume's metadata.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`nsck` *`name`* `report symlinks [`*`path`*`] [share` *`share-name`*`] [norecurse] [outputfile` *`outputfile`*`]`**

name (1-30 characters) specifies a namespace on which to report. Use the show namespace command for a list of configured namespaces.

path (optional, 1-256 characters) narrows the scope of the report to a specific managed volume or path in the namespace (for example, '/claims/common' for the /common directory in the /claims volume).

share *share-name* (optional, 1-64 characters) limits the report to a single share in the volume.

norecurse (optional) specifies to not recurse into subdirectories during the report.

outputfile *outputfile* (optional, 1-255 characters) specifies a prefix for a customized file name (for example, 'jrandom_symlinks.rpt'), as opposed to the default. Use the show reports command to display the file in the maintenance directory.

Default(s) ***path*** - the root of the namespace (all managed volumes).

share-name - all shares in the volume.

outputfile *outputfile* - 'symlinks'

Guidelines The name of the nsck report appears after you issue the command. Use tail to follow the report as it is written. Use show reports *file-name* to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands.

If you want to cancel the report before it finishes, use the cancel nsck report command.

Guidelines: Report Content

The report lists all symlinks (labeled SL) recorded in the volume's metadata. For each symlink, one of the following codes provides some more information:

- RF indicates that the symlink resolves to (or points to) a file.
- RD shows that the symlink resolves to a directory.
- RS means that the symlink resolves to another symlink.
- RX indicates a symlink error: the symlink points to a non-existent file or directory.
- ER indicates an error in querying the filer for symlink information. The report includes specific error information in parentheses.

The NL condition is typically transitory: it indicates that the nsck software is currently unable to get a lock on the parent directory for the given file or directory. This typically indicates high client activity (such as file creates and deletes) in the parent directory, which may lead to temporarily-inaccurate metadata readings in the report. The volume updates its metadata as soon as it is able to get the lock.

Sample

```
bstnA# nsck insur report symlinks outputfile insur_symlinks
Scheduling report: insur_symlinks._claims.rpt on switch bstnA
bstnA# ...
```

shows all symlinks in the 'insur' namespace. The report appears in 'insur_symlinks._claims.rpt.' A sample report appears in Figure 37.8.

Related Commands

cifs deny-symlinks
show namespace
show reports
cancel nsck report
find
tail
grep
copy ftp
copy scp
copy {nfs|cifs}
copy tftp
delete

Figure 37.8 *Sample Report: nsck symlinks*

```
bstnA# show reports insur_symlinks._claims.rpt
**** Symlinks Report: Started at Sat Feb 27 01:27:07 2010 ****
**** Software Version: 5.02.000.12541 (Feb 23 2010 20:12:44) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:
**** Namespace: insur
**** Volume: /claims
**** Path: /claims

Share                Physical Filer
-------------------  ------------------------------------------------------------
[shr1-old          ]  192.168.25.21 NFS:/vol/vol1/NTFS_QTREE/insurance, CIFS:insurance
[shr1-next         ]  192.168.25.51 NFS:/root_vdm_4/patient_records, CIFS:patient_records

**** Legend:
****   SL = Symlink: The reported entry is a symbolic link.
****   RF = File: The symlink resolves to a file in the metadata.
****   RD = Directory: The symlink resolves to a directory in the metadata.
```

```
****   RS = Symlink: The symlink resolves to another symlink in the metadata.
****   RX = Dangling: The symlink is dangling; the object to which it points is missing.
****   ER = Error dereferencing link: Details indicated in parentheses.
****   NL = No Lock: Was unable to lock parent directory during report.

Type        Share                Path
----------  -------------------  ------------------------------------
[SL RD   ]  [shr1-old          ]  /common/reagentLists -> /claims/specs
[SL RX   ]  [shr1-old          ]  /common/acmeIns.txt -> /claims/stats/acmeIns.txt
[SL RD   ]  [shr1-old          ]  /specs/stats -> /claims/stats
[SL RD   ]  [shr1-old          ]  /specs/protos/rejected -> /claims/specs/common

**** Total Symlinks to Files:                 0
**** Total Symlinks to Directories:           3
**** Total Symlinks to Symlinks:              0
**** Total Dangling Symlinks:                 1
**** Total Symlinks:                          4
**** Total Dereferencing Errors:              0
**** Total Locking Errors:                    0

**** Total items:                   0
**** Elapsed time:           00:00:00
**** Symlinks Report: DONE at Sat Feb 27 01:27:07 2010 ****
```

nsck ... sync directories

Purpose If someone goes around a managed volume to directly create a directory on a back-end filer, the volume metadata does not include the new directory. The volume's clients therefore cannot see or use the new directory. (This may be necessary for directories with special features that can only be provisioned at the filer itself, such as disk quotas.) Use the nsck ... sync directories command to find any new subdirectories under a known directory. This is equivalent to the sync directories command.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`nsck`** ***`namespace`*** **`sync directories`** ***`dir-path`*** **`[share`** ***`share-name`*****`]`**

namespace (1-30 characters) specifies the namespace to sync. Use the show namespace command for a list of configured namespaces.

dir-path (1-256 characters) narrows the scope of the sync operation to a specific managed volume and path (for example, "/vol/var/log"). The path must describe a valid volume in the namespace, or one of its subdirectories. Use forward slashes (/) for NFS or CIFS paths.

share ***share-name*** (optional, 1-64 characters) limits the sync and repair to a single share in the volume.

Default(s) **share** ***share-name*** is unspecified, so all shares in the volume are synchronized and repaired.

Guidelines This invokes the same software used by the sync directories command. Follow the Guidelines for that command.

Sample

```
bstnA# nsck medarcv sync directories /lab_equipment path /
Scheduling sync directories operation on switch bstnA, report name:
sync.9._lab_equipment.rpt
bstnA#
```

checks for unrecorded directories in the root of the "medarcv~/lab_equipment" volume.

Related Commands

sync directories
show sync
wait-for sync
cancel sync
clear sync
nsck ... report inconsistencies
show namespace
show reports
nsck ... rebuild

nsck ... sync files

Purpose If a file changes on a back-end filer without the managed volume's awareness, the volume metadata is *inconsistent* for that file. Use the nsck ... sync files command to report on metadata inconsistencies and repair them. This is equivalent to the sync files command.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`nsck` *`namespace`* `sync files` *`dir-path`* `[share` *`share-name`*`] [recurse] [rename-files]`**

namespace (1-30 characters) specifies the namespace to sync. Use the show namespace command for a list of configured namespaces.

dir-path (1-256 characters) narrows the scope of the sync operation to a specific managed volume and path (for example, "/vol/var/log"). The path must describe a valid volume in the namespace, or one of its subdirectories. Use forward slashes (/) for NFS or CIFS paths.

share *share-name* (optional, 1-64 characters) limits the sync and repair to a single share in the volume.

recurse (optional) causes the sync operation to descend into subdirectories.

rename-files (optional) determines that, if a file-collision is discovered, the sync operation can rename the newly-discovered file. All collisions and renames are included in the sync report.

Default(s) **recurse** is off; the sync only examines the ***dir-path*** without descending into any child directories.

rename-files is off; collisions are noted in the sync report, and the metadata is not repaired.

share *share-name* is unspecified, so all shares in the volume are synchronized and repaired.

Guidelines This invokes the same software used by the sync files command. Follow the Guidelines for that command.

Sample

```
bstnA# nsck wwmed sync files /acct recurse
Scheduling sync files operation on switch bstnA, report name:
sync.1._acct.rpt
bstnA#
```

repairs the top-level directory in the "wwmed~/acct" volume. This descends into child directories but does not repair any renamed files that have collided with known files.

Related Commands

sync files
show sync
wait-for sync
cancel sync
nsck ... report inconsistencies
show namespace
show reports
nsck ... rebuild

show nsck

Purpose Use the show nsck command to show the current state of all nsck report, nsck rebuild, and sync operations.

Modes (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show nsck [nsck-id | namespace name | all]`

nsck-id (optional) identifies one nsck or sync job to show. The job ID appears if you run show nsck without any options.

name (optional, 1-30 characters) shows the nsck and/or sync job(s) in one namespace.

all (optional) shows a detailed view of all nsck and sync jobs.

Default(s) None.

Guidelines If you do not include any of the options, the command displays the nsck and sync records for all namespaces.

The show nsck command shows the operation ID (Op ID), the type of operation (Op Type), the namespace and path where the operation is running, and the current Status of the operation. Possible status lines appear below.

The show nsck all and show nsck *nsck-id* commands add detail to the output, including the name of the job report and the switch where the report resides. The latter is useful in redundant pairs, where a report could be generated on either peer.

The report data from the nsck commands is stored in a .rpt file in the "reports" directory. Use show reports to view files from this directory.

Use the clear nsck command to clear one nsck job (or all nsck jobs). The clear sync command clears sync jobs.

Guidelines: Possible Status Conditions

Complete indicates a successful nsck operation.

Pending shows that the nsck operation is starting.

Rebuild in progress shows that the nsck operation is rebuilding a namespace or managed volume (nsck ... rebuild).

Destage in progress appears for the nsck ... destage command.

Report in progress indicates that the nsck job is writing a report, such as an inconsistencies report (nsck ... report inconsistencies) or metadata report (nsck ... report metadata-only).

Cancelled means that someone stopped the nsck operation.

Migrate in progress (*x*%) indicates that someone is migrating a managed-volume's metadata from one metadata share to another. The percentage shows the progress of the migration.

Cancel in progress appears if someone stopped the metadata migration.

Synchronizing appears for the sync files command.

Pending Cancel and then

Canceling appears if someone stopped the synchronization job with the cancel sync command.

Guidelines: Possible Errors

Any of these Status comments indicate that the nsck job failed.

Error: Path not found - A virtual tree has a file or directory that is missing from its physical counterpart.

Error: Out of memory - The namespace software ran out of memory during the nsck operation.

Error: Temporary database - The namespace software encountered an error creating, opening, or deleting a temporary database.

Error: Resolving database entry - An error occurred while resolving an entry in the operation's temporary database.

Error: Bad storage type - The namespace software could not determine whether a file-system entry was a file or a directory.

Error: Attributes not found - The namespace software failed to access file attributes from one of the back-end-directory trees.

Error: Incomplete - The switch rebooted before the nsck job was finished.

Error: Incomplete at force time - An nsck ... rebuild force interrupted the nsck job.

Error: Bad report type specified - The user specified an invalid report name in the nsck command.

Error: Migration failed - The metadata migration failed. All metadata still resides on its original share. You can use the nsck ... migrate-metadata command to invoke a metadata migration.

Error: Unable to open report - This is an internal error. Contact customer support.

Samples

`bstnA# show nsck`

displays the state of the latest nsck and sync operations for all namespaces. Figure 37.9 on page 37-38 shows sample output.

`bstnA# show nsck namespace medarcv`

displays the state of the latest nsck/sync operations for the "medarcv" namespace. Figure 37.10 on page 37-39 shows sample output.

`bstnA# show nsck 1`

shows details about the nsck operation with ID 1. See Figure 37.11 on page 37-39 for sample output.

Related Commands

clear nsck
nsck ... report dir-structure
nsck ... report inconsistencies
nsck ... report metadata-only
nsck ... migrate-metadata
nsck ... destage
nsck ... rebuild
show reports

Figure 37.9 *Sample Output: show nsck*

```
bstnA# show nsck

Op Id  Op Type              Namespace:Path                          Status
------ -------------------- --------------------------------------- ------------------
1      report               insur:/claims                           Complete
2      report               wwmed:/acct                             Complete
3      report               wwmed:/acct                             Complete
4      report               wwmed:/acct/payable                     Complete
5      report               wwmed:/acct                             Complete
6      report               wwmed:/acct                             Complete
7      report               wwmed:/acct/payable                     Complete
8      report               wwmed:/acct                             Complete
9      report               wwmed:/acct                             Complete
10     report               wwmed:/acct                             Complete
11     sync                 wwmed:/acct/payable                     Complete
12     report               medarcv:/lab_equipment                  Complete
13     report               medarcv:/rcrds                          Complete
14     report               medarcv:/lab_equipment                  Complete
15     report               medarcv:/test_results                   Complete
16     report               medarcv:/rcrds                          Complete
17     report               medarcv:/lab_equipment                  Complete
18     report               medarcv:/test_results                   Complete
19     report               medarcv:/rcrds                          Complete
20     report               medarcv:/lab_equipment                  Complete
21     report               medarcv:/test_results                   Complete
22     report               medarcv:/rcrds                          Complete
23     report               medarcv:/lab_equipment                  Complete
24     report               medarcv:/test_results                   Complete
25     report               medarcv:/rcrds                          Complete
26     sync                 medarcv:/rcrds/sprains                  Complete
27     sync                 medarcv:/rcrds/                         Complete
28     sync                 medarcv:/lab_equipment/acme             Complete
29     sync                 medarcv:/lab_equipment/                 Complete
```

```
30     report               insur:/claims                          Complete
31     report               insur:/claims                          Complete
32     report               insur:/claims/tools                    Complete
33     report               insur:/claims                          Complete
34     report               insur:/claims                          Complete
35     report               insur:/claims/tools                    Complete
36     report               insur:/claims                          Complete
37     report               insur:/claims                          Complete
38     report               insur:/claims                          Complete
39     md-migrate           wwmed:/acct                            Complete
40     migrate-volume-group medarcv:/test_results                  Complete
```

Figure 37.10 *Sample Output: show nsck namespace medarcv*

```
bstnA# show nsck namespace medarcv

Op Id  Op Type              Namespace:Path                         Status
------ -------------------  -------------------------------------- ------------------
12     report               medarcv:/lab_equipment                 Complete
13     report               medarcv:/rcrds                         Complete
14     report               medarcv:/lab_equipment                 Complete
15     report               medarcv:/test_results                  Complete
16     report               medarcv:/rcrds                         Complete
17     report               medarcv:/lab_equipment                 Complete
18     report               medarcv:/test_results                  Complete
19     report               medarcv:/rcrds                         Complete
20     report               medarcv:/lab_equipment                 Complete
21     report               medarcv:/test_results                  Complete
22     report               medarcv:/rcrds                         Complete
23     report               medarcv:/lab_equipment                 Complete
24     report               medarcv:/test_results                  Complete
25     report               medarcv:/rcrds                         Complete
26     sync                 medarcv:/rcrds/sprains                 Complete
27     sync                 medarcv:/rcrds/                        Complete
28     sync                 medarcv:/lab_equipment/acme            Complete
29     sync                 medarcv:/lab_equipment/                Complete
40     migrate-volume-group medarcv:/test_results                  Complete
```

Figure 37.11 *Sample Output: show nsck 1*

```
bstnA# show nsck 1

Op Type:     report
Op Id:       1
Report Name: insur_fgns_1._claims.rpt
Switch:      bstnA
Namespace:   insur
Path:        /claims
Status:      Complete
```

show sync

Purpose Use the show sync command to show the results of all sync operations.

Modes (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show sync {files | directories}
show sync {files | directories} namespace
show sync {files | directories} namespace volume vol-path
show sync {files | directories} namespace volume vol-path path
path
```

files | directories is a required choice. This selects the type of sync operation to show.

namespace (optional, 1-30 characters) is the name of a namespace. Use show namespace for a list of all configured namespaces. If you omit this (and all other options), the output includes all recorded sync-files or sync-directories operations.

vol-path (optional, 1-1024 characters) focuses on a particular volume. Use forward slashes (/) for all volumes, including CIFS.

path (optional, 1-1024 characters) narrows the focus further, to a particular path within the volume. As with volumes, use forward slashes in all paths.

Default(s) None.

Guidelines This command shows all sync operations, past and present. They are organized by namespace, volume, and in-volume path where the sync ran:

Namespace is a namespace where at least one sync operation occurred.

Volume is one volume where someone invoked a sync.

Path is the path where a sync operation occurred. This contains details about the operation:

Options are "recurse" and/or "rename-files." This does not appear unless the user selected one of these options. They are options in the sync files command only.

Type appears only for auto-sync operations. An auto-sync occurs only in a volume that is enabled to run them with auto sync files. A client triggers an auto-sync operation by getting an error that indicates metadata inconsistency. The value for this field is "Automatic."

Status is "Pending," "Syncronizing," "Pending Cancel," "Canceling," "Success," or "Failed." A typical sync operation goes from "Pending" to "Syncronizing" to "Success." Use the cancel sync command to cancel the sync operation. For details on the sync operation's progress, you can look at the sync report, described below.

Report Name is the name of a detailed report about the sync operation. You can use the show reports command to view the report.

Share only appears if the share was specified (see the documentation for sync files or sync directories).

Guidelines (Cont.)

Lost Items are the files or directories that were recorded in metadata but are missing from the back-end filer.

Found Items are the reverse: these were found on the filer and needed to be added to the metadata.

Synchronized Items is the total repaired files and directories. Ideally, this should be the sum of the lost and found items above.

Processed Items is the total number of files and directories processed.

This output grows with every run of the sync files or sync directories command. Use clear sync to clear all sync operations from this output (without deleting any of the detailed sync reports; use delete reports to remove those).

Samples

`bstnA# show sync files`

displays the state of all sync-files operations for all namespaces. Figure 37.12 shows sample output.

`bstnA# show sync directories`

displays the state of all sync-directories operations. Figure 37.13 on page 37-42 shows sample output.

`bstnA# show sync files medarcv volume /rcrds`

displays the state of all sync-files operations for a CIFS volume, "medarcv~/rcrds." Figure 37.14 on page 37-42 shows sample output.

Related Commands

sync files
sync directories
clear sync
cancel sync

***Figure 37.12** Sample Output: show sync files*

```
bstnA# show sync files

Namespace: wwmed
    Volume: /acct
        Path: /payable
            Options:            recurse rename-files
            Status:             Success
            Report Name:        sync.1._acct.rpt (bstnA)
            Share:              bills
            Lost Items:         0
            Found Items:        0
            Synchronized Items: 0
            Processed Items:    21


Namespace: medarcv
    Volume: /rcrds
        Path: /sprains
            Options:            recurse rename-files
            Status:             Success
            Report Name:        sync.2._rcrds.rpt (bstnA)
            Lost Items:         0
            Found Items:        0
            Synchronized Items: 0
```

```
            Processed Items:    11

        Path: /
            Options:            recurse rename-files
            Status:             Success
            Report Name:        sync.3._rcrds.rpt (bstnA)
            Lost Items:         1
            Found Items:        0
            Synchronized Items: 1
            Processed Items:    245
```

Figure 37.13 *Sample Output: show sync directories*

```
bstnA# show sync directories

Namespace: medarcv
    Volume: /lab_equipment
        Path: /acme
            Status:             Success
            Report Name:        sync.4._lab_equipment.rpt (bstnA)
            Share:              leased
            Lost Items:         0
            Found Items:        1
            Synchronized Items: 1
            Processed Items:    0

        Path: /
            Status:             Success
            Report Name:        sync.5._lab_equipment.rpt (bstnA)
            Lost Items:         0
            Found Items:        1
            Synchronized Items: 1
            Processed Items:    0
```

Figure 37.14 *Sample Output: show sync files medarcv volume /rcrds*

```
bstnA# show sync files medarcv volume /rcrds

Namespace: medarcv
    Volume: /rcrds
        Path: /sprains
            Options:            recurse rename-files
            Status:             Success
            Report Name:        sync.2._rcrds.rpt (bstnA)
            Lost Items:         0
            Found Items:        0
            Synchronized Items: 0
            Processed Items:    11

        Path: /
            Options:            recurse rename-files
            Status:             Success
            Report Name:        sync.3._rcrds.rpt (bstnA)
            Lost Items:         1
            Found Items:        0
            Synchronized Items: 1
            Processed Items:    245
```

sync directories

Purpose If a directory appears on a back-end filer without the managed volume's awareness, the volume metadata is out-of-sync for that directory. This may occur for new directories that cannot be created from a client interface, such as NetApp qtrees. To use one of these special directories in a managed volume, you can create it on the back-end filer itself and then use the sync directories command to find it and add it to the volume's metadata.

Modes priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `sync directories` ***namespace*** `volume` ***volume*** `path` ***path*** `[share` ***share-name***`]`

namespace (1-30 characters) specifies the namespace to sync. Use the show namespace command for a list of configured namespaces.

volume (1-1024 characters) specifies the namespace volume. All of the namespace's volumes appear in the show namespace output.

path (1-1024 characters) chooses a specific directory in the volume (for example, "var/log"). The sync operation looks for new directories under this path; it does not descend into any already-known subdirectories. This path is relative to the ***volume*** root, above. Use forward slashes (/) for NFS or CIFS paths.

share ***share-name*** (optional, 1-64 characters) limits the sync and repair to a single share in the volume.

Default(s) all shares in the volume are synchronized and repaired.

Guidelines In general, no changes should occur on a filer share that is imported into a managed volume, except changes that go through that volume. However, some directories require features that can only be applied directly on the back-end filer, such as NetApp qtrees. Use this command after one or more empty directories were added behind a managed volume.

As mentioned above, only empty directories can be found with the sync utility.

To find all new directories where they exist, use nsck ... report inconsistencies. This command, sync directories, finds all new subdirectories under a known directory.

Very large directories may take a long time to synchronize. The command returns immediately, providing you with the name of a detailed report. You can continue using the CLI while the synchronization occurs in the background. The show sync command shows the current status of one or more sync operations. You can use the wait-for sync command to wait for the operation to complete. To cancel the operation, use cancel sync.

Each sync operation generates a report named "sync.*sync-job-id*.*volume*.rpt," shown after you issue the command. Use show reports to list all reports, including this sync report. To follow the progress of the sync operation, you can use tail reports *report-name* follow.

If the sync/repair fails, use the nsck ... rebuild command to rebuild the volume (drop all volume metadata and re-import all back-end shares).

Guidelines: Directory Mastership After the Sync

The sync directories command uses the share priority set by the import priority command. That is, if two or more shares have the synchronized directory, and no other instance of that directory is already imported, the import priority determines which share has the master instance of the directory. Whenever there are two or more instances of a volume directory on the volume's back-end shares, one of them is *master* and the rest are called *stripes*. The stripe directory conforms to the attributes of the master.

Sample

```
bstnA# sync directories medarcv volume /lab_equipment path /acme share
leased
Scheduling sync directories operation on switch bstnA, report name:
sync.4._lab_equipment.rpt
bstnA#
```

finds any new subdirectories under a particular directory and on a particular share: the /acme directory in the "medarcv~/lab_equipment~leased" share. Figure 37.15 shows a sample report, where one new directory, "/filters," was found and synchronized.

Related Commands

show sync
wait-for sync
cancel sync
nsck ... report inconsistencies
show namespace
show reports
nsck ... rebuild

Figure 37.15 *Sample Report: sync.4._lab_equipment.rpt*

```
bstnA# show reports sync.4._lab_equipment.rpt
**** Sync Report: Started at Tue Apr 27 01:27:33 2010 ****
**** Software Version: 5.02.000.12561 (Apr 23 2010 20:15:19) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

**** Namespace: medarcv
**** Volume:    /lab_equipment
**** Path:      /acme
**** Share:     leased

Share                Physical Filer
-------------------  ------------------------------------------------
[leased           ]  192.168.25.49:for_lease

**** Legend:
****
****  Actions:
****   SY = File metadata has been synchronized.
****   LF = Lost File, File exists in the metadata, but not on filer.
****   FF = Found File, File exists on filer, but not in the metadata.
****   LL = Lost Symlink, File is a symlink in the metadata, but not on filer.
****   FL = Found Symlink, File is a symlink on the filer, but not in the metadata.
****   IF = Invalid File Handle, Metadata/Filer file handle mismatch.
****   CR = File collided with existing filename, and was renamed.
****   NC = File or directory is now accessible to both NFS and CIFS.
****   SD = Striped directory found on filer; registered in metadata.
****   SC = Directory stripe inconsistency has been cleared.
****   SI = Directory stripe is now inconsistent due to attempted synchronization.
****
```

```
****  Warnings:
****   CC = NFS case-blind name collision, FGN may result under CIFS.
****   IC = Name contains invalid characters for CIFS, FGN will result under CIFS.
****   NI = Name does not exist via both NFS and CIFS.
****   UN = Name contains unmappable unicode characters, FGN will result under NFS.
****   U8 = Name contains invalid UTF-8 characters, not imported.
****   CD = CIFS case blind name collision, directory ignored.
****   RV = Reserved name encountered on filer, not imported.
****   NF = Name is only accessible using NFS.
****   SL = File is a symlink.
****
****  Errors:
****   TO = Operation timed out: An operation to the filer timed out.
****   LK = Failed to acquire Lock.
****   FE = Encountered filer error during sync.
****   CF = File collided with existing filename.
****   CN = CIFS case blind name collision, file renamed.
****   CZ = CIFS case blind name collision, no rename, use rename-files.
****   FD = Found Directory, Directory exists on filer, but not in the metadata.
****   NE = Directory not empty


Synchronization Results:
================================================================================

Type                 Share               Object
-------------------  ------------------  ------------------------------------
[SY    FD         ]  [leased          ]  /acme/filters

================================================================================

Total Found Items:                              1
Total Lost Items:                               0
Total Invalid Filehandles:                      0
File Name Collisions:                           0
File Collision Renames:                         0
Total Migrations Aborted:                       0
Total Items Synchronized:                       1
Directories completely processed:               1
Total Directories:                              1
CIFS Case Blind Collisions:                     0
Non-Empty Directories:                          0

**** Total processed:              0
**** Elapsed time:           00:00:00
**** Sync Report: DONE at Tue Apr 27 01:27:33 2010 ****
```

sync files

Purpose If a file changes on a back-end filer without the managed volume's awareness, the volume metadata is *inconsistent* for that file. Use the sync files command to report on metadata inconsistencies and repair them.

Modes priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`sync files`** ***`namespace`*** **`volume`** ***`volume`*** **`path`** ***`path`*** **`[share`** ***`share-name`*** **`] [recurse] [rename-files]`**

namespace (1-30 characters) specifies the namespace to sync. Use the show namespace command for a list of configured namespaces.

volume (1-1024 characters) specifies the managed volume. All of the namespace's volumes appear in the show namespace output.

path (1-1024 characters) narrows the scope of the repair to a specific path in the volume (for example, "var/log"). This path is relative to the ***volume*** root, above. Use forward slashes (/) for NFS or CIFS paths.

share ***share-name*** (optional, 1-64 characters) limits the sync and repair to a single share in the volume.

You can enter the final two options in any order, or you can omit them:

recurse causes the sync operation to descend into subdirectories.

rename-files determines that, if a file-collision is discovered, the sync operation can rename the newly-discovered file. All collisions and renames are included in the sync report, described below.

Default(s) **recurse** is off; the sync only examines the ***dir-path*** without descending into any child directories.

rename-files is off; collisions are noted in the sync report, and the metadata is not repaired.

all shares in the managed volume are synchronized and repaired.

Guidelines In general, no changes should occur on a filer share that is imported into a managed volume, except changes that go through that volume. However, some situations make this impossible:

- On reboot recovery, some filers replace files that were deleted shortly before the reboot.
- A filer-integrated Anti-Virus application may rename, remove, or create files on the filer itself.
- Some customers prefer to perform restore operations directly on their back-end filers.

Guidelines (Cont.)

Whenever a client tries to access a file that has been removed by one of the above operations, the client system reports an internal error. For CIFS operations, messages appear in the syslog (labeled UNXFINDNONE, UNXNOPATH, UNXOPENTYPE and/or UNXCRETYPE.), and the volume software sends an SNMP trap (vcifsNotfound or vcifsTypeMismatch). Use grep to search the syslog for a string. If a CIFS-only volume is configured to launch synchronizations automatically (with auto sync files), the manual sync files command is unnecessary.

Use this command after any of the above operations (filer reboot, filer restore, or virus detection at the filer). You can also use this in response to client complaints. To confirm any inconsistencies and find all directories where they exist, use nsck ... report inconsistencies.

Very large directories may take a long time to synchronize. The command returns immediately, providing you with the name of a detailed report. You can continue using the CLI while the synchronization occurs in the background. The show sync command shows the current status of one or more sync operations. You can use the wait-for sync command to wait for the operation to complete. To cancel the operation, use cancel sync.

Each sync operation generates a report named "sync.*sync-job-id*.*volume*.rpt," shown after you issue the command. Use show reports to list all reports, including this sync report. To follow the progress of the sync operation, you can use tail reports *report-name* follow.

If the sync/repair fails, use the nsck ... rebuild command to rebuild the managed volume (drop all volume metadata and re-import all back-end shares).

Guidelines: Discovering Symlinks

This command also probes for NFS symbolic links (or symlinks) in a multi-protocol namespace. NFS symlinks are recorded in a multi-protocol volume's metadata so that the volume can support them for its CIFS clients. (You can disable this CIFS-side de referencing with the cifs deny-symlinks command, if desired).

Symlink discovery is unnecessary for most installations because symlinks are correctly flagged during import, and whenever an NFS client creates a new symlink. It is only necessary after a software upgrade adds CIFS symlink support to an existing multi-protocol volume, or if an NFS client created symlinks directly on a back-end filer.

Guidelines: Directory-Stripe Synchronization

Before the sync operation can sync any file, it must first verify that the file's parent directory has synchronized attributes (permissions, ACLs, ownership, and so on). The parent directory's attributes must be identical for all instances of the directory, in all of the shares behind the managed volume. If any of the managed volume's shares are nearly out of space, this attribute-sync process may fail; this is because CIFS ACLs may be part of the directory attributes, and they require disk space.

This directory-attribute sync may result in SC or SI flags in the report:

- SC indicates that a previously unsynchronized directory was successfully synchronized by this command. The managed-volume software does this in the background, so it is rare to see this in the report for a manual sync operation.
- SI indicates that the sync operation attempted to sync a directory's attributes, and the directory-attribute synchronization failed. This indicates that one of the managed volume's shares is out of space. You can use show namespace to find the full share. You can access the filer directly to add more space, or use a place-rule with a source share to drain some files off of the full share. Then the managed-volume software synchronizes the directory attributes in the background, and this condition clears.

Guidelines: File Mastership After the Sync

The sync files command uses the share priority set by the import priority command. That is, if two or more shares have the same file at the same path, and no other instance of that file is already imported, the share with the better import priority wins the master instance of the file. The volume renames the other file(s) or fails the sync operation, as determined by the **rename-files** flag.

This is only relevant if you synchronize multiple shares. If you synchronize a single share, its files can only collide with already-imported files. An already-imported file is always master over a newly-imported file.

Samples

```
bstnA# sync files wwmed volume /acct path /
Scheduling sync files operation on switch bstnA, report name:
sync.1._acct.rpt
bstnA#
```

repairs the top-level directory in the "wwmed~/acct" volume. This does not descend into child directories, nor does it repair any renamed files that have collided with known files.

```
bstnA# sync files wwmed volume /acct path /jsmith recurse
Scheduling sync files operation on switch bstnA, report name:
sync.2._acct.rpt
bstnA#
```

recursively repairs the entire /jsmith directory in the "wwmed~/acct" volume.

```
bstnA# sync files medarcv volume /rcrds path / recurse rename-files
Scheduling sync files operation on switch bstnA, report name:
sync.3._rcrds.rpt
```

recursively repairs the root directory (/) in a CIFS volume, "medarcv~/rcrds." Figure 37.16 on page 37-48 shows a sample report, where a missing file, "/copyRandomx.exe," was found and synchronized. Two additional files were also newly-discovered and synchronized.

Related Commands

show sync
wait-for sync
cancel sync
nsck ... report inconsistencies
show namespace
show reports
nsck ... rebuild

Figure 37.16 *Sample Report: sync.3.rcrds.rpt*

```
bstnA# show reports sync.3._rcrds.rpt
**** Sync Report: Started at Tue Apr 27 01:27:33 2010 ****
**** Software Version: 5.02.000.12561 (Apr 23 2010 20:15:19) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

**** Namespace: medarcv
**** Volume:    /rcrds
**** Path:      /
**** Options:   recurse rename-files

Share                Physical Filer
-------------------  ------------------------------------------------
[rx               ]  192.168.25.29:prescriptions
```

```
[charts            ]  192.168.25.20:histories
[bulk              ]  192.168.25.27:bulkstorage

**** Legend:
****
****  Actions:
****   SY = File metadata has been synchronized.
****   LF = Lost File, File exists in the metadata, but not on filer.
****   FF = Found File, File exists on filer, but not in the metadata.
****   LL = Lost Symlink, File is a symlink in the metadata, but not on filer.
****   FL = Found Symlink, File is a symlink on the filer, but not in the metadata.
****   IF = Invalid File Handle, Metadata/Filer file handle mismatch.
****   CR = File collided with existing filename, and was renamed.
****   NC = File or directory is now accessible to both NFS and CIFS.
****   SD = Striped directory found on filer; registered in metadata.
****   SC = Directory stripe inconsistency has been cleared.
****   SI = Directory stripe is now inconsistent due to attempted synchronization.
****
****  Warnings:
****   CC = NFS case-blind name collision, FGN may result under CIFS.
****   IC = Name contains invalid characters for CIFS, FGN will result under CIFS.
****   NI = Name does not exist via both NFS and CIFS.
****   UN = Name contains unmappable unicode characters, FGN will result under NFS.
****   U8 = Name contains invalid UTF-8 characters, not imported.
****   CD = CIFS case blind name collision, directory ignored.
****   RV = Reserved name encountered on filer, not imported.
****   NF = Name is only accessible using NFS.
****   SL = File is a symlink.
****
****  Errors:
****   TO = Operation timed out: An operation to the filer timed out.
****   LK = Failed to acquire Lock.
****   FE = Encountered filer error during sync.
****   CF = File collided with existing filename.
****   CN = CIFS case blind name collision, file renamed.
****   CZ = CIFS case blind name collision, no rename, use rename-files.
****   FD = Found Directory, Directory exists on filer, but not in the metadata.
****   NE = Directory not empty


Synchronization Results:
================================================================================

Type                 Share               Object
-------------------  ------------------  ------------------------------------
[SY LF             ]  [charts          ]  /copyRandomx.exe
[SY    FF          ]  [charts          ]  /RECORD~1/test.dat
[SY    FF          ]  [charts          ]  /RECORD~1/test2.dat

================================================================================

Total Found Items:                          2
Total Lost Items:                           1
Total Invalid Filehandles:                  0
File Name Collisions:                       0
File Collision Renames:                     0
Total Migrations Aborted:                   0
Total Items Synchronized:                   3
Directories completely processed:          32
Total Directories:                         32
CIFS Case Blind Collision Renames:          0
```

```
**** Total processed:             279
**** Elapsed time:           00:00:00
**** Sync Report: DONE at Tue Apr 27 01:27:33 2010 ****
```

sync subshares from-namespace

Purpose A *subshare* is a CIFS share (along with its ACL) that is nested under an imported CIFS share. Use the sync subshares from namespace command to discover all of the CIFS subshares behind a managed volume and then share them from a front-end service. This synchronizes the CIFS subshares between back-end filers and the front-end service.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax

```
sync subshares from-namespace ns volume vol to-service fqdn
    [expose-hidden] [tentative]
```

ns (1-30 characters) is the namespace for the sync operation. The namespace must support CIFS.

vol (1-1024 characters) specifies a managed volume with subshares. This command finds all CIFS subshares in all of the back-end shares behind this volume. It also replicates any subshares that are not already identical on those back-end shares. The command replicates the subshare name, directory path (relative to the share root), and ACL on each back-end filer that needs it.

fqdn (1-255 characters) is the target CIFS service to export all of the subshares found in the above volume.

expose-hidden (optional) causes the command to use a different front-end name for any subshares that are "hidden" on their back-end filers. A CIFS share is hidden from Windows-network views by adding a dollar-sign ($) to the end of the share name (for example, "mybooks$"). This option causes the operation to use an "exposed" front-end subshare for all such hidden subshares, without the dollar suffix ("mybooks"). Clients of the CIFS service can therefore see the shares.

This option does not expose any hidden subshares that may have been generated by the ARX, such as "_acopia_mybooks_4$."

tentative (optional) finds all of the subshares behind the volume and prints them in a report, but does not share them from the CIFS service.

Default(s) None.

Guidelines

This operation is only effective in volumes that support CIFS and filer subshares. Subshares require special processing and configuration in order to be supported by the ARX. By default, if a CIFS client connects to a subshare on the ARX, the managed volume connects to the root of the filer share and then changes to the subshare directory from there. Since the volume connects to the root share, the ACL at that share is used to determine client-access privileges. This bypasses any share-level ACL that may be defined for the filer subshare.

You can configure a managed volume so that it passes a client connection from a front-end subshare directly to the back-end subshare, thereby allowing the filer to enforce its subshare ACLs. Use the gbl-ns-vol filer-subshares command to prepare a CIFS volume for subshare support, and then use this command to share all filer subshares out through a front-end CIFS service. (You can also use the export (gbl-cifs) ... filer-subshare command to share out a single subshare at a time.)

This command replicates all subshares, so that they have the same share names, ACLs, and directory names on all of the volume's shares. The *subshare-replication* process creates a sync report with the results. The name of the report appears after you issue the command. Use show reports *report-name* to see it.

To add a new subshare to a running CIFS service, follow the procedure outlined in the *Guidelines* of the filer-subshares command.

Guidelines: Preparation for NetApp/EMCs Where You Want Free-Space Quotas

If the volume is backed by any NetApp filers or EMC servers and you plan to support free-space quotas on them (see freespace cifs-quota), prepare the subshares directly on the filer *before* you invoke this command. The subshare-replication process does not create NetApp qtrees or EMC quota trees, and those special directories are required to support free-space quotas. Before you run subshare replication with this command, access the filers directly and

- (on a NetApp) create one qtree per subshare, or
- (on an EMC) ensure that the imported share is a File System with one quota tree per subshare.

See the *Guidelines: Subshare Replication with Free-Space Quotas* section of the filer-subshares documentation. Then invoke this command to synchronize their share-level ACLs and other attributes.

Guidelines: Naming for Replicated Subshares

This operation replicates subshares between filers as needed. Subshare replication also occurs during share import, so it is not typically necessary for this command to repeat it. For cases where it is necessary, the volume replicates every subshare name, directory path (relative to the import-share root) and ACL. Replicated subshares use the same name as the source subshare, called the *native* subshare name, whenever possible.

The volume cannot use a native subshare name if any other CIFS share on the back-end filer is already using that name. This is called a subshare-name *collision*. For subshare-name collisions, the managed volume names the subshare with the following syntax:

_acopia_*subshare_id*$

where

- *subshare* is the native name of the original subshare.
- *id* is unique integer to identify the subshare. This is unique for every subshare behind the ARX and its redundant peer.
- $, at the end of the share name, hides the subshare name from most network clients.

To prevent any such generated names on a volume's filers, you can use the native-names-only option in the filer-subshares command. The native-names-only option causes subshare replication to fail.

Guidelines: Repairing a Failed Subshare Replication

If this replication process fails for any subshare, the front-end subshare is "degraded:" clients cannot access the files in the unsynchronized back-end subshare(s). The show cifs-service all command shows the current status of the subshare-replication process, and whether or not any subshare is degraded. You can use this command's sync report to understand the problem, resolve the problem at the filer, and then use sync subshares from-service to re-run the subshare replication.

Sample

```
bstnA# sync subshares from-namespace medarcv volume /rcrds to-service
ac1.medarch.org expose-hidden

% INFO: Sync subshares operation completed. See report
'syncSshrVolToService_20100706010239.rpt' for more detail.
```

replicates all subshares in the "medarcv~/rcrds" volume, then shares them out to clients through the "ac1.medarch.org" service. Figure 37.17 on page 37-53 shows a sample report.

Related Commands

filer-subshares
export (gbl-cifs) **... filer-subshare**
show namespace
show reports
sync subshares from-service

Figure 37.17 *Sample Report: syncSshrVolToService_....rpt*

```
bstnA# show reports syncSshrVolToService_20100706010239.rpt
**** Synchronize Subshares from Volume to Service Report: Started at Tue Jul  6 01:02:39 2010 ****
**** Software Version: 5.02.000.12591 (Jul  3 2010 13:41:14) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

Operation Parameters
```

```
--------------------------------------------------------------------------------
Service: ac1.MEDARCH.ORG
Namespace: medarcv
Volume: /rcrds (modifiable, no access-based-enum)
Options:
    None.


ABE - access-based-enumeration.
ACL - security descriptor access control list.

Operation Execution Detail
--------------------------------------------------------------------------------

Action: Exported filer subshare.
    Share Name: CELEBS$
    Relative Path: /rcrds/VIP_wing
    Remark:
    Exported as: CELEBS

Action: Exported filer subshare.
    Share Name: MP3S
    Relative Path: /rcrds/2011/mp3downloads
    Remark:
    Exported as: MP3S

Action: Exported filer subshare.
    Share Name: Y2004
    Relative Path: /rcrds/2004
    Remark:
    Exported as: Y2004

Action: Exported filer subshare.
    Share Name: Y2010
    Relative Path: /rcrds/2010
    Remark:
    Exported as: Y2010

Information: Existing Export for filer subshare.
    Share Name: Y2005
    Relative Path: /rcrds/2005
    Remark:
    Exported as: Y2005

Summary:
--------------------------------------------------------------------------------

4 exports were created.
1 exports already exist.


**** Total processed:               5
**** Elapsed time:           00:00:03
**** Sync Subshares from Volume to Service Report: DONE at Tue Jul  6 01:02:42 2010 ****
```

sync subshares from-service

Purpose

A *subshare* is a CIFS share (along with its ACL) that is nested under an imported CIFS share. Use the sync subshares from-service command to send CIFS subshares from a front-end CIFS service to all of the filers behind it. This is useful for repairing a degraded subshare.

Modes

priv-exec

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
sync subshares from-service fqdn to-namespace ns volume vol
    [tentative]
```

fqdn (1-255 characters) is the FQDN for a CIFS service with subshares. Use the show cifs-service subshares all command for a complete list of CIFS services with subshares.

ns (1-30 characters) is the namespace behind the above service.

vol (1-1024 characters) is a particular managed volume with subshares. If the CIFS service exports any subshares from this volume, this command replicates the subshare definitions at all filers behind the volume.

tentative (optional) prints a report about all of the changes the command would make, but does not replicate any subshare information from filer to filer.

Default(s)

None.

Guidelines

This operation is only effective in volumes that support CIFS and filer subshares. When you configure subshare support at the volume (with filer-subshares) and the volume's front-end export (with export (gbl-cifs) ... filer-subshare), a client that connects to the front-end subshare passes through the volume to the back-end subshare, so that the back-end filer can impose its subshare ACL. Without subshare support, a client that connects to a front-end subshare is passed through to the root of the back-end share, and the filer uses the root share's ACL.

The ACL must be identical on all back-end instances of a subshare, so that clients have the same access rules no matter which back-end subshare they need.

You can use this command, sync subshares from-service, to correct a possible problem with subshares. When you first enable a managed volume with subshare support, it discovers and replicates all of the subshares on its back-end shares. After you export a subshare through a front-end service at some *fqdn*, the show cifs-service *fqdn* command indicates whether or not this subshare-replication process succeeded. If the subshare is degraded, you can get a detailed report about the problem by running this sync command with its tentative flag. Then you can correct the problem at the back-end filers and run this sync command again, this time without the tentative flag, to retry the replication.

The command creates a sync report with the results. The name of the report appears after you issue the command. Use show reports *report-name* to see it.

To export subshares that are currently defined at the back-end filers, use the sync subshares from-namespace command. This is the reverse of the current command; it sends subshare definitions from the back-end filers to the front-end service.

Guidelines: Preparation for NetApp/EMCs Where You Want Free-Space Quotas

If the volume is backed by any NetApp filers or EMC servers and you plan to support free-space quotas on them (see freespace cifs-quota), prepare the subshares directly on the filer *before* you invoke this command. The subshare-replication process does not create NetApp qtrees or EMC quota trees, and those special directories are required to support free-space quotas. Before you run subshare replication with this command, access the filers directly and

- (on a NetApp) create one qtree per subshare, or
- (on an EMC) ensure that the imported share is a File System with one quota tree per subshare.

See the *Guidelines: Subshare Replication with Free-Space Quotas* section of the filer-subshares documentation. Then invoke this command to synchronize their share-level ACLs and other attributes.

Guidelines: Naming for Replicated Subshares

This operation replicates subshares between filers as needed. The volume replicates every subshare name, directory path (relative to the import-share root) and ACL. Replicated subshares use the same name as the source subshare, called the *native* subshare name, whenever possible.

The volume cannot use a native subshare name if any other CIFS share on the target filer is already using that name. This is called a subshare-name *collision*. For subshare-name collisions, the managed volume names the subshare with the following syntax:

_acopia_*subshare_id*$

where

- *subshare* is the native name of the original subshare.
- *id* is unique integer to identify the subshare. This is unique for every subshare behind the ARX and its redundant peer.
- $, at the end of the share name, hides the subshare name from most network clients.

To prevent any such generated names on a volume's filers, you can use the native-names-only option in the filer-subshares command.

Sample

```
newptA# sync subshares from-service provmed.MEDARCH.ORG to-namespace
provMed volume /mds tentative

% INFO: Sync subshares operation completed. See report
'syncSshrServToVolume_201110030248.rpt' for more detail.
```

tentatively syncs the subshares in the "provmed.MEDARCH.ORG" service. This produces a report that lists the subshares that require synchronization: see the sample in Figure 37.18 on page 37-57. The sample report shows missing subshares.

Related Commands

filer-subshares
show cifs-service
sync subshares from-namespace
export (gbl-cifs) **... filer-subshare**
show namespace
show reports

Figure 37.18 *Sample Report: syncSshrServToVolume_....rpt*

```
newptA# show reports syncSshrServToVolume_201110030248.rpt
**** Synchronize Subshares from Service to Volume Report: Started at 10/03/2011 02:48:51 -0400
****
**** Software Version: 6.02.000.14265 (Sep 28 2011 20:09:46) [nbuilds]
**** Hardware Platform: ARX-500
**** Report Destination:

Operation Parameters
--------------------------------------------------------------------------------
Service: provmed.MEDARCH.ORG
Namespace: provMed
Volume: /mds (modifiable, no access-based-enum)
Options:
    Tentative. Report, but do not perform the operation.


ABE - access-based-enumeration.
ACL - security descriptor access control list.

Operation Execution Detail
--------------------------------------------------------------------------------

Problem: Share missing.
    Share Name: BLAKE
    Relative Path: blake
    Import Share: MDS
    Filer Name: ntap-npt2
    Filer IP: 192.168.8.98
    Status: Synchronize was not attempted.

Problem: Share missing.
    Share Name: BLAKE
    Relative Path: blake
    Import Share: SURGEONS
    Filer Name: ntap-nwpt
    Filer IP: 192.168.8.90
    Status: Synchronize was not attempted.

Problem: Share missing.
    Share Name: HOGTROFF
    Relative Path: hogtroff
    Import Share: MDS
    Filer Name: ntap-npt2
    Filer IP: 192.168.8.98
    Status: Synchronize was not attempted.

Problem: Share missing.
    Share Name: HOGTROFF
    Relative Path: hogtroff
    Import Share: SURGEONS
    Filer Name: ntap-nwpt
    Filer IP: 192.168.8.90
    Status: Synchronize was not attempted.

Problem: Share missing.
    Share Name: MCCOY
    Relative Path: mccoy
    Import Share: MDS
    Filer Name: ntap-npt2
    Filer IP: 192.168.8.98
    Status: Synchronize was not attempted.
```

```
Problem: Share missing.
    Share Name: MCCOY
    Relative Path: mccoy
    Import Share: SURGEONS
    Filer Name: ntap-nwpt
    Filer IP: 192.168.8.90
    Status: Synchronize was not attempted.

Problem: Share missing.
    Share Name: MHOWARD
    Relative Path: mhoward
    Import Share: MDS
    Filer Name: ntap-npt2
    Filer IP: 192.168.8.98
    Status: Synchronize was not attempted.

Problem: Share missing.
    Share Name: MHOWARD
    Relative Path: mhoward
    Import Share: SURGEONS
    Filer Name: ntap-nwpt
    Filer IP: 192.168.8.90
    Status: Synchronize was not attempted.

Problem: Share missing.
    Share Name: PIERCE
    Relative Path: pierce
    Import Share: MDS
    Filer Name: ntap-npt2
    Filer IP: 192.168.8.98
    Status: Synchronize was not attempted.

Problem: Share missing.
    Share Name: PIERCE
    Relative Path: pierce
    Import Share: SURGEONS
    Filer Name: ntap-nwpt
    Filer IP: 192.168.8.90
    Status: Synchronize was not attempted.

Summary:
--------------------------------------------------------------------------------
No share synchronize operations were executed.
10 shares would have been added.


**** Total processed:              10
**** Elapsed time:           00:00:00
**** Sync Subshares from Service to Volume Report: DONE at 10/03/2011 02:48:51 -0400 ****
```

wait-for nsck

Purpose Use the wait-for nsck command to wait for an nsck operation to complete.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`wait-for nsck`** ***`namespace`*** **`[`*`vol-path`*`] [timeout`** ***`timeout`*****`]`**

namespace (1-30 characters) is the name of a namespace.

vol-path (optional, 1-4096 characters) identifies a particular managed volume in the namespace.

timeout (optional, 1-2096) is the timeout value, specified in seconds. If this timer expires before the nsck operation completes, the command exits with a warning.

Default(s) ***timeout*** - 0 (zero, meaning "wait indefinitely")

Guidelines The wait-for nsck command waits for an nsck operation on a specified namespace or volume to complete.

To interrupt the wait-for nsck command before the nsck operation completes, press **<Ctrl-C>**.

If the command is waiting for nsck ... rebuild, it exits after the volume(s) are destaged and have started importing files and directories from back-end shares. This command only waits until clients have full access to their files, not until import is complete.

Sample

```
bstnA(gbl-ns[medarcv])# wait-for nsck medarcv timeout 60
```

waits for a nsck operation to complete in the "medarcv" namespace. If 60 seconds elapse before the nsck operation is finished, the command exits.

Related Commands

clear nsck
nsck ... report dir-structure
nsck ... report inconsistencies
nsck ... report metadata-only
nsck ... rebuild
show nsck
show reports
terminal beta

wait-for sync

Purpose Use the wait-for sync command to block the CLI until one or more sync operations completes.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
wait-for sync {files | directories} [timeout timeout]
wait-for sync {files | directories} namespace [timeout timeout]
wait-for sync {files | directories} namespace volume vol-path
[timeout timeout]
wait-for sync {files | directories} namespace volume vol-path
path path [timeout timeout]
```

files | directories is a required choice. This selects the type of sync operation to await.

The ***namespace***, ***vol-path***, and ***path*** arguments mimic the ones used in the sync files or sync directories command. These identify one or more sync operations; if you omit all of them, the command waits for all sync-files or sync-directories operations on the switch.

namespace (1-30 characters) specifies a namespace where one or more sync operations are occurring.

vol-path (optional, 1-1024 characters) specifies a volume. Use forward slashes (/) for all volumes, including CIFS.

path (optional, 1-1024 characters) identifies a particular path that is being synchronized. As with volumes, use forward slashes in all paths.

timeout (optional, 1-2096) is the timeout value, specified in seconds. If this timer expires before the sync operation completes, the command exits with a warning.

Default(s) ***timeout*** - 0 (zero, meaning "wait indefinitely")

Guidelines The wait-for sync command waits for a specified sync operation to complete.

To interrupt the wait-for sync command, press **<Ctrl-C>**. Use cancel sync to cancel the sync operation altogether. The show sync command shows all sync-files or sync-directories operations.

Samples

```
bstnA# wait-for sync files timeout 60
```

waits for all sync-files operations on the switch to complete. If they do not complete in 60 seconds, the command returns with a warning.

```
bstnA(gbl-ns[medarcv])# wait-for sync directories medarcv timeout 300
```

waits for all sync-directories operations in the "medarcv" namespace to complete. If they do not complete in 300 seconds, the command returns with a warning.

```
bstnA> wait-for sync files wwmed timeout 120
```

waits for all sync-files operations in the "wwmed" namespace to complete. If they do not complete in 120 seconds, the command returns with a warning.

Related Commands

sync files
sync directories
show sync
cancel sync

38

Basic Troubleshooting Tools

This chapter contains reference information for using and tuning the System Logging (syslog) utility and other troubleshooting tools.

Log Components

A *log component* is a source of syslog messages, typically an internal process or group of processes. The table below is an alphabetical list of all log components, with a brief description and an indication of the board(s) where the software runs.

Log Component	Description
AFFINITY	Affimgrd and libaffin, which manage the assignment of internal processes to CPU cores.
AFN_NET	IPC Abstraction Layer, for communication between internal ARX processes.
API	API Services.
APP_RON_RTMD	RON Route Table Maintenance Daemon.
AT	AT Library. This is deprecated, so it is unlikely to yield any log messages.
AUTH_CONSOLE	Console Authentication. This component performs authentication for administrators who log into the Console (serial) port.
AUTH_HTTP	HTTP Authentication. This component performs authentication for administrators who access the GUI through HTTP.
AUTH_HTTPS	HTTPS Authentication. This component performs authentication for administrators who access the GUI through HTTP over SSL, or HTTPS.
AUTH_SSH	SSH Authentication. This component performs authentication for administrators who access the CLI through Secure SHell (SSH).
AUTH_TELNET	Telnet Authentication. This component performs authentication for administrators who access the CLI through Telnet.
AUTO_DIAG_CLI	Auto Diagnostics, which gathers input/output statistics and monitors for service degradation.
BESPD	Back-end Share Probing daemon, which periodically probes the shares behind ARX volumes and checks their health.
BOOTPD	BOOTP Daemon, or the boot-protocol server, which manages ARX boots and reboots.
CATALOG_INTERNAL	Message Catalog Internal Errors.
CHASS	Chassis Manager, which implements chassis operations such as 'reload,' 'firmware upgrade,' and hardware monitoring.
CHASS_LIB	Chassis Manager Library. The chassis manager implements chassis operations such as 'reload,' 'firmware upgrade,' and hardware monitoring.
CIFS_CLI	CIFS CLI.

Log Component	Description
CIFS_SHIM_LIB	Interface to metadata storage on CIFS filers. Metadata is data about managed-volume files, such as their locations on back-end filers.
CIFSCONFD	CIFS Configuration Daemon, which manages one front-end CIFS service.
CLI	Command Line Interface.
CLI_NFS_COPY	Copy to/from NFS namespace.
CLI_RUN_EXIM	Run exim4 from the CLI..
CLI_WIZARD	Initial Interview, which prompts with questions and performs initial configuration the first time the ARX boots.
CLUSTER	Not supported.
COMMON_DNAS	Common DNAS Messages.
COMMON_NIS	Common Network Information System (NIS) messages. NIS servers provide a lookup service for NFS servers, to manage large NFS access lists.
COMMONLIB	Not supported.
CORE	Core collector.
CRMD	Critical Resource Monitor Daemon, which makes HA failover decisions based on external filers, routers, etc.
DDBD	Distributed Database Library. This is for the internal database, which stores the ARX configuration. You can see the full configuration with 'show running-config' and 'show global-config'.
DEBUG	Internal debugging - used for development only.
DFPD	DNAS Fast Path Manager Daemon, which configures managed volumes in the NSM data plane.
DME	Policy Data Mover Engine (DME). This implements file migrations from one back-end filer to another.
DMOUNTD	Distributed Mount Daemon for NFS mounts.
DMSG_LIB	DNAS Message Library.
DNAS_IMPORT	DNAS Import Messages. An import occurs when a filer share is enabled in an enabled managed volume. This is the process of walking the directory tree of the share and recording all of its files and directories in managed-volume metadata.
DNAS_MD_MIGRATE	DNAS Metadata Migration, which moves managed-volume metadata from one external filer to another.
DNAS_PGRP	DNAS Process Group and NFS. DNAS, or Distributed NAS, is the internal name for namespace software.

Log Component	Description
DNAS_POLICY	Inline policy events that occur during normal volume processing, such as file creates, directory renames, and other client actions. This is off by default because it generates a great deal of log messages, enough to potentially affect client traffic.
DNAS_SINIT	DNAS Storage Initialization.
DNAS_SREMOVE	DNAS Storage Removal.
DNC_LIB	DNAS Configuration Library for NFS. DNAS is the internal name for namespace software.
DOMAIN_JOIN	Active Directory domain joining components, which join an ARX-CIFS service to a Windows domain.
DR	Disaster Recovery, which involves loading a configuration on a backup-ARX site and activating the configuration on command.
DT	Direct Transport processes. This is a high-speed inter-process communication (IPC) library.
EMAILHOME	Email home daemon and CLI commands.
ENVLIB	Environmental Library, for monitoring the hardware platform. This only applies to the ARX-500.
EVAL_NET_RESULT	The results of joining an Active Directory (AD) domain.
EXIMD	Exim Daemon, which sends RPC requests to the exim message-transfer agent. That agent sends SMTP (email) messages from the ARX to external devices.
FCN	File Change Notification, for API access to snapshotted configuration and metadata.
FCND	File Change Notification daemon which provides API access to snapshotted configuration and metadata.
FFP	Fast Filter Processor.
FFP_UTIL	Fast Filter Processor Library.
FFPD	Fast Filter Processor Daemon.
FILER_MGMT_LIB	Snapshot daemon filer operations library.
FREESPACEMONITOR	Free Space Monitor.
FT_CIFS_SHIM_LIB	Interface to metadata storage on CIFS filers. Metadata is information about files and directories in managed volumes, such as their locations on back-end filers.
FT_CLI	File Tracking CLI. These are the CLI operations for setting up a file-history archive and using it to track the location of managed-volume files over time.

Log Component	Description
FTD	File Tracking daemon. This is the process of sending volume metadata and configuration to an external file-history archive, and using that archive to track the location of managed-volume files over time.
GBL_LIB	Global Services Library, for NFS and CIFS service management operations.
GBL_SVCMGR	Global Service Manager, which starts and stops NFS and CIFS front-end services.
GBLSVC_CLI	Global Service CLI.
GUI	Graphical User Interface (GUI), also known as the ARX Manager.
HA_CLI	High Availability CLI, such as the 'redundancy' command.
HWAGENT	Hardware Agent, for startup, shutdown, and monitoring of hardware components.
IPC_EVENT_LIB	IPC Event Library, for communication between internal ARX processes. This is deprecated, so it is unlikely to yield any log messages.
IPC_EVENT_MGR	IPC Event Manager, for communication between internal ARX processes.
IPC_LIB	Internal software inter-process messaging library.
IPCMOND	IPC service monitoring daemon, for monitoring the communication between internal ARX processes.
IPMIUTIL	ipmiutil event collector. All log levels for this logging component go to the kernel.log file, not syslog.
KRBCONFD	Active Directory (AD) Configuration Daemon, for periodically monitoring DCs and selecting the active DCs for each Windows domain.
KRBDDNS	Dynamic DNS registration helper. This registers dynamic-DNS aliases for an ARX-CIFS service with a remote DNS server.
L2SW_DRV	Layer2 - Internal Switch Hardware. This only applies to ARX-500, ARX-2000, and ARX-4000 systems.
L2SW_LVL7	Layer2 - Internal Switch Software. This controls packet traffic over network interfaces, channel traffic and LACP, spanning-tree, and other low-level networking. This only applies to ARX-500, ARX-2000, and ARX-4000 systems.
LIBCIFS	CIFS client library, used for non-client access to CIFS file servers. For example, this library is used for policy-invoked migrations.
LIBCIFSAUTH	CIFS authentication library, for authenticating front-end CIFS clients and internal ARX processes that access back-end filers.
LIBCIFSEXT	CIFS utility library, used in conjunction with LIBCIFS.

Log Component	Description
LIBEXIM	Exim Utilities messages.
LIBKDCLBSTATS	Kerberos load-balancing statistics.
LIBSDDL	CIFS Security Descriptor Definition Language translation library, which transforms binary Security IDs (SIDs) and descriptors into text strings. Callers to this library include the SID_CACHE, POLICY_SHADOW, and SUBSHRMGT processes.
LIBXSD	CIFS security-descriptor translation library, which processes can use to access the SID_CACHE daemon. The SID_CACHE daemon translates between Security IDs (SIDs) on various file servers and the group/user names in a Windows domain.
LICENSE	License monitoring daemon and CLI.
LIP_LIB	Logical IP Address Library. A logical IP address is an internal address that processes use to communicate to one another.
LOCBROKERD	IPC Location Broker for internal (process-to-process) RPC requests.
LOGMGRD	Logging Manager Daemon.
MANUF	Manufacturing Purposes only - Internal Switch Software.
MERGED	Layer2 - Merge FPGA Daemon.
MGMT_CLI	Management SLM - not supported.
MGT_DNAS_CONFIG	DNAS Configuration. DNAS, or Distributed NAS, is the internal name for namespace software.
MGT_RON_RTMD	Route Table Control Daemon. This checks gateway health and loads the best available routes into the OS route tables.
MGT_SVC_AGENT	Service Agent, which acts as a local agent for the central Service Manager (MGT_SVCMGR) process.
MGT_SVCMGR	Service Manager, used to start (and possibly restart) ARX processes.
MLBD	Metalog Block Device (MLBD) Controller, which manages the internal metalog subsystem. The metalog subsystem stores database-transaction logs for managed volumes.
MS_CLIENT	Microsoft RPC Services - High-level library used to interact with CIFS filers.
MS_RPC	Microsoft RPC Library - Low-level library used to interact with CIFS filers.
MTLDAEMON	Metalog Daemon. This is the process that records metalog (namespace-recovery log) data, on the local NVRAM hardware. This only runs on the ARX-500, ARX-2000, and ARX-4000.
NET	Network, or data-plane, processes.
NET_APP	Miscellaneous Application Events in the network software.

Log Component	Description
NET_CIFS	NSM CIFS Proxy.
NET_ERR_TRAP	Fastpath resource errors.
NET_HA	NSM HA.
NET_IPC_PROXY	NSM IPC Proxy, for communication between internal ARX processes on the control plane and on the data plane.
NET_LOG_PROXY	NSM Log Proxy.
NET_MSG	NSM Messaging between network processes.
NET_NFS	NSM NFS Proxy.
NET_OS	NSM Embedded Operating System.
NET_SLB	Server Load Balancing in the network software - not supported.
NET_STATS	NSM Statistics Collection.
NET_TCPSP	TCP Splicing.
NETBIOS	Netbios name service, which supports WINS.
NETSNMP	SNMP management on ARX-1500 and ARX-2500.
NFSCONFD	NFS Configuration Daemon, which manages one front-end NFS service.
NISD	NIS daemon messages. NIS servers provide a lookup service for NFS servers, to manage large NFS access lists. The NIS daemon queries all the configured NIS servers and caches the results of those lookups.
NLA	NETLOGON Agent, supports NTLM authentication of CIFS clients, but only for ARX CIFS services configured for constrained delegation.
NLM	NFS Network Lock Manager (NLM) Daemon. This implements the NFS NLM protocol, which manages file locks for NFS clients.
NMPIPE	Named Pipe, for communication between internal ARX processes (MS-RPC operations on CIFS IPC$ named pipes).
NSCK	DNAS Namespace Check (NSCK), for examining and repairing managed-volume metadata. Metadata is managed-volume information about its files and directories, such as locations on back-end filers.
NSCKRPT	NSCK Report Messages.
NSM	NSM Shared Memory.
NSM_RON	NSM RON.
NTLMAGENTAPI	NTLM/NTLMv2 authentication library using ARX Secure Agent.
NTP	Network Time Protocol (NTP) subsystem to keep the ARX clock(s) synchronized with an external clock on the IP network.

Log Component	Description
NVRDAEMON	Non Volatile RAM (NVRAM) Daemon, which manages one or more mtlDaemons (see MTLDAEMON). An mtlDaemon records metalog (namespace-recovery log) data, on the local NVRAM hardware. This only runs on the ARX-500, ARX-2000, and ARX-4000.
OBJ_MGR	Object Manager Database.
OBJ_MGR_LIB	Object Manager Library.
OM	Object Manager (Distributed Database) process. This is for the internal database, which stores the ARX configuration. You can see the full configuration with 'show running-config' and 'show global-config'.
OMIO	Object Manager Record I/O Tracing. This tracing is for the internal database, which stores the ARX configuration. (You can see the full ARX configuration with 'show running-config' and 'show global-config.').
OMTT	Object Manager Transaction Tracing. This tracing is for the internal database, which stores the ARX configuration. (You can see the full ARX configuration with 'show running-config' and 'show global-config.').
PERSONALITY	Personality, which parses XML data transferred between internal processes.
POLICY	not supported.
POLICY_ACTION	Migration requests and responses. You can use this to diagnose migration issues, such as file matching and mismatching.
POLICY_CLI	not supported.
POLICY_PDP	Policy Decision Point - policy scheduler, which uses schedule configuration to trigger scans and other policy events.
POLICY_PEP	Policy Enforcement Point - policy scanning, events, rule-state processing, and rule info.
POLICY_SHADOW	Policy Shadow Volume, a read-only volume that contains replicas of all the files and directories on another volume.
PROBE	Probes for Monitoring and Controlling Subsystems.
PTW	Parallel Treewalk Library.
RELMGR	Release Manager, which stores software-release (.rel) files, unpacks them, prepares them for use on the ARX, and otherwise manages them.
RELMGR_CLI	Release Manager CLI commands. This is deprecated, so it is unlikely to yield any log messages.
REMOTE_CLI	Remote SSH CLI command execution, invoked through the CLI (with the 'rconsole' command) or the GUI.
REPLMGR	Replication Manager - Directs Data-Mover Engine (DME) file-migrations.

Log Component	Description
REPORT	Common subsystem for generating and viewing reports.
RESTARTDETECT	This determines if internal processes can be restarted after a failure.
RON	Resilient Overlay Network (RON) CLI operations.
RON_SPFD	Resilient Overlay Network (RON) Routing Daemon. This manages RON tunnels and determines the best route to remote ARX devices.
RON_VPND	Not supported.
ROOTD	Root Daemon, which performs internal Unix management that requires root-level privileges.
SCM_DSMD	DNAS Switch Monitor.
SEC_CONFIG	Security Configuration.
SEC_LOCAL	Security Local Enforcement.
SECURITY_CLI	Security CLI.
SECURITYD	Security Daemon.
SHMEM	Common, shared, memory-management subsystem. This is primarily used for gathering statistics.
SHRTOP	Share topology library - Used for various filer-subshare computations.
SID_CACHE	Security descriptor translation daemon (XSDD) and its caches. XSDD translates between Security IDs (SIDs) on various file servers and the principal (group or user) names in a Windows domain.
SLMCONFD	SLM Configuration Daemon - not supported.
SMGMT	Storage Management Library.
SMTP	Used for file transfer (through email), email events, and other email-related features.
SNAPD	Snapshot daemon, which coordinates snapshots on multiple filers behind an ARX volume.
SQMD	Redundant Pair Site Quorum Manager, which makes failover decisions based on cluster integrity.
SSB_LIB	SSB Library. The SSB is an internal mechanism for processes to communicate with one another.
SSBEVENTD	SSB Event Daemon. The SSB is an internal mechanism for processes to communicate with one another.
SSH	SSH configuration.
SSRM	Simple Share Resource Monitor, which monitors filers on behalf of the HA subsystem.

Log Component	Description
STATSD	Statistics Daemon, which exposes all ARX statistics to SNMP, the CLI, and the ARX manager (GUI).
STATSMON	Statistics monitoring.
SUBSHRMGT	Subshare management daemon and its client library. This manages filer subshares. A filer subshare is any CIFS share that is inside another CIFS share.
SVCMGR_LIB	SvcMgr Client Library. This is deprecated, so it is unlikely to yield any log messages.
SYNC_CLOCK	Log messages related to sync_clock and sntpc on ASMs - not supported.
SYNC_LIB	Shadow Synchronization Library, which supports a shadow volume (see POLICY_SHADOW).
SYSLOG_NG_GUARD	syslog-ng_guard, which monitors and (if necessary) reanimates the syslog-ng process. The syslog-ng process maintains the syslog file, possibly duplicating it to a remote peer.
TASK_LIB	Task Dispatcher Library - Internal framework used by internal processes for managing event-driven tasks and scheduled tasks.
TCMD	NSM Configuration Manager.
TCORED	NSM Core Dump Manager.
TIME	Not supported.
TRANFS	Transitional File System in DNAS.
TRANFSD	Not supported.
TXNMOND	Object Manager Transaction Monitor. This monitors the internal database, which stores the ARX configuration. (You can see the full ARX configuration with 'show running-config' and 'show global-config.').
UTIL	Fast Filter IPC Utilities, for communication between internal ARX processes and the Chassis Manager (see CHASS).
VCIFS	dNAS CIFS Proxy, which handles requests to volume software on behalf of CIFS clients.
WDOG	Watchdog, which performs internal health checks on ARX hardware such as the SSB. The SSB is an internal mechanism for processes to communicate with one another.
WMA	Windows Management Authorization (WMA) CLI Commands and Library.
WMI	Software Library for managing snapshots on back-end Windows servers using Windows Remote Management protocol (WinRM) and Windows Management Instrumentation (WMI).

additional-command

Purpose You can set up *auto-diagnostics* to regularly gather system information and send it to F5 Support. F5 can use this information to monitor your ARX from off site, watching for gradual loss of service or resource degradation. The information sent to F5 is the output from a series of CLI show commands. You can use the additional-command operation to add another show command to this collection of diagnostics.

Use no additional-command to remove an extra show command from auto-diagnostics.

Mode gbl-auto-diag

Security Role(s) storage-engineer or crypto-officer

Syntax

```
additional-command "show-command-string"
no additional-command "show-command-string"
```

"*show-command-string*" (1-128 characters) is a valid show command in the ARX CLI, such as "show schedule." Surround the string with quotation marks (as shown) if it contains any spaces.

Default(s) No additional commands configured.

Guidelines This command adds an additional show command to the following list, which is always included in the auto-diags collection:

- show clock
- show chassis
- show processors
- show processors usage
- show health
- show version
- show running-config
- show global-config
- show ron
- show namespace all
- show arp all
- show redundancy all
- show cores
- show statistics namespace ... fastpath all fastpath
- show policy details
- show cifs-service user-sessions all summary
- show statistics cifs authentication all
- show active-directory status detailed
- show statistics filer detailed
- show external-filer all

Guidelines (Cont.)

You can add up to 20 additional commands to the above list.

The no form of this command cannot remove any of these default show commands.

In a redundant pair of ARX machines (see redundancy), both peers collect and send auto-diagnostics information through E-mail. All of the commands in gbl mode are shared between the peers, so only the active peer collects gbl-mode information (such as show global-config and show namespace all). Both peers collect and send commands related to cfg mode (such as show running-config and show ron). This also applies to any commands you add to the collection with additional-command; if the command relates to gbl mode, it is only included in collections from the currently-active peer. If the command relates to cfg mode, it is included in auto-diagnostic collections from both peers.

Sample

```
prtlndA(gbl-auto-diag)# additional-command "show channel summary"
```

adds a new show command to the auto-diags collection for the "prtlndA" chassis.

Related Commands

auto-diagnostics

auto-diagnostics

Purpose

F5 Support can monitor your ARX from off site, watching for gradual loss of service or resource degradation. This is accomplished by the ARX gathering diagnostic information and sending it (encrypted) as an E-mail attachment to F5. You can set the schedule, along with additional E-mail recipients and other options. Use the auto-diagnostics command to begin setting up auto monitoring by F5 Support.

Use the no auto-diagnostics command to disable auto-diagnostics and remote monitoring.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
auto-diagnostics
no auto-diagnostics
```

Default(s)

```
no auto-diagnostics
```

Guidelines

You must configure SMTP on the ARX before you enable auto-diagnostics, which uses E-mail as a delivery mechanism. Start with the smtp command in cfg mode.

We recommend activating this feature so that F5 Support can monitor your ARX and watch for impending issues. This also provides F5 Support with you updated configuration data, so that you will not be required to provide this data at the beginning of a support call. This eases the burden of backing up your running-config and global-config; this provides a second backup for both configuration files.

This command places you in gbl-auto-diag mode, where you must use the schedule (gbl-auto-diag) command to assign a schedule for collecting and sending diagnostic information. You can optionally use the mail-to (gbl-auto-diag) command to add additional E-mail recipients; the attachment is only encrypted for the E-mail to F5 Support. You can also use additional-command to add more show commands to the E-mail. To test the configuration, you can use the auto-diagnostics test command in priv-exec mode. The show auto-diagnostics command shows the current configuration for this utility, along with the status of the most-recent auto-diagnostics collection.

In a redundant pair of ARX machines (see redundancy), both peers collect and send auto-diagnostics information through E-mail. All of the commands in gbl mode are shared between the peers, so only the active peer collects gbl-mode information (such as show global-config and show namespace all). Both peers collect and send commands related to cfg mode (such as show running-config and show ron). This also applies to any commands you add to the collection with additional-command; if the command relates to gbl mode, it is only included in collections from the currently-active peer. If the command relates to cfg mode, it is included in auto-diagnostic collections from both peers.

To address an immediate issue, use the collect command to gather a larger amount of diagnostic information and possibly send it to F5 Support.

Guidelines: Show Commands Included in the Auto Diagnostics Collection

The process always includes the following show commands in its E-mail attachment:

- show clock
- show chassis
- show processors
- show processors usage
- show health
- show version
- show running-config
- show global-config
- show ron
- show namespace all
- show arp all
- show redundancy all
- show cores
- show statistics namespace ... fastpath all fastpath
- show policy details
- show cifs-service user-sessions all summary
- show statistics cifs authentication all
- show active-directory status detailed
- show statistics filer detailed
- show external-filer all

Sample

```
bstnA(gbl)# auto-diagnostics
bstnA(gbl-auto-diag)# ...
```

starts configuring auto diagnostics (and remote monitoring) for the “bstnA” switch.

Related Commands

schedule (gbl-auto-diag)
mail-to (gbl-auto-diag)
additional-command
collect
smtp
auto-diagnostics test
show auto-diagnostics

auto-diagnostics test

Purpose F5 Support can monitor your ARX from off site, watching for gradual loss of service or resource degradation. This is accomplished by the ARX gathering diagnostic information and sending it (encrypted) as an E-mail attachment to F5. After you set up the schedule for this auto-diagnostics data collection, you can use the auto-diagnostics test command to perform a test run.

Mode priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax `auto-diagnostics test [local]`

local (optional) sends the E-mail to local recipients only, defined by the mail-to (gbl-auto-diag) command. This does *not* send any E-mail to F5.

Default(s) None.

Guidelines This command gathers auto diagnostics (as determined by the auto-diagnostics command and its sub-mode commands), compresses them into a Unix Tape Archive (tar) file, and sends the file through smtp as an E-mail attachment. This tests both the auto-diagnostics configuration and the smtp configuration.

By default, the E-mail goes to F5 Support and all configured mail-to (gbl-auto-diag) addresses. The E-mail attachment for F5 Support is encrypted, as it must travel outside your local network. The E-mail attachment is un encrypted for each of the local mail-to (gbl-auto-diag) addresses, for convenience. You can use the local option to send the test message only to the local addresses.

The E-mail messages from this test have a different "subject" line than standard messages from scheduled runs: the subject line begins with "Test:".

Sample

```
bstnA# auto-diagnostics test local
```

collects and sends auto diagnostics, un encrypted, to the local E-mail addresses set by the mail-to (gbl-auto-diag) command.

Related Commands

auto-diagnostics
mail-to (gbl-auto-diag)
smtp

clear health

Purpose Certain SNMP traps raise persistent alarms on the system. For each "raise" trap, there is a corresponding "clear" trap that clears the alarm condition. You can use the clear health command to manually clear a raised alarm condition.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear health`** ***`trap`***

trap (1-64 characters) identifies the trap(s) to be cleared. Use show health to see a list of traps with open alarm conditions. Use the Event name from the output of show health.

Default(s) None

Guidelines The show health command shows all currently-raised traps on the system. For cases where a raised trap is unlikely to every be cleared by its corresponding "clear" trap, you can use this command to manually clear the trap. This clears all instances of the given trap. Then the alarm condition no longer appears in the show health output.

Sample

```
bstnA# clear health diskFail
```

clears all diskFail traps.

Related Commands show health

collect

Purpose Some machine problems require more diagnosis than is possible on-site. Use the collect command to collect diagnostic data on the switch and either upload it or store it locally.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax: Uploading to an FTP Site

```
collect {all | diag-info | config | cores | logs | stats-logs |
  reports |  state | capture}
  ftp://[user:password@]server/file.tgz  [async]
```

all | diag-info | config | cores | logs | stats-logs | report | state | capture chooses the type of diagnostics to be collected:

all collects all diagnostic files on the switch.

diag-info collects all diagnostic files on the switch except reports, which can be excessively large.

config collects the configuration database, database logs, and the output from show running-config and show global-config.

cores gathers any core-dump files on the system, along with all shared-object libraries used by the system.

logs assembles the syslog, the procdat log (process data; a more-detailed log file for diagnosis by engineers), upgrade and installation logs, disk-usage logs, Kerberos state files, and logs from the system's service manager. For options to collect logs from specific times, see the documentation for collect logs.

stats-logs selects all stats-log files, .csv files with performance data relevant to the ARX, its filers and servers, and other external devices. You can use the show stats-logs command to list all of the stats-log files currently available.

reports chooses all reports.

state collects individual log files from every processor in the switch, three system log files (syslog, procdat, and traplog), a chronological list of all CLI commands used since the switch booted, the latest of each of the stats-log files, and the output from a series of useful show commands.

capture selects only the packet-capture files, if there are any. These files show the flow of IP traffic through the ARX during a capture session. You can use show capture to list all of the capture files currently available.

ftp://[*user*[*:password*]@] *server/file*.tgz (1-255 characters) is the destination URL:

***user*[*:password*]@** (optional) is the username and password (for example, "random:pw@"). If you omit the password, the CLI prompts you for one.

server is the machine name (for example, "mymachine.myco.com").

file*.tgz** is the desired path name for the file. Lead with an extra slash (/) if the path is absolute: for example, "...myserver//var/log/acopiaDiags.tgz" specifies /var/log/acopiaDiags.tgz on "myserver." Use only a single slash if the path is local to the ***user's home directory.

Syntax: Uploading to an FTP Site (Cont.)

async (optional) makes this command return immediately, rather than waiting for the collection process to finish. The CLI generates a report that you can use to chart the collection process; the report name appears after the process starts.

Syntax: Secure Upload through SCP

```
collect {all | diag-info | ... | capture}
  scp://[user@]server:file.tgz [accept-host-key] [async]
```

{all | ... | capture} define the collection process, as explained above.

scp://[*user*@]*server*:*file*.tgz (1-255 characters) is the URL for the destination file:

> ***user***@ (optional if someone created an ip scp-user) is the username to present to the other end of the SCP connection. This user must be valid at the remote host. If you omit this and an ip scp-user is defined, it defaults to the username set by that command.
>
> ***server*:** is the IP address or hostname for the SCP host. End with a colon (:).
>
> ***file*.tgz** is the destination-file path. Lead with a slash ("/") if the path is absolute (for example, "scp://root@10.1.1.5:/var/diags/arx.tgz"). Use no slash if the path is local to the home directory for ***user*** (for example, "scp://root@10.1.1.5:diags.tgz").

accept-host-key (optional) indicates that if the other end of the connection has an unknown SSH host key (that is, if it is new, or if its key has changed since the last time the host was contacted), the ARX should accept the new host key and continue with the upload. Otherwise, the ARX stops the upload if the host presents an unknown key.

async (optional) was explained above.

Syntax: Sending to an ARX Volume

```
collect {all | diag-info | ... | capture}
  {nfs|cifs} namespace vol file.tgz [async]
```

{all | ... | capture} are explained above.

nfs | cifs is a required choice. This chooses the protocol for the file transfer.

namespace (1-30 characters) identifies the namespace to hold the collect file.

path (1-1024 characters) is the volume name.

***file*.tgz** is the path to the collect file, starting at the root of the volume.

async (optional) was explained above.

Syntax: Uploading to a TFTP Server

```
collect {all | diag-info | ... |capture} tftp://server/file.tgz
  [async]
```

{all | ... | capture} are explained above.

tftp://*server*/*file*.tgz (1-255 characters) is the URL for the collect file:

> ***server*** is the machine name (for example, "mymachine.myco.com").
>
> ***file*.tgz** is the desired path name for the file. Lead with an extra slash (/) if the path is absolute: for example, "...myserver//var/log/acopiaDiags.tgz" specifies /var/log/acopiaDiags.tgz on "myserver." Use only one slash if the path is local to the server's "tftpboot" directory. This conforms with the specification for FTP URLs in RFC 1738.

async (optional) was explained above.

Syntax: Sending to an E-mail Address

```
collect {all | diag-info | ... | capture}
  smtp://[e-mail-address/]file.tgz [async]
```

{all | ... | capture} are explained above.

smtp://[*e-mail-address*/]*file*.tgz (1-255 characters) is an E-mail destination for the collect file:

smtp:// declares that the destination is an E-mail address.

e-mail-address (optional) is the recipient of the E-mail in *username@host* format (for example, "jsmith@myco.com"). If you omit this, the CLI uses the default address set by the cfg-smtp to command.

***file*.tgz** is the name of the copy. This is sent as an attachment to the outbound E-mail message. This is not a path, so do not use slashes (/). You must use a .tgz extension. For example, "...acopiaDiags.tgz" is a valid file name.

async (optional) was explained above.

Syntax: Collecting Locally

```
collect {all | diag-info | ... |capture} local-file.tgz [async]
```

{all | ... | capture} are described above.

***local-file*.tgz** (1-255 characters) saves the file locally, to the diag-info directory. To preserve disk space, you can only save one diag-info file at a time. You must use a .tgz extension.

async (optional) was explained above.

Default(s)

(FTP only) If you omit the *user*:*password* from the FTP URL, the CLI uses the username and password set by the ip ftp-user command.

Guidelines

This file contains multiple log files and possibly core files, so it can be very large. If you upload it, choose a host with ample disk space. If you send the message through SMTP, configure the local E-mail server (chosen with the mail-server command) and the recipient E-mail server to allow very large E-mail messages. (The E-mail server at F5 Support allows for large messages.)

For SCP uploads, the CLI prompts for a password; enter the password for the ***user*** that you identified in the URL. The CLI also prompts for a password for FTP uploads where you do not supply one in the command.

The CLI prompts for confirmation before collecting the diagnostics. Answer **yes** to continue. The collection process can be time-consuming; prompts appear to show the progress of the operation as it occurs.

You must configure SMTP on the ARX before you use the "smtp://" syntax for E-mail. Start with the smtp command in cfg mode.

For saves to the local disk, you can only save one "collect" file at a time. If one already exists, the CLI prompts for permission to overwrite the previously-collected file. Use show diag-info to show that the file was created on disk.

Samples

```
bstnA# collect diag-info ftp://jpublic@arxftp.f5.com/10-24-03.tgz
Password: jpassword

Collect diagnostic information? [yes/no] yes
```

sends all diagnostic information to the FTP server at arxftp.f5.com.

```
bstnA# collect diag-info cifs medarcv /rcrds admin/collect.tgz

Collect diagnostic information? [yes/no] yes
```

sends all diagnostic information to a CIFS volume on the ARX, "medarcv~/rcrds."

```
bstnA# collect config ftp://arxftp.f5.com/10-24-03.tgz

Collect diagnostic information? [yes/no] yes
```

sends only config information. This uses the FTP credentials established by the ip ftp-user command.

```
bstnA# collect diag-info investigate.tgz

Only one copy of diagnostic information may be stored on the switch at
one time.  This automatically removes any saved diagnostic information.

Are you sure? [yes/no] yes
```

collects all diagnostic information in a local file.

```
prtlndA# collect logs smtp://logsAndReports.tgz

Collect diagnostic information? [yes/no] yes
```

collects all logs and sends them in an E-mail attachment.

Related Commands

collect logs
copy ftp
copy scp
copy smtp
copy {nfs|cifs}
copy tftp
ip ftp-user
ip scp-user
smtp
show diag-info

collect logs

Purpose You can collect all log files on the system at once, or you can focus on the logs from a specific time. This is an extension of the collect command that applies to log files only.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
collect logs [start-time start [utc]] [end-time end [utc]]
destination [async]
```

start-time *start* [utc] (optional) sets a starting time for collected logs. If you omit this, the CLI collects all of its logs up to the **end-time**.

start is a time stamp in one of two formats:

- ***mm/dd/yyyy:HH:MM:SS*** is designed for manual entry of a date and time.
- ***yyyy-mm-ddTHH:MM:SS*** is designed for a copy-and-paste of a time stamp from a log file. Time stamps from log files are in UTC instead of local time.

utc (optional) declares that the start is in UTC instead of local time. This is typically used for time stamps copied from a log file (the second format from above).

end-time *end* [utc] (optional) sets an end time for collected logs. If you omit this, the CLI collects all of its logs after the **start-time**. The details of this time stamp are otherwise the same as above.

destination is a local-file name, an ARX volume, or a URL. The URL can use FTP, SCP, TFTP, or SMTP. The documentation for the collect command describes the exact syntax.

async (optional) makes this command return immediately, rather than waiting for the collection process to finish. The CLI generates a report that you can use to chart the collection process; the report name appears after the process starts.

Default(s) If you omit the optional flags, the CLI collects all logs.

Guidelines This is an alternative syntax for the collect command. It is useful for collecting logs from a specific time frame; this can reduce the size of the collection as well as focus a diagnosis.

Samples

```
bstnA# collect logs start-time 06/05/2007:12:00:00 end-time
06/05/2007:13:00:00 ftp://jpublic@arxftp.f5.com/1hr
Password: jpassword

Collect diagnostic information? [yes/no] yes
```

collects logs from noon until 1PM on June 5, 2007. It then sends the collected logs to an FTP site.

```
prtlndA# collect logs start-time 2007-01-07T05:00:34 utc sinceBoot.tgz

Only one copy of diagnostic information may be stored on the switch at
one time.  This automatically removes any saved diagnostic information.

Are you sure? [yes/no] yes
```

collects all logs since the given time (in UTC format). All of the logs are placed in a local file named "sinceBoot.tgz."

Related Commands collect

expect monitor

Purpose You can use the expect monitor command to repeat a show command at regular intervals.

Mode priv-exec

Security Role(s) crypto-officer, storage-engineer, network-engineer, or network-technician

Syntax **`expect monitor`** ***`seconds “show-command”`*** **`[`**`“search-string”`**`]`**
`[timeout` ***`timeout-seconds`*** **`]`**

seconds is the number of seconds to wait between each invocation of the command.

“*show-command*” is a valid CLI show command, surrounded by quotation marks (““). Use any of the show commands documented in this manual.

“*search-string*” (optional) also requires quotes. If you use this, the command stops repeating as soon as the *search-string* appears in the output. Otherwise, the command repeats until you use **<Ctrl-C>** to stop it.

timeout-seconds (optional, 1-2096) sets a time limit on this monitor process. The show commands stop repeating when this time expires. You can also use <Ctrl-C> to stop the process at any time.

Default(s) ***timeout-seconds*** - 21,600 (6 hours)

Guidelines You can use this command to watch the progress of a configuration change from the CLI. For example, this command can repeat the show redundancy command after you enable (cfg-redundancy) redundancy for a pair of ARXes, to monitor their initial rendezvous.

You cannot issue other CLI commands in this session while the expect monitor command is running. Use **<Ctrl-C>** to stop the repetition at any time.

To run a full CLI script, copy the script onto the ARX (using copy ftp, copy scp, copy {nfslcifs}, or copy tftp), and then use run to invoke it. To perform any command in the future (possibly on a schedule), use the at command.

Samples

```
bstnA# expect monitor 5 “show namespace wwmed”
```

repeats the show namespace wwmed command every 5 seconds, until you use <Ctrl-C> to stop it.

```
bstnA# expect monitor 10 “show namespace medarcv” “/rcrds”
```

repeats the show namespace medarcv command every 10 seconds until “/rcrds” appears in the output.

Related Commands
at
run

logging destination

Purpose Use the logging destination command to define an external syslog server as an external logging destination.

Use the no form of the command to remove a syslog server.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
logging destination server-name [udp | tcp] [port port]
no logging destination server-name
```

server-name (1-80) is the IP address or hostname of the external syslog server.

udp | tcp (optional) is the transport protocol over which to send the syslog messages.

port *port* (optional, 1-65,535) is the UDP or TCP port number where the external server is listening.

Default(s)

```
udp
port 514
```

Guidelines The external server receives log messages as they are generated and appends them to its local syslog file.

Use the show logging destination command to list all configured external syslog servers.

Samples

```
bstnA(cfg)# logging destination 10.1.1.76 tcp
```

sends syslog messages over TCP to an external syslog daemon at 10.1.1.76.

```
bstnA(cfg)# no logging destination 172.16.44.7
```

stops sending syslog messages to 172.16.44.7.

Related Commands show logging destination

logging level

Purpose

Each group of log messages, known as a log component, has a separately-tunable *logging level*. Critical level is the most terse, displaying critical log messages only; debug level is the most verbose, displaying all log messages from critical down to debug. The messages appear in the syslog file. Use the logging level command to set a logging level.

Use the no form of this command to revert a logging level back to the default.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
logging level component {critical | error | warning | notice |
info | debug | disable}
logging level all {critical | error | warning | notice | info |
debug | disable}
no logging level {component | all}
```

component (up to 255 characters) is the component to tune. See Log Components, on page 38-3 for a complete list of components.

all selects all components.

critical | error | warning | notice | info | debug | disable is a required choice. This is the logging level you choose for the component. The **disable** option stops all logging from the chosen ***component***.

Default(s)

info

Guidelines

Use show logging levels to verify the log-level settings for each component.

From any mode, use show logs syslog or grep *pattern* logs syslog to view the log messages in the syslog file. See the manual, *ARX Log Catalog*, for a full list of log messages.

This command applies to components on the ACM. To enable log messages for the NSM, use logging fastpath processor and logging fastpath component.

Samples

```
bstnA(cfg)# logging level NSCKRPT debug
```
sets the logging level to debug (as verbose as possible) for NSCKRPT messages.

```
bstnA(cfg)# logging level all error
```
sets the logging level to error (fairly terse) for all messages.

```
bstnA(cfg)# no logging level all
```
sets the logging level to the default for all messages.

```
bstnA(cfg)# logging level RON disable
```
disables logging for the RON component.

```
bstnA(cfg)# no logging level POLICY_ACTION
```
sets the logging level to the default for POLICY_ACTION messages.

Related Commands

show logging levels
show logs
grep
logging fastpath processor
logging fastpath component

mail-to (gbl-auto-diag)

Purpose The auto-diagnostics utility gathers performance data and sends it to F5 Support. It sends the data as an E-mail attachment. You can use the mail-to command to add an additional destination address for these E-mails.

Use the no mail-to command to remove an E-mail recipient.

Mode gbl-auto-diag

Security Role(s) storage-engineer or crypto-officer

Syntax

```
mail-to recipient
no mail-to recipient
```

recipient (1-128 characters) is one E-mail recipient (for example, "juser@nemed.com").

Default None

Guidelines This sets a destination for auto-diagnostics E-mail messages. This E-mail recipient receives a copy of every auto-diagnostics collection that the ARX assembles. This is in addition to the copy sent to F5 Support. The attachment delivered to F5 Support is encrypted, but the attachment sent to this E-mail recipient is not.

The maximum number of E-mail recipients is limited by the total bytes required to store them. If they are concatenated into a single string, separated with commas, the total string length cannot exceed 1024 bytes.

You cannot remove F5 Support as an E-mail recipient.

Samples

```
bstnA(gbl-auto-diag)# mail-to juser@wwmed.com
```

sets up one E-mail recipient for the "bstnA" switch. Whenever the bstnA switch sends out auto-diagnostics to F5 Support, it also sends an un encrypted copy of the diagnostics to juser@wwmed.com.

```
prtlndA(gbl-auto-diag)# no mail-to ex@nemed.com
```

removes an E-mail recipient.

Related Commands auto-diagnostics

management source

Purpose

Use the management source command to set a source address for sending syslog messages.

Use the no form to remove the management-source configuration.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`management source {mgmt | vlan` ***`vlan-id`*****`}`**
`no management source`

mgmt is the switch's out-of-band management interface, the default. You specify the out-of-band management IP address at switch boot-up.

vlan specifies to use a configured VLAN.

vlan-id (1-4096) identifies the VLAN to use as the source. Use the show vlan summary command for a list of all configured VLANs.

Default(s)

mgmt on platforms that support this command

vlan 1 on platforms that do not support this command.

Platforms

any *except* ARX-VE

Guidelines

Use logging destination to send syslog messages to an external syslog server. This command sets the source address for those log messages.

If you remove a VLAN interface that is configured as the management source, the syslog does not function properly. You must configure another source (mgmt or VLAN).

The in-band management address for VLAN 1 (the default VLAN) is always the management source for an ARX-VE. The ARX-VE has no out-of-band management interface.

Use the show logging destination command to view the current management source.

Sample

```
bstnA(cfg)# management source vlan 9
```

configures VLAN 9 as the management source for sending syslog traffic and SNMP traps.

Related Commands

logging destination
show logging destination
show vlan summary

moving-average

Purpose

**** This is a beta-level command. It does not appear in the CLI unless you use the* terminal beta *command to unlock all beta-level commands. ****

The ARX records the round-trip times between itself and external devices, such as CIFS clients and back-end filers. The *stats-monitor* process keeps moving averages of these round-trip times and other statistics, and records them in .csv files in the stats-log directory. The stats monitor also can log *notification* messages if the number of errors and/or response time from an external device increases dramatically. The stats monitor only logs a notification if the device has these increases multiple times in a row. Use the moving-average command to define the percentage of increase required, and the number of times that the errors or response times must increase that much to merit a notification.

Use no moving-average to return to the default thresholds.

Mode

gbl-stmon-nfy

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
moving-average {response-time | errors} threshold pct-threshold
moving-average {response-time | errors} interval num-samples

no moving-average {response-time | errors} threshold
no moving-average {response-time | errors} interval
```

response-time | errors is a required choice. This determines the type of threshold or interval you are setting, response times or errors.

threshold ***pct-threshold*** (optional, 1-65,535) is the percentage above the moving average that could trigger an alarm. For example, if you set **moving-average response-time threshold 200**, the stats monitor may trigger a notification if a device's response time is 200% above (or, 3 times) its current moving average. As another example, if you set **moving-average errors threshold 150**, stats monitor may trigger a notification if the device's number of errors is 150% above its current average. The stats monitor process only triggers the notification if this occurs for some number of samples, as set by the **interval** option.

interval ***num-samples*** (optional, 1-65,535) is the number of high samples in a row that are required to trigger a notification. If this many samples in a row are at least ***pct-threshold*** above the current average, the stats-monitor process logs a notification.

Default(s)

either **threshold** - 100%

either **interval** - 6

Guidelines

This is a beta-level command. Use the terminal beta command to unlock all beta-level commands, including this one.

The defaults are sufficient for most sites. Use this command only on the advice of F5 Support.

This command changes the moving-average thresholds for one set of filers or ARX front-end services, chosen when you use the notify command to enter the mode.

The *moving average* is the average for a number (for example, 8) of the most-recent samples. The 8th time the software calculates a moving average, it takes the average of samples 1 to 8; the 9th time it runs, it averages samples 2 to 9; the 10th time, it takes the average of samples 3 to 10; and so on.

If a group of samples deviates from the current moving average, the stats monitor generates a notification. Use this command, moving-average, to decide how many samples in a row it takes to log the notification, and to decide how far over the current average a sample must be to qualify.

These notifications are recorded in the syslog; you can use show logs syslog to view them along with all other syslog messages. Look for the log component named "SMON_ANALYSIS," or use grep to search for it. To create an additional SNMP trap for each notification, use the trap command.

The stats monitor collects data for a long period of time, called a *moving-average window*, before it starts using the moving average for notifications. This window is calculated from the moving average and sampling interval commands: it is 8 * *moving-average interval* * *sampling interval*. By default, that is 8 * 6 * 5 minutes, or four hours. The moving-average window restarts if you use the sampling interval command to change the sampling interval.

Samples

```
bstnA(gbl-stmon-nfy)# moving-average response-time threshold 10
bstnA(gbl-stmon-nfy)# moving-average response-time interval 3
```

sets the following tolerances for response times: if a device's response time is **10**% (or more) slower than its current average for **3** samples in a row, stats monitor logs a notification message. It also sends an SNMP trap if trap is enabled.

```
bstnA(gbl-stmon-nfy)# moving-average errors threshold 5
bstnA(gbl-stmon-nfy)# moving-average errors interval 3
```

sets the following tolerances for errors: if a device returns **5**% (or more) errors than its current average for **3** samples in a row, stats monitor logs a notification message. As above, stats monitor also sends an SNMP trap if trap is enabled.

Related Commands

notify
trap

notify

Purpose

**** This is a beta-level command. It does not appear in the CLI unless you use the* terminal beta *command to unlock all beta-level commands. ****

The ARX records the round-trip times between itself and external devices, such as CIFS clients and back-end filers. You can use the notify command to trigger analysis for one group of external devices: CIFS-service clients, NFS-service clients, CIFS filers, and or NFS filers. Whenever an issue is discovered for one of the analyzed device groups, the stats-monitor process logs a *notification* in the syslog. Optionally, the stats monitor can also send an SNMP trap. This command also brings you into a submode, where you can alter the notification settings for the chosen group of devices.

You can use the negative form of the command, no notify, to deactivate this stats-monitor process for front-end or back-end services.

Mode

gbl-stmon

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
notify {cifs-service | nfs-service}
notify filer-share {cifs | nfs}

no notify {cifs-service | nfs-service}
no notify filer-share {cifs | nfs}
```

cifs-service | nfs-service is a required choice. The nfs-service option starts analyzing all nfs front-end services, and the cifs-service option starts analyzing all cifs front-end services.

cifs | nfs is a required choice if you choose to monitor filer shares. This chooses between analyzing CIFS or NFS on back-end filers.

Default(s)

```
no notify cifs-service
no notify nfs-service
no notify filer-share cifs
no notify filer-share nfs
```

Guidelines

This is a beta-level command. Use the terminal beta command to unlock all beta-level commands, including this one.

The notify command places you in gbl-stmon-nfy mode, where you can use other CLI commands to edit the notifications for the set of devices you have chosen. The notifications only occur when some number of samples deviates from the moving average by some percentage; for example, a notification occurs when 6 samples from a particular CIFS service have twice as many NetworkErrors as the moving average, or when 6 samples from a particular back-end NFS export have progressively-higher round-trip times for GETATTR requests. You can use the moving-average command to reset the number of samples to trigger a notification, and/or to change the threshold for each sample to qualify as "too high." By default, a notification goes to the syslog only; the trap command enables SNMP traps in addition to those syslog messages.

The stats monitor records its notifications in the syslog; you can use show logs syslog to view them along with all other syslog messages. Look for the log component named "SMON_ANALYSIS," or use grep to search for it. To create an additional SNMP trap for each notification, use the trap command.

Sample

```
bstnA(gbl-stmon)# notify cifs-service
bstnA(gbl-stmon-nfy)#
```

enables analysis of all front-end cifs services.

Related Commands

stats-monitor
moving-average
trap
show health

sampling interval

Purpose

**** This is a beta-level command. It does not appear in the CLI unless you use the* terminal beta *command to unlock all beta-level commands. ****

The ARX records the round-trip times between itself and external devices, such as CIFS clients and back-end filers. The stats-monitor process keeps running averages of these round-trip times and other statistics, and records them in .csv files in the stats-log directory. You can use the sampling interval command to determine the interval at which these samples are recorded.

Use no sampling interval to return to the default interval.

Mode

gbl-stmon

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
sampling interval seconds
no sampling interval
```

seconds (30-86,400) is the number of seconds between each sample. 86,400 seconds is one day.

Default(s)

300 (seconds, or 5 minutes)

Guidelines

This is a beta-level command. Use the terminal beta command to unlock all beta-level commands, including this one.

The default is sufficient for most sites. Use this command only on the advice of F5 Support.

If you select a sampling interval that is shorter than the default, the stats-monitor process writes .csv files more frequently. These raw .csv files are kept (together with all other log files) on a separate disk partition from configuration files and other important data. Each hour, the stats-monitor also creates an hourly file with the hour's first set of samples; this is a summary of the data in the raw file. Use show stats-logs to see a list of these .csv files, or show stats-logs *file-name* to view the contents of one of them. The stats-monitor files are not permitted to consume more than 50% of the log partition. If the .csv files approach their size limit, the oldest raw-data files are removed and the hourly-summary files remain.

This command restarts the calculation of all moving averages. After you invoke this command, stats monitor must run for the full *moving-average window* before it can log any more notifications (or, possibly, send any more traps). This window is calculated from the sampling interval command as well as the moving-average command: it is 8 * *moving-average interval* * *sampling interval*. By default, that is 8 * 6 * 5 minutes, or four hours.

Sample

```
bstnA(gbl-stmon)# sampling interval 60
```

reduces the sampling interval (and raises the sampling frequency) to once every minute.

Related Commands

stats-monitor
show stats-logs
moving-average

schedule (gbl-auto-diag)

Purpose

The auto-diagnostics utility gathers data about the ARX, encrypts it, and sends it in an E-mail message to F5 Support. Use this schedule command to assign a schedule for this operation.

Use no schedule to remove the auto-diag schedule, thereby disabling auto diagnostics.

Mode

gbl-auto-diag

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
schedule name
no schedule
```

name (1-64 characters) identifies the schedule for auto diagnostics. Use show schedule for a list of configured schedules. Choose a schedule that has a frequency measured in days, weeks, or a longer time measure; for example, every day, every 3 days, every 1 week, or every 2 months. This command does not accept a schedule with a higher frequency, such as every 12 hours.

Default(s)

None.

Guidelines

The auto diagnostics utility gathers performance information about your system and regularly sends it to F5 Support. This command sets the schedule for sending this information. The auto diagnostics utility requires a schedule in order to run automatically.

To create a schedule, use the gbl-mode schedule command. The schedule cannot run more frequently than once per day; you can define this with the every command.

Sample

```
minturnA(gbl-auto-diag)# schedule every2am
```

uses a schedule named “every2am” for auto-diagnostic collection on “minturnA.”

Related Commands

auto-diagnostics
schedule
show schedule

show auto-diagnostics

Purpose F5 Support can monitor your ARX from off site, watching for gradual loss of service or resource degradation. This is accomplished by the ARX gathering diagnostic information and sending it (encrypted) as an E-mail attachment to F5. After you set up the schedule for this auto-diagnostics data collection, you can use the show auto-diagnostics command to show the auto-diagnostics configuration along with the status of the latest auto-diagnostics run.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show auto-diagnostics`

Guidelines The output contains several tables of configuration information, followed by a table showing the status of the last auto-diagnostics run.

Schedule is the schedule's configured name, chosen with the schedule (gbl-auto-diag) command. The fields in this table show the configuration of the schedule.

Description only appears if set by the gbl-schedule description (gbl-schedule) command.

Start Time can be reset by the gbl-schedule start command. This is the date and time of day used for all time-based calculations: for example, if a schedule runs every Saturday, this determines the time of day on Saturdays when the schedule fires.

Stop Time only appears if set by the gbl-schedule stop command. This is the expiration time for the schedule, after which the schedule never fires.

Interval is set by the gbl-schedule every command. This is the time between scheduled runs.

Duration appears only if set by the gbl-schedule duration command. This is not applicable to auto-diagnostics.

Status is "Paused" (the schedule is between run times), "Running," or "Waiting for start time" (the schedule has never been invoked).

Previous is a sub-table showing the Run time (when the run started) and End time of the most-recent run.

Current appears if the rule is currently running. This has the same fields as the Previous table, Run time and End time.

Next shows when the schedule's next run will begin. It also has the Run time and End time fields.

Mail-to is a list of local E-mail addresses, added with the mail-to (gbl-auto-diag) command. Every time the above schedule fires, the system sends a tar file with auto-diagnostic output to each of these addresses. It also sends an encrypted version of the same tar file to F5 Support.

Guidelines (Cont.) Additional commands are extra CLI-show commands that are collected along with the default auto-diagnostics. You can add or remove these with the [no] additional-command command.

Status describes the most-recent auto-diagnostics collection and delivery operation since the last reboot. The following fields appear only if an auto-diagnostics collection has occurred:

Last time executed is the most-recent time that the schedule fired and the collection process started.

Status of last execution is either "Success" or "Failure" with a report name in square brackets. You can use show reports *report-name* to read the report and find the specific failure.

Date last sent is a time stamp for the most-recent E-mail delivery.

Number of times sent since last reboot shows how many times the system has collected and delivered auto-diagnostics data since the last ARX reboot.

Sample `bstnA(cfg)# show auto-diagnostics`

shows the auto-diagnostics configuration and status for the "bstnA" switch. See Figure 38.1 for sample output.

Related Commands auto-diagnostics

Figure 38.1 *Sample Output: show auto-diagnostics*

```
bstnA(cfg)# show auto-diagnostics

  Schedule:         daily4am
    Description:      two hours between 4 and 6 AM
    Start Time:       Sun Sep  4 04:00:00 2005
    Stop Time:        Wed Jan  7 04:00:00 2015 (Expires in 1681 d 03:27:36)
    Interval:         1 days
    Duration:         02:00:00
    Status:           Paused (runs in 02:27:36)

    Previous:
    Run Time:         Mon May 31 04:00:00 2010
    End Time:         Mon May 31 06:00:00 2010

    Next:
    Run Time:         Tue Jun  1 04:00:00 2010
    End Time:         Tue Jun  1 06:00:00 2010


  Mail-to:   juser@wwmed.com, mmorrison@wwmed.com

  Additional commands:
    show global service
    show share status

  Status:
    Last time executed: 01 Jun 2010 04:45:23 AM UTC
    Status of last execution: Success [ collect_diag_201006010445.rpt DONE ]
    Date last sent: 01 Jun 2010 05:27:42 AM UTC
    Number of times sent since last reboot: 2
```

show documentation

Purpose Many syslog and CLI-response messages have further documentation to help with problem assessment and diagnosis. Use the show documentation command to display the documentation for a given message.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show documentation`** ***`msg-catalog-id`***

msg-catalog-id (1-255 characters) identifies the message. Every message-catalog message has an uppercase ID that appears with it in the syslog and in the CLI.

Guidelines The output shows the documentation for a particular message-catalog message. The CLI, the syslog, and most processes share a common message catalog. The full message catalog is documented in an alphabetical reference, the *ARX Log Catalog*.

Sample `bstnA(cfg)# show documentation SQM_CANNOT_BE_PROMOTED_GRACE_PERIOD`
shows the documentation for one message ID. See Figure 38.2 for sample output.

Related Commands show logs **`syslog`**

Figure 38.2 *Sample Output: show documentation*

```
bstnA(cfg)# show documentation SQM_CANNOT_BE_PROMOTED_GRACE_PERIOD

SQM_CANNOT_BE_PROMOTED_GRACE_PERIOD documentation:

  In order to force a seniority promotion, the remote peer switch must be
  offline and the local switch must be stalled waiting for it to come online.
  This condition is indicated by the presence of the "Peer has more recent
  configuration..." (SQM_PTN_IN_TIME_STALL) message in the 'show redundancy
  history' output. Unless the switch has been in this stalled state for at
  least five minutes, it cannot be promoted.
```

show health

Purpose Certain SNMP traps raise persistent alarms on the system. For each "raise" trap, there is a corresponding "clear" trap that clears the alarm condition. Use the show health command to see all SNMP-related alarms that have been raised but not cleared.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show health`

Guidelines This shows a table with one row per active alarm. Each row has the following columns:

Date is the date and time when the event occurred.

ID is the last three digits in the trap's OID.

Event is the name of the trap.

Description explains the alarm condition.

Refer to the ARX SNMP Reference for a full list of all traps, along with explanations for each trap and procedures to address the issue.

To manually clear a raised trap, you can use the clear health command.

Follow the Guidelines for the command, snmp-server traps, to set up a remote host to receive SNMP traps. For an overall view of the system hardware, use show chassis.

Samples

```
prtlndA(cfg)# show health
```

shows any raised alarms on the "prtlndA" chassis. See Figure 38.3 on page 38-40 for sample output.

```
bstnA> show health
```

shows any raised alarms on the "bstnA" chassis. See Figure 38.4 on page 38-41 for sample output.

Related Commands

clear health
snmp-server traps
show chassis

Figure 38.3 *Sample Output: show health (No Active Alarms)*

```
prtlndA(cfg)# show health

System Health Information
Date                     ID      Event                       Description
-------------------------------------------------------------------------------
Thu Jan 25 04:27:29 2007 (0)   - No active alarms.
```

Figure 38.4 *Sample Output: show health (with Active Alarms)*

```
bstnA> show health

System Health Information
Date                      ID      Event                       Description
--------------------------------------------------------------------------------
Thu Jan 20 00:49:14 2011 (540) - powerFail                    Slot power 2/1 failed
Thu Jan 20 00:49:24 2011 (540) - powerFail                    Slot power 1/2 failed
Thu Jan 20 00:50:57 2011 (185) - gatewayOffline               Client gateway 10.46.28.218 is
offline.
Thu Jan 20 00:53:29 2011 (411) - nsmStandby                   Processor 2.2 is in standby.
Thu Jan 20 00:53:40 2011 (411) - nsmStandby                   Processor 2.4 is in standby.
Thu Jan 20 00:53:40 2011 (411) - nsmStandby                   Processor 2.3 is in standby.
Thu Jan 20 00:53:40 2011 (411) - nsmStandby                   Processor 2.1 is in standby.
Thu Jan 20 00:59:32 2011 (826) - virtualServiceAclUpdateFail 192.168.25.15:/acct:NIS information
has not been fully resolved by the NIS daemon.
Thu Jan 20 01:00:59 2011 (148) - downRevAdForestLevelRaise   The forest functional level of domain
MEDARCH.ORG is not capable of supporting constrained delegation
Thu Jan 20 01:08:40 2011 (19)  - archiveFreeSpaceThresholdRaise fileRecordsMed,
fs4:arx_file_archv:is running low on freespace (<10 GB Free).
Thu Jan 20 01:18:44 2011 (826) - virtualServiceAclUpdateFail 192.168.25.15:/claims:NIS information
has not been fully resolved by the NIS daemon.
```

show health time-skew

Purpose Several ARX applications require clock synchronization between the ARX and its network peers. These applications include Kerberos, which frequently uses time stamps, and the policy engine, which often places files based on their last-modified times. Use the show health time-skew command to see the clock differences between the ARX and other stations with which it communicates.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show health time-skew [timeout seconds]`**

timeout *seconds* (optional, 1-3600) is the maximum number of seconds to wait for each server response. The default timeout is 3 seconds. The ARX makes two or more attempts in some cases, and this timeout applies to each attempt. Also, the timeout applies to each individual server; the overall time for all servers to respond may be significantly longer. You can use <Ctrl-C> to stop the command at any time.

Guidelines We strongly recommend using NTP to keep the ARX clock synchronized with other system clocks on the network. Use the ntp server command to connect the current switch to one or more NTP servers.

The output shows a table with one row per network peer. Each row has the following columns:

IP Address identifies the network device.

Role shows the relationship of this device to the ARX. Possible values here are "External Filer" or "Domain Controller." Use the show external-filer command for a full list of external filers used by this ARX. Use the show active-directory command for a full list of DCs and a general understanding of how the ARX is using them.

Observed Time Skew ... shows a time value in *HH*:*MM*:*SS* format, followed by "ahead," "behind," or (for a time value of "00:00:00") nothing.

Sample `bstnA> show health time-skew timeout 30`

shows the time skews between the "bstnA" chassis and its network peers. See Figure 38.5 on page 38-42 for sample output.

Related Commands ntp server
show external-filer
show active-directory

Figure 38.5 *Sample Output: show health time-skew*

```
bstnA> show health time-skew
IP Address            Role                Observed Time Skew vs. ARX
--------------------  ------------------  ----------------------------
192.168.74.82         External Filer      00:18:41  ahead
192.168.25.24         External Filer      00:00:00
192.168.25.25         External Filer      04:16:19  behind
192.168.25.48         External Filer      3 d 03:59:15 ahead
192.168.25.22         External Filer      00:00:00
192.168.25.44         External Filer      00:00:00
```

```
192.168.25.19         External Filer        00:00:00  ahead
192.168.25.71         External Filer        00:00:00
192.168.25.27         External Filer        00:01:29  behind
192.168.25.20         External Filer        00:01:29  behind
192.168.25.28         External Filer        00:01:29  behind
192.168.25.29         External Filer        00:01:29  behind
192.168.25.21         External Filer        00:00:00
192.168.25.47         External Filer        00:00:00
192.168.74.90         External Filer        00:00:00
192.168.25.102        Domain Controller     00:01:29  behind
192.168.25.104        Domain Controller     00:01:29  behind
192.168.202.9         Domain Controller     00:01:30  behind
192.168.202.10        Domain Controller     00:01:30  behind
192.168.202.11        Domain Controller     00:01:29  behind
172.16.124.73         Domain Controller     00:01:29  behind
192.168.25.103        Domain Controller     00:01:29  behind
172.16.168.21         Domain Controller     00:01:29  behind
10.19.230.94          Domain Controller     00:01:30  behind
192.168.25.109        Domain Controller     00:01:30  behind
192.168.25.110        Domain Controller     00:01:30  behind
172.16.108.139        Domain Controller     00:00:41  behind
172.16.110.5          Domain Controller     00:00:41  behind
172.16.120.5          Domain Controller     00:00:41  behind
172.16.213.9          Domain Controller     00:01:06  behind
172.16.240.88         Domain Controller     00:01:07  behind
172.16.210.7          Domain Controller     00:01:06  behind
172.16.74.88          Domain Controller     00:01:30  behind
```

show id-mappings

Purpose Syslog messages use numeric IDs to identify namespaces, volumes, shares, and file-history archives. Use the show id-mappings command to map these IDs to the names used in the CLI.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show id-mappings`**

Guidelines The output shows four tables:

Namespace lists all namespaces on the switch, one per row. The ID column contains the numeric ID.

Volume lists all volumes on the switch, one per row. The Namespace column shows the volume's namespace, and the ID column contains the volume's numeric ID.

Share lists all shares on the switch, one per row. The Volume and Namespace columns show the share's context, and the ID column contains the share's numeric ID. Contact F5 Support if "Not Found" appears at the end of any of these rows; this indicates an issue in the configuration database.

Archive Name is a table of file-history archives on the ARX. A *file-history archive* is a repository for a managed volume's file-location records. You can use the file-history archive command to create one. The ID column in this table contains the archive's numeric ID.

Sample `bstnA(cfg)# show id-mappings`
shows the ID mappings on the current switch. See Figure 38.6 for sample output.

Related Commands show logs **`syslog`**

Figure 38.6 *Sample Output: show id-mappings*

```
bstnA(cfg)# show id-mappings

Namespace                          ID
-----------------------------  -----
medco                            1
wwmed                            2
medarcv                          3
insur                            4

Volume                    Namespace                    ID
------------------------  -------------------------  -----
/vol                      medco                        1
/acct                     wwmed                        2
/rcrds                    medarcv                      3
/lab_equipment            medarcv                      4
/test_results             medarcv                      5
/claims                   insur                        7

Share                Volume                Namespace                ID
-------------------  --------------------  ---------------------  -----
corporate            /vol                  medco                      1
```

```
sales                /vol                 medco                    2
generic              /vol                 medco                    3
budget               /acct                wwmed                    5
bills                /acct                wwmed                    6
bills2               /acct                wwmed                    7
it5                  /acct                wwmed                    8
metadata-share       /acct                wwmed                    4
rx                   /rcrds               medarcv                 11
charts               /rcrds               medarcv                 12
bulk                 /rcrds               medarcv                 13
metadata-share       /rcrds               medarcv                 10
equip                /lab_equipment       medarcv                 15
leased               /lab_equipment       medarcv                 16
backlots             /lab_equipment       medarcv                 19
scanners             /lab_equipment       medarcv                 20
metadata-share       /lab_equipment       medarcv                 14
equipSnap            /lab_equipment       medarcv                 26
leasedSnap           /lab_equipment       medarcv                 28
chemistry            /test_results        medarcv                 21
hematology           /test_results        medarcv                 22
2005_charts          /test_results        medarcv                 23
shr1-old             /claims              insur                   31
shr1-next            /claims              insur                   32
metadata-share       /claims              insur                   30

Archive Name                           ID
-------------------------------------  -----
fileRecordsMed                           29
```

show logging destination

Purpose Use the show logging destination command to view a list of configured external syslog servers.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show logging destination`**

Guidelines The output contains the Source Interface used for communicating with the syslog servers, followed by a list of syslog servers.

Use the logging destination command to define a syslog server.

Use the management source command to define a source address for sending syslog messages and traps.

Sample

```
bstnA# show logging destination

Source Interface:   mgmt (10.1.1.7)

External Syslog Destination(s)
------------------------------
172.16.202.8
```

Related Commands logging destination
management source

show logging levels

Purpose Log messages are divided into groups based on the system components that generate them. You can set a logging level for each component, from critical (terse; shows only the most-urgent messages) to debug (verbose). Use the show logging levels command to list all of the log components and the current logging level for each.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show logging levels [`*`component`*`]`**

component (optional) specifies that you only want the logging level for the given component.

Guidelines The log components are listed in an earlier section; see Log Components, on page 38-3. This command shows the logging levels in a graph: for each log component, an X appears in the column of its logging level. Use the logging level command to change any or all component-logging levels.

Samples `bstnA(cfg)# show logging levels`

shows the logging levels for all components. See Figure 38.7, below.

`bstnA(cfg)# show logging levels POLICY`

shows the logging levels for one component, POLICY. See Figure 38.7, below.

Related Commands logging level

Figure 38.7 *Sample Output: show logging levels*

```
bstnA(cfg)# show logging levels


Configured Log Levels:
                    disable  debug   info   notice  warning  error critical
AFN_NET             |       |       |       |    X  |       |       |      |
APP_RON_RTMD        |       |       |       |    X  |       |       |      |
AUTH_CONSOLE        |       |       |       |    X  |       |       |      |
AUTH_HTTP           |       |       |       |    X  |       |       |      |
AUTH_HTTPS          |       |       |       |    X  |       |       |      |
AUTH_SSH            |       |       |       |    X  |       |       |      |
AUTH_TELNET         |       |       |       |    X  |       |       |      |
BOOTPD              |       |       |       |    X  |       |       |      |
CHASS               |       |       |       |    X  |       |       |      |
CHASS_LIB           |       |       |       |    X  |       |       |      |
CIFSCLI             |       |       |       |    X  |       |       |      |
CIFSCONFD           |       |       |       |    X  |       |       |      |
CIFS_INIT           |       |       |       |    X  |       |       |      |
CIFS_SHIM_LIB       |       |       |       |    X  |       |       |      |
CLI                 |       |       |       |    X  |       |       |      |
CLI_WIZARD          |       |       |       |    X  |       |       |      |
...
```

Figure 38.8 *Sample Output: show logging levels POLICY*

```
bstnA(cfg)# show logging levels POLICY


Configured Log Levels:
                    disable  debug   info   notice  warning  error critical
POLICY              |       |       |       |    X  |       |       |      |

bstnA(cfg)# ...
```

show stats-monitor

Purpose The stats-monitor process keeps round-trip-time (RTT) and error statistics for communication between the ARX and external devices, as well as internal performance statistics. It also analyzes this data, watches for sudden changes in RTT or error rates, and notifies you of these changes. Use the show stats-monitor command to see the current administrative state of the stats monitor.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show stats-monitor`

Guidelines The output contains a configuration table, followed by a table of configured notifications. The notification table only appears if at least one notification is set up with the notify command.

Statistics Monitor Configuration is a table with the high-level configuration for the stats monitor. You can edit this configuration with the stats-monitor command and its sub commands.

Status is always Enabled.

Sampling Interval is the time between the data samples (such as average RTTs and error rates) recorded in the stats-log. The stats monitor records these samples in a series of .csv files. Each .csv file represents one group of external devices (such as back-end CIFS filers or NFS clients for a particular NFS service), or an internal process (such as file migrations or an internal CIFS work queue). You can use show stats-logs to view these .csv files. The default sampling interval is adequate for most sites, but you can use the sampling interval command to change it if necessary.

Max Data Size is the storage limit for all statistics (.csv) files stored in the stats-logs partition.

Max Raw Data Size is the percentage of Max Data Size permitted for .csv files with "raw" (not "hourly") in their names.

Used Data Size and

Used Raw Data Size are the amounts of the above storage that are already used. These are for all statistics (hourly and raw) and raw statistics, respectively.

Notification Parameters only appears if at least one notification is active; use the notify command to enable analysis and notification for front-end services or back-end filers. This shows each notification in one row of the table, with the following fields:

Applies To identifies the type of service, process, or device that is configured for notifications. You set this with the notify command. This is :

- "CIFS Services" (related to a cifs front-end service and its clients; these statistics are similar to the CIFS output from show statistics global server),
- "NFS Services" (for an nfs service and its clients, similar to the NFS output from show statistics global server),
- "CIFS Shares" or "NFS Shares" (back-end shares behind an ARX volume, viewable with show statistics namespace ... summary on an individual share),

Guidelines (Cont.)

Statistics shows the type of notification. This is Errors for a notification of an unusual number of errors, or RespTime for a notification of an increase in response time. You can use the moving-average command to enable or disable notifications for this type of event.

Traps indicates whether or not the stats monitor sends an SNMP trap with this notification. This is Enabled or Disabled, and you can use the trap command to change the setting.

Metric is always Mov.Avg.

Parameters shows the threshold settings that trigger a notification. You can use the moving-average command to change these thresholds.

Sample `bstnA# show stats-monitor`

shows the configuration of the stats monitor. See Figure 38.9, below.

Related Commands stats-monitor
sampling interval
show stats-logs

Figure 38.9 *Sample Output: show stats-monitor*

```
bstnA# show stats-monitor

Statistics Monitor Configuration
--------------------------------
  Status               Enabled
  Sampling Interval    300 seconds
  Max Data Size        500 MB
  Max Raw Data Size    80% of Max
  Used Data Size       327 kB
  Used Raw Data Size   274 kB

  Notification Parameters
  -----------------------

  Applies To       Statistics    Traps        Metric    Parameters
  ----------       ----------    -----        ------    ----------
  CIFS Services    Errors        Disabled     Mov.Avg   >100%  for 6 Samples

  CIFS Services    RespTime      Disabled     Mov.Avg   >100%  for 6 Samples

  NFS Services     Errors        Disabled     Mov.Avg   >100%  for 6 Samples

  NFS Services     RespTime      Disabled     Mov.Avg   >100%  for 6 Samples

  CIFS Shares      Errors        Disabled     Mov.Avg   >100%  for 6 Samples

  CIFS Shares      RespTime      Disabled     Mov.Avg   >100%  for 6 Samples

  NFS Shares       Errors        Disabled     Mov.Avg   >100%  for 6 Samples

  NFS Shares       RespTime      Disabled     Mov.Avg   >100%  for 6 Samples
```

show system tasks

Purpose Use the show system tasks command to see a list of subsystem tasks running on the ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show system tasks`**

Guidelines

Proc - is the module slot (number) and processor on which this task is running. This column only appears for platforms with more than one physical processor, such as the ARX-4000.

Subsystem - is the name of the subsystem task. If the task is a script, that fact is noted in parentheses: "(script)."

Instance - is how many instances of the same task are running on this processor. This column does not appear for the ARX-1500 or ARX-2500.

Memory (Kb) - is the amount of memory used by this task.

CPU% - measures the CPU cycles that the task is consuming. This percentage is measured since the last invocation of show system tasks. For a measurement of overall CPU usage, use show processors and/or show processors usage.

CPU-Time - only appears on an ARX-1500 or an ARX-2500. This shows the CPU time used by the current subsystem task since the process last started. The time is expressed in *HH*:*MM*:*SS* format.

Samples

`bstnA>` **`show system tasks`**

shows all subsystem tasks. See Figure 38.10 for sample output.

`stoweA>` **`show system tasks`**

shows all subsystem tasks running on an ARX-2500 named "stoweA." Figure 38.11 on page 38-52 shows sample output.

Related Commands

show processors
show processors usage
show logging levels

Figure 38.10 *Sample Output: show system tasks*

```
bstnA> show system tasks

 Proc  Subsystem                  Instance  Memory (Kb)   CPU%
 ----  -------------------------  --------  ------------  ----
 1.1   afnEmailHomed              1         6384          0
 1.1   afnEximMgrd                1         5640          0
 1.1   afnFTd                     1         11404         0
 1.1   afnFfpMgrd                 1         4884          0
 1.1   afnSlmConfd                1         11288         0
 1.1   afncrmd                    1         5268          0
```

```
 1.1   afnipceventmgrd              1         5876          0
 1.1   afnloggermgrd                1         5936          0
 1.1   afnmerged                    1         4824          0
 1.1   afnnetd                      1         6304          0
 1.1   afnpdpd                      1         6112          0
 1.1   afnrootd                     1         8420          0
 1.1   afnsnapd                     1         14112         0
 1.1   afnsqmd                      1         9412          0
 1.1   afnssbdevd                   1         5376          0
 1.1   afnssrmd                     1         5884          0
 1.1   afnstatsd                    1         5336          0
 1.1   bespd                        1         29512         0
 1.1   bootpd                       1         4780          0
 1.1   chassnew                     1         5460          0
 1.1   cifsconfd                    1         6504          0
 1.1   cifsconfd                    1         6528          0
 1.1   cli_daemon                   1         46200         0
 1.1   coreCollector (script)       1         1740          0
 1.1   coreTimer (script)           1         1680          0
...
 1.1   rameccReporter (script)      1         1340          0
 1.1   relmgr                       1         5476          0
 1.1   rtmd                         1         5760          0
 1.1   scmcluster                   1         3968          0
 1.1   securityd                    1         5768          0
 1.1   shmemctld                    1         3668          0
 1.1   shmemd                       1         3888          0
 1.1   spfd                         1         5464          0
 1.1   sqmd_sa_wdog                 1         2520          0
 1.1   subshrmgtd                   1         8684          0
 1.1   svc_agent                    2         8140          0
 1.1   svcmgr                       1         7188          0
 1.1   syslog-ng                    1         1088          0
 1.1   syslog-ng_guard              1         3760          0
 1.1   tcmd                         1         6184          0
 1.1   tcored                       1         4880          0
 1.1   temipcproxyd                 1         4980          0
 1.1   temstatsd                    1         6668          0
 1.1   txnmond                      1         7444          0
 1.1   xsdd                         1         7244          0
 1.1   Other Tasks                  34        77772         3
```

Figure 38.11 *Sample Output: show system tasks (ARX-2500)*

```
stoweA> show system tasks

Subsystem                    Memory (Kb)   CPU%    CPU-Time
---------------------------  ------------  ----    ---------
afnEmailHomed                6768          0       00:00:00
afnEximMgrd                  5824          0       00:00:00
afnFTd                       7436          0       00:00:00
afnFcnd                      24348         0       00:00:00
afnSlmConfd                  12452         0       00:00:00
afncrmd                      4988          0       00:00:00
afnipceventmgrd              4640          0       00:00:00
afnlicd                      12400         0       00:00:00
afnloggermgrd                5776          0       00:00:00
afnnetd                      6212          0       00:00:00
...
```

stats-monitor

Purpose

**** This is a beta-level command. It does not appear in the CLI unless you use the* terminal beta *command to unlock all beta-level commands. ****

The ARX records the round-trip times and errors between its own software and other external devices (such as back-end filers and front-end clients), as well as time required for migrations and internal operations. You can use the stats-monitor command to begin configuring an analysis process that keeps a running average of this data and monitors it for sudden changes. Whenever an issue is discovered, the stats-monitor process logs a notification of the issue. Optionally, it also sends an SNMP trap and raises and alarm that is visible with the show health command.

You can use the stats-monitor command to enter gbl-stmon mode and edit the analysis parameters.

Mode

gbl

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
stats-monitor
```

Default(s)

None

Guidelines

This is a beta-level command. Use the terminal beta command to unlock all beta-level commands, including this one.

This command puts you into gbl-stmon mode, where you can use commands to change the stats-monitor configuration. The stats monitor takes periodic samples of errors, computation times, and round-trip times for analysis; you can use the sampling interval command to change the time between these samples. To change the set of external devices that are analyzed (such as front-end client machines and/or back-end filers), you can use the notify command. The notify command also leads to a submode, where you can enable traps for the group of devices and change the thresholds for the device set's notifications.

The stats monitor records its notifications in the syslog; you can use show logs syslog to view them along with all other syslog messages. Look for the log component named "SMON_ANALYSIS," or use grep to search for it.

The raw monitoring data from the stats monitor is kept in .csv files, which you can list and examine with the show stats-logs command.

Sample

```
bstnA(gbl)# stats-monitor
bstnA(gbl-stmon)#
```

enters the configuration mode for monitoring statistics. From here, you can change parameters for the stats-monitor process.

Related Commands

sampling interval
notify
show stats-logs
show stats-monitor
show health

trap

Purpose

**** This is a beta-level command. It does not appear in the CLI unless you use the* terminal beta *command to unlock all beta-level commands. ****

The ARX records the round-trip times between itself and external devices, such as CIFS clients and back-end filers. The stats-monitor process keeps running averages of these round-trip times and other statistics, and raises alarms in the syslog if the number of errors from an external device increases dramatically, or if the response time from an external host slows dramatically. Use the gbl-stmon-nfy trap command to trigger an SNMP trap for these alarms.

Use no trap to stop sending an SNMP trap for these alarm conditions.

Mode

gbl-stmon-nfy

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
trap {response-time | errors}
no trap {response-time | errors}
```

response-time | errors is a required choice. This determines the type of error that triggers an SNMP trap: slow response times or excessive errors.

Default(s)

```
no trap response-time
no trap errors
```

Guidelines

This is a beta-level command. Use the terminal beta command to unlock all beta-level commands, including this one.

This command enables traps for one set of filers or ARX front-end services, chosen when you use the notify command to enter the mode.

Refer to the ARX SNMP Reference for a full list of all traps, along with explanations for each trap and procedures to address the issue. Follow the *Guidelines* for the command, snmp-server traps, to set up a remote host to receive SNMP traps. To manually clear a raised trap, you can use the clear health command.

To send these traps in E-mail to a set of interested recipients, you can use the email-event command and its sub-commands to create an e-mail event group.

Sample

```
bstnA(gbl-stmon-nfy)# trap response-time
bstnA(gbl-stmon-nfy)# trap errors
```

causes the stats monitor to send SNMP traps for any slow response times or excessive errors.

Related Commands

notify
snmp-server traps
clear health
email-event -> group ... event or group (cfg-email-event) stats-monitor

39

Network Troubleshooting Tools

Network-Log Components

A *Network-log component* is a source of fastpath-log messages, typically an internal process or group of processes responsible for network traffic. On the original ARX platform, these processes all resided on a single module called the Network Services Module, or NSM. The table below is an alphabetical list of all Network-log components, with a brief description of each.

Log Component	Description
NSM_BSD	Board Support Package.
NSM_CACHE	NSM Cache.
NSM_CIFS	CIFS Proxy on the data plane.
NSM_CM	Connection Manager.
NSM_DGRAM	Datagram.
NSM_DME	Data Mover Engine (DME) proxy on the data plane. This migrates files from one back-end filer to another.
NSM_DMS	NSM DMS.
NSM_DT	Direct Transport.
NSM_ECHO	Packet Echo Service.
NSM_GWMON	Gateway Monitor.
NSM_HA	High-Availability.
NSM_HOT	Hot files.
NSM_IPA	AIPC Server and Client, for internal communication between network processes.
NSM_LIB	Networking Library Functions.
NSM_LOG	System Logging.
NSM_LOOP	Packet Loop Service.
NSM_MSG	Inter-NSM Messaging.
NSM_NAT	The Start of tem message categories.
NSM_NAT	NAT Subsystem.
NSM_NET_HEALTH	Network Health Checks.
NSM_NFS	NFS Proxy on the data plane.
NSM_NFS_CLIENT	NSM NFS client.
NSM_NFS_DNAS	The NFS proxy interactions with DNAS (becoming obsolete).

Log Component	Description
NSM_NFS_FC	NFS File Cache Management Proxy.
NSM_NFS_IPC	The NFS Proxy IPC (obsolete?).
NSM_NFS_MOUNT	Obsolete.
NSM_NFS_SERVER	NFS file server state machine.
NSM_NFS_STATE	The NFS proxy state machine (becoming obsolete).
NSM_NLM	NFS Lock Manager (NLM) proxy on the data plane.
NSM_PING	PING Application.
NSM_PORT	Port Management.
NSM_PRESTO	Presto Console Log.
NSM_RON	RON Fastpath Proxy.
NSM_SMB2	CIFS Proxy.
NSM_SNTP	SNTP.
NSM_STATS	Statistics from the data plane.
NSM_TEME	Fastpath environment-kernel.
NSM_TFTP	TFTP Client.
NSM_TTCP	TTCP Service.
NSM_VIP	VIP Processing.
NSM_VIP_FENCE	VIP Fencing, which can occur when an ARX volume is creating a coordinated snapshot. The VIP fence, while it is raised, prevents clients from changing any volume state while its back-end filers are each taking their snapshots.
NSM_XID_MAP	NFS RPC XID Mapping.

capture merge

Purpose You can use the capture session command to capture IP traffic and store it in multiple files, where each file is closed and the next is opened as it reaches a size limit. The packet-capture uses the same format as WireShark, an open-source packet analyzer. When you end the capture session, the ARX merges these files into a single file by default. You have the option to end the session without merging the files, or the process may not be able to merge them if the session is somehow interrupted. To merge some capture files in one of these cases, use the capture merge command.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`capture merge file`** ***`prefix`***

prefix (1-256 characters) is the common prefix of the files to be merged. Use show capture to see a list of available capture files; choose a prefix that is common to two or more of them from the same capture session. The full file names are "*prefix*[*id_timestamp*].cap," where the *id* and the *timestamp* differentiate one file from the other.

Default(s) None.

Guidelines This command merges multiple output files from a single capture session. This only works when the session is finished; use show capture sessions to see any or all currently-active capture sessions.

The merge makes a single file in chronological order, with the earliest captured packet at the top of the file. To analyze the file from the CLI, you can use show capture *filename* to see the full contents or show capture *filename* summary to see a summary.

Use grep to search through the file for a string. You can also use collect, copy ftp, or a similar copy command to copy a capture file to another machine. You can use WireShark, Tshark, tcpdump, or some other packet analyzer to analyze the file on another host, or you can send it back to F5 Support for their analysis.

Sample

```
bstnA# capture merge file proxyTraffic

% INFO: Merging 2 capture files into proxyTraffic.cap; this may take up
to 1 minute to complete.
```

merges all packet-capture files with the prefix, "proxyTraffic."

Related Commands

show capture sessions
show capture
grep
collect
copy ftp
copy scp
copy {nfs|cifs}
copy smtp
copy tftp

capture session

Purpose This command captures IP traffic and stores it in a file. The packet-capture uses the same format as WireShark, an open-source packet analyzer.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
capture session session-id ip ip-address [and-ip ip2]
  vlan vlan-id file prefix [filesize kilobytes]
  [filecount count] [protocol {cifs | non-cifs}]

no capture session session-id [no-merge]
```

session-id (1-4; limit of 2 on the ARX-VE or ARX-500) is a unique ID for the session, to be used for stopping the capture in the no form of the command.

ip-address is the address to match against. This selects the IP packets to capture; any packet with a matching source or destination IP address is included in the capture, in chronological order.

and-ip *ip2* (optional) adds a second address to the filter. If you enter this, the capture includes all bidirectional traffic between *ip-address* and *ip2*.

vlan *vlan-id* (1-4095) focuses on traffic over the specified VLAN.

file *prefix* (1-255 characters) is the prefix you choose for the output file. The full file name is "*prefix*[*id_timestamp*].cap," where the *id* and the *timestamp* are only used if the session creates multiple capture files. The file(s) go to the "capture" directory: use show capture to see the file listing, or show capture *file-name* to view the file.

filesize *kilobytes* (optional, 1-50,000; 1-1,000,000 for the ARX-4000) truncates the capture at the specified file size. One kilobyte is 1000 bytes, not 1024.

filecount *count* (optional, 1-10) limits the number of capture files. Each file is no larger than the filesize. If you set this count to 2 or more, the capture process rotates the capture files indefinitely. With multiple, rotating capture files, you can use no capture session *session-id* to stop capturing packets.

protocol {cifs | non-cifs} (optional) filters the captured packets. If you choose **cifs**, the capture file only includes packets to or from CIFS-related ports: UDP/88, TCP/88, UDP/137, UDP/138, TCP/139, and/or TCP/445. If you choose **non-cifs**, the capture file includes packets to or from any *other* ports.

no-merge (optional, for the no form of the command) only applies to a capture session where the ***count*** is two or more files. Without this, the no form of the command merges all files from a multi-file session. This option prevents the merge.

Syntax: Capture All Proxy-IP Traffic

```
capture session session-id proxy-all file prefix [filesize
kilobytes] [filecount count] [protocol {cifs | non-cifs}]
```

proxy-all captures all traffic with any proxy-IP address as its source or destination. Use show ip proxy-addresses to show all proxy-IP addresses on the switch. No VLAN ID is required for this syntax; the VLAN is implicit. The ARX-VE only has a single proxy-IP address, so it does not support this option.

Default(s)

filesize - 16,000 kilobytes

filecount - 1

Guidelines

This command starts a capture session, similar to the monitor command. Unlike the monitor command, this does not require a network analyzer.

The preparation for a packet capture can be time-consuming, especially on a busy system. The CLI prints a message after you enter this command, warning of the possible delay.

To see any or all currently-active capture sessions, use show capture sessions. Use the show capture command to view the all capture files in a directory listing, or to view a capture file's contents. Use grep to search through the file for a string. The tail ... follow command displays the capture file as it grows. You can also use collect, copy ftp, or another copy command to copy a capture file to another machine, where you can examine it with WireShark, TShark, tcpdump, or some other packet analyzer.

The no capture command immediately stops the current capture. If you use the no-merge option with a multi-file capture, the session creates multiple capture files. Multiple capture files also result when a multi-file session is interrupted by an ARX reboot. To merge all the files from a particular capture session into a single file, use the capture merge command.

Jumbo frames are not included in any capture files. The jumbo mtu command enables jumbo frames. If you use jumbo frames, you can use the monitor command with a network analyzer to capture your network traffic. The monitor command requires an ARX with multiple network interfaces.

Samples

```
bstnA# capture session 1 ip 192.168.25.19 vlan 25 file clientCap
```

captures all traffic to and from a particular IP address. The file name is "clientCap."

```
bstnA# capture session 2 proxy-all file proxyTraffic filesize 150
filecount 2
```

captures all traffic to or from the proxy-IP address(es). This capture session is limited to two files of 150 Kbytes each. The capture session continues indefinitely, rotating the two files, until someone uses no capture session 2 to stop it.

```
bstnA# no capture session 2
```

immediately stops capture session 2. This command merges the two files produced by the capture. The resulting file is named "proxyTraffic.cap."

Related Commands

show capture sessions
show capture
capture merge
monitor
tail
grep
copy ftp
copy scp
copy {nfs|cifs}
copy smtp
copy tftp

clear statistics filer connections

Purpose Use this command to clear the connection statistics for a given back-end filer, or for all of them.

Mode priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax

```
clear statistics filer connections
clear statistics filer ext-filer-name connections
```

ext-filer-name (1-64 characters) specifies a single filer where you want to clear all connection statistics. If you omit this, the command clears the connection statistics for all filers. This identifies the external filer by its configured name on the ARX. Use show external-filer to display all configured external filers.

Default(s) None

Guidelines The show statistics filer connections command shows current connection counts between the ARX's software and its back-end filers. Use this command to clear those connection counters for one filer, or all of them.

The CLI prompts for confirmation before clearing the statistics; enter **yes** to proceed.

Sample

```
bstnA# clear statistics filer fs2

Clear the connection statistics on filer 'fs2'.

Are you sure? [yes/no] yes
```

clears the connection statistics for the external filer named "fs2."

Related Commands show statistics filer connections
show external-filer

clear statistics global server

Purpose Use this command to clear the client-traffic statistics for a global server.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear statistics global server`** ***`fqdn`***

fqdn (1-64 characters) is the fully-qualified domain name (for example, "myserver.organization.org") for a global server. Use show global server to see a list of global servers.

Default(s) None

Guidelines The show statistics global server command shows the high-level traffic between a global server and its clients. Use this command to clear the traffic counters for one global server.

The CLI prompts for confirmation before clearing the statistics; enter **yes** to proceed.

Sample

```
bstnA# clear statistics global server ac1.medarch.org
Clear statistics of global server ac1.medarch.org.

Are you sure? [yes/no] yes
```

Related Commands show statistics global server
show global server

drop filer-connections

Purpose Use the drop filer-connections command to drop all connections to a back-end filer or server. This command may produce a noticeable interruption for your clients; use it only for a filer that is overwhelmed with ARX connections.

Modes priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `drop filer-connections` ***`filer`*** `[processor` ***`slot.processor`***`]`

filer (1-64 characters) identifies the filer to be disconnected (for example, "smb-1"). The show external-filer command lists all filers defined on the ARX.

slot.processor (for example, 2.7) focuses on one network (or fastpath, or NSM) slot and processor. If you use this option, only the identified network processor drops its connections to the filer.

Default(s) None

Guidelines The CLI prompts for confirmation before dropping any connections; enter **yes** to proceed.

This is useful for a Tier-2 filer that is currently overwhelmed with TCP traffic, or is experiencing connectivity issues that may be related to a high number of TCP connections. To set a limit on the number of CIFS connections from the ARX to this filer, use the cifs connection-limit command. If you set a lower limit than the current one, you have the option to wait for CIFS clients to disconnect gracefully. This causes the ARX to block any new CIFS connections to the filer until enough currently-connected clients drop off. You can use this command to reverse that decision, immediately dropping all connections to the filer and going to the lower limit.

You can use nfs tcp connections to set a limit on the number of NFS/TCP connections to the filer.

You can use the show filer connections command to monitor the current connections to a filer.

◆ Note

The NSM re-establishes its NFS connections as soon as they are dropped. Similarly, CIFS-client applications may re-establish connections as soon as they are disconnected. Therefore, this command may not have any noticeable affect on the show filer connections *output.*

To drop a CIFS-client connection, you can use drop cifs-service user-session.

Sample

```
bstnA# drop filer-connections smb1 processor 2.6
Drop the connections from processor 2.6 to filer smb1.

Proceed? [yes/no] yes
bstnA#
```

disconnects all CIFS/TCP connections between NSM processor 2.6 and the "smb1" filer.

Related Commands

show filer connections
drop cifs-service user-session
cifs connection-limit

expect nslookup

Purpose Use this command to perform an **nslookup** from the ARX, thereby testing the DNS lookup configuration.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `expect nslookup` ***ip-or-hostname*** `[timeout` ***seconds***`]`

ip-or-hostname (1-128 characters) is the IP address or host name to look up (for example, "172.16.46.2," "myserver3," or "juser-pc.myco.com").

seconds (optional, 1-2096) sets a time limit on the nslookup operation. If you omit this and the command takes excessive time, you can use <Ctrl-C> to stop it.

Default(s) ***seconds*** - 21,600 (6 hours)

Guidelines This command calls a remote DNS server to look up an IP address or machine name. The output is the IP address and fully-qualified domain name (FQDN) for the given host.

The lookup fails if none of the ARX's DNS servers are reachable. Use the show ip domain command to show all configured DNS servers. To add a new one, use the ip name-server command.

Sample

```
bstnA# expect nslookup bboard.wwmed.com

bboard.wwmed.com is an alias for rh1.wwmed.com.
rh1.wwmed.com has address 192.168.25.19
```

looks up the host named "bboard.wwmed.com."

Related Commands show ip domain
ip name-server

expect show firewall

Purpose The firewall configuration determines which IP packets are accepted into the control-plane processor (processor 1.1) and which packets are dropped. Use expect show firewall to show the current firewall configuration.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`expect show firewall [timeout seconds]`**

seconds (optional, 1-2096) sets a time limit on the show-firewall operation. If you omit this and the command takes excessive time, you can use <Ctrl-C> to stop it.

Default(s) ***seconds*** - 21,600 (6 hours)

Guidelines These firewall rules apply to out-of-band management (MGMT) port. This is on the ACM or control plane.

This command shows three ordered lists of firewall rules: INPUT (packets coming in), FORWARD (packets to be forwarded to another station), and OUTPUT (packets to be sent out). Every packet is compared to one of these lists depending on its path (for example, all incoming packets are compared to the INPUT list).

Within each list, the packet is compared to the rules in order. Each row is one rule:

pkts is the number of packets accepted by the rule.

bytes is the total number of bytes from the above packets.

target (typically ACCEPT or DROP) determines what to do with a matching packet.

The remaining flags are match criteria.

Use expect show ifconfig to show configurations for all interfaces (internal and external) on the control-plane processor.

Sample `bstnA# expect show firewall`
shows the firewall rules for the ACM processor. See Figure 39.1 for sample output.

Related Commands expect show ifconfig

Figure 39.1 *Sample Output: expect show firewall*

```
bstnA# expect show firewall

Chain INPUT (policy DROP 528 packets, 61360 bytes)
 pkts bytes target     prot opt in     out     source               destination
    0     0 DROP       tcp  --  *      *       0.0.0.0/0            0.0.0.0/0           tcp
flags:0x03/0x03
 4169  505K ACCEPT     0    --  lo     *       0.0.0.0/0            127.0.0.1           state
NEW,RELATED,ESTABLISHED
16187 1244K OOB_IN     0    --  *      *       0.0.0.0/0            10.1.1.7
47035 8393K IB_IN      0    --  *      *       0.0.0.0/0            169.254.70.32
 152K   49M IB_IN      0    --  *      *       0.0.0.0/0            169.254.70.33
   34 11152 ACCEPT     udp  --  *      *       0.0.0.0/0            0.0.0.0/0           udp dpt:67
```

```
    0     0 ACCEPT     udp  --  *      *       0.0.0.0/0            0.0.0.0/0           udp dpt:68
1495K  504M ACCEPT     0    --  *      *       0.0.0.0/0            169.254.70.0/24

Chain FORWARD (policy DROP 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source               destination
    0     0 ACCEPT     0    --  *      lo      0.0.0.0/0            127.0.0.1           state
NEW,RELATED,ESTABLISHED

Chain OUTPUT (policy DROP 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source               destination
 4169  505K ACCEPT     0    --  *      lo      127.0.0.1            0.0.0.0/0           state
NEW,RELATED,ESTABLISHED
34724 7411K OOB_OUT    0    --  *      *       10.1.1.7          0.0.0.0/0
49214   18M IB_OUT     0    --  *      *       169.254.70.32        0.0.0.0/0
 167K   42M IB_OUT     0    --  *      *       169.254.70.33        0.0.0.0/0
    0     0 ACCEPT     udp  --  *      *       0.0.0.0/0            0.0.0.0/0           udp spt:67
    0     0 ACCEPT     udp  --  *      *       0.0.0.0/0            0.0.0.0/0           udp spt:68
1461K  494M ACCEPT     0    --  *      *       169.254.70.0/24      0.0.0.0/0

Chain IB_IN (2 references)
 pkts bytes target     prot opt in     out     source               destination
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0            0.0.0.0/0           tcp dpt:443
    0     0 ACCEPT     udp  --  *      *       0.0.0.0/0            0.0.0.0/0           udp dpt:161
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0            0.0.0.0/0           tcp dpt:22
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0            0.0.0.0/0           tcp dpt:23
 199K   58M MGMT_IN    0    --  *      *       0.0.0.0/0            0.0.0.0/0

Chain IB_OUT (2 references)
 pkts bytes target     prot opt in     out     source               destination
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0            0.0.0.0/0           tcp spt:443
    0     0 ACCEPT     udp  --  *      *       0.0.0.0/0            0.0.0.0/0           udp spt:161
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0            0.0.0.0/0           tcp spt:22
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0            0.0.0.0/0           tcp spt:23
 216K   60M MGMT_OUT   0    --  *      *       0.0.0.0/0            0.0.0.0/0

Chain MGMT_IN (2 references)
 pkts bytes target     prot opt in     out     source               destination
 1671  159K ACCEPT     icmp --  *      *       0.0.0.0/0            0.0.0.0/0           icmp !type
13
 198K   58M ACCEPT     0    --  *      *       0.0.0.0/0            0.0.0.0/0           state
RELATED,ESTABLISHED
    8   440 DROP       0    --  *      *       0.0.0.0/0            0.0.0.0/0

Chain MGMT_OUT (2 references)
 pkts bytes target     prot opt in     out     source               destination
 1283 89824 ACCEPT     icmp --  *      *       0.0.0.0/0            0.0.0.0/0           icmp !type
14
 221K   61M ACCEPT     0    --  *      *       0.0.0.0/0            0.0.0.0/0           state
NEW,RELATED,ESTABLISHED
    0     0 ACCEPT     udp  --  *      *       0.0.0.0/0            0.0.0.0/0           udp dpt:162
    0     0 DROP       0    --  *      *       0.0.0.0/0            0.0.0.0/0

Chain OOB_IN (1 references)
 pkts bytes target     prot opt in     out     source               destination
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0            0.0.0.0/0           tcp dpt:443
    0     0 ACCEPT     udp  --  *      *       0.0.0.0/0            0.0.0.0/0           udp dpt:161
15370 1134K ACCEPT     tcp  --  *      *       0.0.0.0/0            0.0.0.0/0           tcp dpt:22
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0            0.0.0.0/0           tcp dpt:23
  817  109K MGMT_IN    0    --  *      *       0.0.0.0/0            0.0.0.0/0

Chain OOB_OUT (1 references)
 pkts bytes target     prot opt in     out     source               destination
```

```
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp spt:443
    0     0 ACCEPT     udp  --  *      *       0.0.0.0/0           0.0.0.0/0           udp spt:161
28375 6068K ACCEPT     tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp spt:22
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp spt:23
 6349 1344K MGMT_OUT   0    --  *      *       0.0.0.0/0           0.0.0.0/0
```

expect show ifconfig

Purpose Use expect show ifconfig to show configuration details for all ACM-processor interfaces.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`expect show ifconfig [timeout seconds]`**

seconds (optional, 1-2096) sets a time limit on this command. Whether or not you set this timeout, you can use <Ctrl-C> to stop the process at any time.

Default(s) ***seconds*** - 21,600 (6 hours)

Guidelines These interfaces are on the ACM.

The output is labeled with interface names in the left column. Next to each interface name are the configuration details for the interface. Most of these interfaces are internal to the ARX, used for communication with other internal processors.

The interface names are in the following format: *interface*.*vlan*:*num*. The *interface* is the name of the interface (such as eth0), the *vlan* is the VLAN number, and the *num* is different for every IP alias that is used by the *interface*.*vlan*.

Use expect show firewall to show the firewall rules used by the ACM processor.

Sample `prtlndA# expect show ifconfig`
shows the interface configurations for the ACM processor on "prtlndA." See Figure 39.2 on page 39-16 for sample output.

Related Commands expect show firewall

Figure 39.2 *Sample Output: expect show ifconfig*

```
prtlndA# expect show ifconfig

bond0     Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
          RX packets:954160 errors:0 dropped:0 overruns:0 frame:0
          TX packets:907659 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:670060549 (639.0 MiB)  TX bytes:698558518 (666.1 MiB)

bond0.100 Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          inet addr:169.254.104.1  Bcast:169.254.104.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
          RX packets:954179 errors:0 dropped:0 overruns:0 frame:0
          TX packets:907680 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:649072348 (619.0 MiB)  TX bytes:691289093 (659.2 MiB)

bond0.100 Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          inet addr:169.254.104.32  Bcast:169.254.104.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
```

```
bond0.100 Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          inet addr:169.254.104.33  Bcast:169.254.104.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1

bond0.100 Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          inet addr:169.254.104.68  Bcast:169.254.104.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1

bond0.100 Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          inet addr:169.254.104.69  Bcast:169.254.104.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1

bond0.100 Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          inet addr:169.254.104.70  Bcast:169.254.104.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1

bond0.100 Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          inet addr:169.254.104.71  Bcast:169.254.104.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1

eth0      Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
          RX packets:498428 errors:0 dropped:0 overruns:0 frame:0
          TX packets:633677 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:261440790 (249.3 MiB)  TX bytes:433161035 (413.0 MiB)
          Base address:0x4020 Memory:b1260000-b1280000

eth1      Link encap:Ethernet  HWaddr 00:15:17:6B:C3:2D
          inet addr:10.1.23.11  Bcast:10.59.255.255  Mask:255.255.0.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:8193 errors:0 dropped:0 overruns:0 frame:0
          TX packets:7777 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:693807 (677.5 KiB)  TX bytes:1808206 (1.7 MiB)
          Base address:0x5000 Memory:b0800000-b0820000

eth2      Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
          RX packets:339809 errors:0 dropped:0 overruns:0 frame:0
          TX packets:187641 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:391236999 (373.1 MiB)  TX bytes:207820730 (198.1 MiB)
          Base address:0x4000 Memory:b1220000-b1240000

eth3      Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
          RX packets:50203 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1148 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:4316117 (4.1 MiB)  TX bytes:179668 (175.4 KiB)
          Base address:0x3020 Memory:b0f60000-b0f80000

eth4      Link encap:Ethernet  HWaddr 00:0A:49:17:B1:09
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
          RX packets:65720 errors:0 dropped:0 overruns:0 frame:0
          TX packets:85193 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:13066643 (12.4 MiB)  TX bytes:57397085 (54.7 MiB)
          Base address:0x3000 Memory:b0f20000-b0f40000

lo        Link encap:Local Loopback
```

```
inet addr:127.0.0.1  Mask:255.0.0.0
UP LOOPBACK RUNNING  MTU:16436  Metric:1
RX packets:351539 errors:0 dropped:0 overruns:0 frame:0
TX packets:351539 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:83063134 (79.2 MiB)  TX bytes:83063134 (79.2 MiB)
```

expect show netstat

Purpose Use expect show netstat to show the current network status at the ACM processor.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`expect show netstat [timeout`** ***`seconds`*****`]`**

seconds (optional, 1-2096) sets a time limit on this command. Whether or not you set this timeout, you can use <Ctrl-C> to stop the process at any time.

Default(s) ***seconds*** - 21,600 (6 hours)

Guidelines Use expect show ifconfig to show configurations for all interfaces (internal and external) on the ACM processor. The expect show firewall command shows the processor's firewall rules, which determine whether to drop or accept IP packets.

Sample `bstnA# expect show netstat`
shows the network status for the ACM processor. See Figure 39.3 for sample output.

Related Commands expect show ifconfig
expect show firewall

Figure 39.3 *Sample Output: expect show netstat*

```
bstnA# expect show netstat

Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address           Foreign Address         State
tcp        0      0 169.254.70.1:35264      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.78:44160     0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.1:43104      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.1:32800      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.1:46080      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.80:33185     0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.1:33153      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.83:4450      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.84:2050      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.83:2050      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.80:4450      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.79:2050      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.78:2050      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.77:2050      0.0.0.0:*               LISTEN
tcp        0      0 169.254.70.76:2050      0.0.0.0:*               LISTEN
...
```

expect traceroute

Purpose Use this command to show all IP-router hops from the ACM processor to a given IP address.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax `expect traceroute ip-destination [timeout seconds]`

ip-destination (1-128 characters) is the destination-IP address (for example, 172.16.46.2).

seconds (optional, 1-2096) sets a time limit on this command. Whether or not you set this timeout, you can use <Ctrl-C> to stop the process at any time.

Default(s) ***seconds*** - 21,600 (6 hours)

Guidelines This command sends a packet from the ACM processor and shows each IP router that forwards the packet. The packet starts at the MGMT interface (use show interface mgmt to find this address) or an inband (VLAN) management interface (show interface vlan).

If the packet does not reach its next hop, it prints asterisks (*) to mark the time. Use **<Ctrl-C>** to escape the command.

Use the ping command to see if an address is reachable from various processors. Use expect show ifconfig to show configurations for all interfaces (internal and external) on the ACM processor. The expect show firewall command shows the processor's firewall rules, which determine whether to drop or accept IP packets.

Sample

```
bstnA# expect traceroute 192.168.25.19
```

```
 1  192.168.25.1  0.587 ms  0.553 ms  0.491 ms
 2  192.168.25.19  0.367 ms  0.363 ms  0.367 ms
```

shows that there are two hops from the ACM processor to the machine at 192.168.25.19.

Related Commands
ping
expect show ifconfig
expect show firewall

expect ttcp

Purpose Test TCP (TTCP) is an open-source benchmarking tool for testing TCP (and UDP) connections. Use this command to run a TTCP test between two ARXes (over a RON tunnel) or between an ARX and a remote TTCP server.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
expect ttcp server [timeout seconds]
expect ttcp transmit ttcp-server-ip [timeout seconds]
```

seconds (optional, 1-2096) sets a time limit on this command. Whether or not you set this timeout, you can use <Ctrl-C> to stop the process at any time.

ttcp-server-ip (1-128 characters) is the IP address of a TTCP server (for example, 172.16.46.2). If the TTCP server is another ARX on the RON, use the .1 address on the switch's private subnet (for example, 169.25.100.1). Use show ron route to find the private subnet.

Default(s) ***seconds*** - 21,600 (6 hours)

Guidelines To test the throughput between two ARXes, you must first prepare the receiving switch to serve TTCP. Run the expect ttcp server command on the receiving switch. This blocks the CLI until the optional **timeout** expires, or until you press <Ctrl-C>. Then go to the sending switch to start the TTCP test.

The expect ttcp transmit command invokes a 10-second TTCP test to any server that supports TTCP. This can be any filer, client, or other station on the network that can receive TTCP transmissions, or it can be another ARX on the same Resilient Overlay Network (RON; see interface ron). You must identify the server first, as described above.

The transmission can hang for longer than one minute if the IP address is incorrect. You can use the optional **timeout** or <Ctrl-C> to cancel the test.

The transmit test ends with a one-line report containing the following fields:

nnnn.nnn MB / 10.00 Sec = xxx.xxxx Mbps is the amount of data transmitted in 10 seconds.

aa %TX is the percent-CPU utilization at the transmitting switch's ACM processor.

bb %RX is the percent-CPU utilization reported from the TTCP server.

When a switch-to-switch test is over, you can go back to the server switch and press <Ctrl-C>.

The ARX can only receive TTCP from another ARX in the same RON.

Use the ping command to see if an address is reachable from various processors. Use expect show ifconfig to show configurations for all interfaces (internal and external) on the ACM processor. The expect show firewall command shows the processor's firewall rules, which determine whether to drop or accept IP packets.

Sample

```
prtlndA# expect ttcp server timeout 600
Waiting for TTCP client. Hit ^C to abort...
```

uses the “prtlndA” switch as the TTCP server. It blocks the CLI for up to 600 seconds (10 minutes), waiting for a transmission from another ARX.

This takes place on the transmitting switch, named “bstnA:”

```
bstnA# show ron route

Default Policy
--------------

  Destination            Subnet           via Tunnel        Milliseconds
  ------------------------------------------------------------------------
  gffstnA            169.254.104.0/24     toGoffstown            0.1
  minturnA            169.254.80.0/24     toMinturn              0.1
  provA              169.254.127.0/26     toProvidence           0.1
  prtlndA             169.254.66.0/24     toPortland             0.1
  prtlndB             169.254.96.0/24     toPortlandB            0.1

bstnA# expect ttcp transmit 169.254.66.1

Starting 10 second TCP transmit test. Hit ^C to abort...
  586.5138 MB /  10.00 sec =  491.9229 Mbps 12 %TX 23 %RX

bstnA#
```

Shows the private subnet for the “prtlndA” switch, then uses the subnet to send a TTCP test to the server.

To unblock the CLI on the server switch, “prtlndA,” you can wait for the timeout or press <Ctrl-C>:

```
Waiting for TTCP client. Hit ^C to abort... <Ctrl-C>
prtlndA#
```

Related Commands

ping
expect show ifconfig
expect show firewall

logging fastpath component

Purpose Each group of network-log messages, known as an network-log *component*, has a separately-tunable *logging level*. Level 1 is the most terse, displaying non-recoverable errors only; level 10 is the most verbose, displaying all levels of messages including per-packet logs. The messages appear in a log file named "fastpath." Use the logging fastpath component command to set a network-logging level.

Use the no form of this command to disable network logging for a component.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
logging fastpath component nsm-component level
no logging fastpath component nsm-component
```

nsm-component (1-128 characters) is the network component to tune. See Network-Log Components, on page 39-3 for a complete list of network components.

level (0-10) sets the logging level for the component. Level 0 disables all logs from the component. Level 1, as mentioned above, logs only non-recoverable errors. Levels 2 and 3 include warnings and recoverable errors. Level 4 adds logs about internal-configuration changes. Levels 5-10 include per-packet logs, where level 10 is the most verbose.

Important

Any level greater than 1 can significantly degrade networking performance on the ARX. Use this feature only under the strict guidance of F5 personnel.

Default(s) ***level*** - 1

Guidelines The first time you enter any logging fastpath command, the CLI issues a warning about the performance impact of network (also called "NSM") logging. If this appears, enter **yes** to proceed.

This is similar to the logging level command, which sets logging levels for software components on the or ACM. This affects the logging level on all network processors where logging is enabled; use the logging fastpath processor command to enable logging on a processor. The show fastpath logging command shows which network processors are currently enabled for logging, along with the log-level settings for each component.

If you set the logging level higher than zero, you can filter the logs for a given IP address, host name, or any other search string. This reduces the size of the fastpath file and focuses the logs on a particular problem. Use the logging fastpath component ... filter command to implement filtering for an network-log component.

From any mode, use show logs fastpath or grep pattern logs fastpath to view the log messages in the fastpath file.

Samples

```
bstnA(cfg)# logging fastpath component NSM_CIFS 6

Enabling fastpath logging adversely impacts the performance of the
switch.

Proceed? [yes/no] yes
```

sets the NSM_CIFS component to log level 6. On every network processor where logging is enabled, the NSM_CIFS component starts logging messages up to and including level-6 messages. This can be verbose on a busy system.

```
bstnA(cfg)# no logging fastpath component NSM_CIFS
```

disables log messages from the NSM_CIFS component.

Related Commands

logging fastpath processor
logging fastpath component ... filter
show fastpath logging

logging fastpath component ... filter

Purpose

Each source of network-log messages, known as an network-log *component*, can filter its messages before adding them to the log file. This reduces the stress on network processors and can help with diagnosing a network problem. Use the logging fastpath component ... filter command to filter the log messages from a particular network component.

Use the no form of this command to remove one match string from the log filter.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

`logging fastpath component` ***`nsm-component`*** **`filter`** ***`match-string`*** **`{include | exclude}`**

`no logging fastpath component` ***`nsm-component`*** **`filter`** ***`match-string`***

nsm-component (1-128 characters) is the network component to filter. See Network-Log Components, on page 39-3 for a complete list of network components.

match-string (optional, 1-80 characters) is a string to match against. Quote the string if it contains any spaces. Any messages that match this string are added to the "fastpath" log file. To include log messages that match multiple strings (for example, multiple IP addresses), repeat this command with each desired string. If a log message matches *any* of the entered strings, the logging component adds it to the "fastpath" file.

include | **exclude** is a required choice. The **include** choice causes the filter to include any network-log messages that match the *match-string*. The **exclude** option reverses the filter; a message is *excluded* from the log if it matches the *match-string*.

Default(s)

no filter

Guidelines

The first time you enter any logging fastpath command, the CLI issues a warning about the performance impact of network logging. If this appears, enter **yes** to proceed.

Use logging fastpath processor to activate logging for one or more network processors, then use the logging fastpath component command to set the logging level for an network-log component. This command filters the messages from the component; the filter is ineffective for any component(s) where the logging level is 0 (zero).

Use show fastpath logging to verify the filter settings for each component.

From any mode, use show logs fastpath or grep pattern logs fastpath to view the log messages in the fastpath file.

Samples

```
bstnA(cfg)# logging fastpath component NSM_CIFS filter 172.16.22.100
include
bstnA(cfg)# logging fastpath component NSM_CIFS filter 192.168.25.31
include
bstnA(cfg)# logging fastpath component NSM_CIFS filter 192.168.25.32
include
```

filters the NSM_CIFS component for three IP addresses. Any NSM_CIFS log message that matches any of these IP addresses is included in the fastpath file.

```
bstnA(cfg)# no logging fastpath component NSM_CIFS filter 192.168.25.32
```

removes one IP address from the NSM_CIFS filter.

```
bstnA(cfg)# logging fastpath component NSM_VIP filter 192.168.25.10
exclude
bstnA(cfg)# logging fastpath component NSM_VIP filter 192.168.25.12
exclude
```

filters the log messages from the NSM_VIP component. This excludes all log messages that match either “192.168.25.10” or “192.168.25.12;” all other messages from the component are included.

Related Commands

logging fastpath processor
logging fastpath component
show fastpath logging

logging fastpath processor

Purpose

A network processor can send log messages to the "fastpath" log file for a detailed diagnosis of networking issues. Use the logging fastpath processor command to activate logging for one network processor.

Use the no form of this command to stop the processor from adding any log messages to the file.

Mode

cfg

Security Role(s)

network-engineer or crypto-officer

Syntax

```
logging fastpath processor slot.processor
no logging fastpath processor slot.processor
```

slot.processor (for example, 1.4, 2.7, or 3.1) identifies a slot and network processor. Use the show processors command for a full list of processors on the ARX.

Default(s)

no logging

Platforms

ARX-500, ARX-2000, and ARX-4000

Guidelines

The first time you enter any logging fastpath command, the CLI issues a warning about the performance impact of network logging. If this appears, enter **yes** to proceed.

Several network *components* run on a network processor, where each component can generate its own set of log messages in the "fastpath" log file. The network components are listed earlier in the chapter; see Network-Log Components, on page 39-3. Each component has a tunable logging level that you can set with the logging fastpath component command. By default, all of the logging components are set at a very terse logging level, so as not to overwhelm the network processor(s) that have logging enabled. If you raise the logging level to diagnose a problem, you should limit the volume of logging messages with the logging fastpath component ... filter command.

Important

Too much logging may overwhelm a network processor, possibly resulting in a service interruption. Only use network logging under the guidance of F5 personnel.

Use show fastpath logging to show which processors are enabled for logging. This also shows all configured logging levels and filter settings.

From any mode, use show logs fastpath or grep pattern logs fastpath to view the network-log messages in the fastpath file.

Samples

```
bstnA(cfg)# logging fastpath processor 3.1

Enabling fastpath logging adversely impacts the performance of the
switch.

Proceed? [yes/no] yes
```

enables logging for processor 1 in slot 3.

```
bstnA(cfg)# no logging fastpath processor 4.4
```

disables logging for processor 4 in slot 4.

Related Commands

logging fastpath component
logging fastpath component ... filter
show fastpath logging
show logs
grep

monitor

Purpose Use the monitor command to configure port mirroring. *Port mirroring* mirrors the Ethernet traffic on one or more ports onto another port, where the destination port typically has a network analyzer attached. The network analyzer can therefore see all traffic going through the source interface(s) in real time.

Important

A service interruption can occur if the destination interface is over subscribed. Use this command only as advised by F5 Technical Support.

Use the no form of the command to stop port mirroring.

Mode cfg

Security Role(s) network-engineer or crypto-officer

Syntax

```
monitor {module | diagnostic}
  source-interface slot/port {rx | tx | both}
  destination-interface slot/port

no monitor {module | diagnostic}
  source-interface slot/port {rx | tx | both}
  destination-interface slot/port
```

module | diagnostic is a required choice.

module enables you to mirror from one external port to another on the same module.

diagnostic is for internal use only.

source-interface *slot/port* (2/1-14 on ARX-4000; 1/1-12 on ARX-2000) identifies a port to be monitored (for example, 2/3 or 1/6).

rx | tx | both is a required choice, where you specify the direction(s) of the packets to monitor (rx = received packets; tx = transmitted packets; both = both received and transmitted packets).

destination-interface *slot/port* (same ranges as for **source-interface**) identifies an external port where the network analyzer is located; for example, 2/1. Choose a destination port with equal or greater bandwidth than the source port.

Default(s) None

Platforms ARX-2000 or ARX-4000

Guidelines

Before you use this command, you typically connect a network analyzer at the destination port.

The monitor module command may overwhelm the destination interface, causing client/server packets to be dropped. The CLI prompts you with a warning before mirroring any packets; enter **yes** to continue.

Use show interface summary to show all ports and their slots.

You can issue this command more than once to mirror traffic from more than one port. If there is more traffic from the source port than the destination port can carry, frames are dropped; there is no buffering at the destination port.

Remember to shut down a monitor session when you finish diagnosing the problem. Port mirroring causes a heavy load on the monitored port, which must duplicate every frame and forward it to the network analyzer's port. Use the show monitor command to see if there is an active monitor session.

Samples

```
bstnA(cfg)# monitor module source-interface 2/6 both
destination-interface 2/2
Warning: Oversubscription of the destination monitor port may result in
a service interruption. Proceed only upon advice from Technical
Support.

Are you sure? [yes/no] yes
```

monitors port 2/6. A network analyzer at port 2/2 will see all packets, both transmit and receive, going through port 2/6.

```
bstnA(cfg)# no monitor module
```

stops monitoring on the port.

Related Commands

show chassis
show monitor
show interface **summary**

ping

Purpose Use the ping command to send one or more pings (ICMP ECHO requests) to a specified IP address.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
ping destination-ip [count number] [framesize bytes]
ping destination-ip from slot.processor [count number]
  [framesize bytes]
ping destination-ip source source-ip [count number]
  [framesize bytes]
ping destination-ip from slot.processor source source-ip
  [count number] [framesize bytes]
```

destination-ip is the IP address to receive the ping.

count *number* (optional, 1-10,000) limits the number of pings to send.

from *slot.processor* (optional) is not available on the ARX-1500, ARX-2500, or ARX-VE. This identifies a slot and processor to send the ping (for example, 2.3). Use the show processors command for a full list of processors on the ARX. You can ping from processor 1.1 to test the out-of-band (MGMT) network processor and connection.

source *source-ip* (optional) is the source-IP address to send in the ICMP ECHO request. If you omit the **from** clause, the CLI chooses an appropriate processor. The ICMP ECHO response will be returned to this address; if you set the source processor with **from** and the IP does not reside on that processor, the ping output shows "no response."

framesize *bytes* (optional) is the size of the packet that you want to send. This is an extra payload that the software adds to the ICMP ECHO header.

Default(s) **count *number*** - infinite; use **<Ctrl-C>** to stop the pings.

from *slot.processor* -

if you choose a ***source-ip***: the processor that is associated with the source IP.

without a ***source-ip***: the first available processor. This starts with the first network processor (for example, 1.2 in an ARX-2000). If no network processors are available, the management processor (1.1) sends the ping.

source *source-ip* - the best source IP to reach the *destination-ip*, as chosen by the sending processor's routing table.

framesize *bytes* - 0 (zero)

Guidelines

If you do not use the count clause to limit the number of pings, use **<Ctrl-C>** to stop them.

For a list of proxy-IP addresses, which network processors use to communicate with filers and servers, use the show ip proxy-addresses command. The output maps each proxy-IP address to a particular processor. For a list of virtual-IP addresses, which processors use to communicate with the client side, use the show global server command. The show interface vlan command produces a list of in-band (VLAN) management interfaces, and the show interface mgmt command shows the IP for the out-of-band Mgmt interface on the front panel. Use the show ip route command to view the IP routing tables for all processors on the switch.

The output is typical for ping implementations. The last column shows the source processor used for the ping, in *slot.processor* format, for platforms that support a specific processor source. If you specified only the ***source-ip*** in the command, the last column shows the source address instead. A summary of all pings appears at the end of the output; see *Samples*, below.

You can use the expect traceroute command to show all the IP-router hops from the ACM processor to any given IP address.

Samples

```
bstnA> ping 172.16.100.83
PING 172.16.100.83 (172.16.100.83) 0 data bytes
8 bytes from 172.16.100.83: icmp_seq=1 ttl=62 time=3 ms. from 2.3
8 bytes from 172.16.100.83: icmp_seq=2 ttl=62 time=<1 ms. from 2.3
8 bytes from 172.16.100.83: icmp_seq=3 ttl=62 time=<1 ms. from 2.3
8 bytes from 172.16.100.83: icmp_seq=4 ttl=62 time=<1 ms. from 2.3
8 bytes from 172.16.100.83: icmp_seq=5 ttl=62 time=<1 ms. from 2.3
8 bytes from 172.16.100.83: icmp_seq=6 ttl=62 time=<1 ms. from 2.3
8 bytes from 172.16.100.83: icmp_seq=7 ttl=62 time=<1 ms. from 2.3
8 bytes from 172.16.100.83: icmp_seq=8 ttl=62 time=<1 ms. from 2.3
8 bytes from 172.16.100.83: icmp_seq=9 ttl=62 time=<1 ms. from 2.3
8 bytes from 172.16.100.83: icmp_seq=10 ttl=62 time=<1 ms. from 2.3
8 bytes from 172.16.100.83: icmp_seq=11 ttl=62 time=<1 ms. from 2.3

<Ctrl-C>

-------172.16.100.83 ping statistics
11 packets transmitted, 11 packets received, 0% packet loss
round-trip min/avg/max 0/0/3 ms
bstnA>
```

sends several pings.

```
bstnA> ping 10.53.2.10 count 4
```

sends 4 pings.

```
bstnA> ping 10.1.1.1 source 10.1.1.7 count 4
```

sends 4 pings from source IP 10.1.1.7 (the out-of-band management IP) to 10.1.1.1 (the out-of-band management gateway).

Related Commands

show processors
show ip proxy-addresses
show interface vlan
show interface mgmt
show ip route
show global server
expect traceroute

show capture sessions

Purpose You can invoke an IP capture session that collects certain IP packets and streams them into a file. Use the show capture sessions command to view all such sessions.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show capture sessions`

Guidelines You can capture IP traffic into a file with the capture session command. This command shows a list of all currently-running sessions.

The output is split into two tables. Each session appears on its own row. The first table shows the choices made with the capture session command. It has the following columns:

Session identifies the session by its ID.

Ip is the address to match against, if any. Any packet or source or destination address matches this is included in the capture file. This is "proxy-all" if the session is capturing traffic from all proxy-IP addresses at once.

Additional Ip is another address to match against, if any. Any packet exchanged between this address and the address above is included in the capture. As above, "proxy-all" indicates that the session is capturing traffic to all proxy-IP addresses at once.

VLAN is the VLAN ID that is being scanned for matching IP packets. If this is 1, a frame with no specified VLAN ID can also match.

File Size is the maximum size of the file, if any. You can set this with an option in the capture session command.

File Count is the maximum number of files that the session can produce.

The second table shows the current state of the each capture and the prefix for its output files:

Session identifies the session by its ID.

State is either "Capturing" or "Complete."

File Name is the prefix of all the output files. Use the show capture command to show a list of all capture files, or to show the contents of any of them. To maintain this directory, use the copy, rename, grep, and delete commands.

Sample `bstnA> show capture sessions`

shows all active capture sessions. See Figure 39.4 on page 39-35 for sample output.

Related Commands capture session
show capture

Figure 39.4 *Sample Output: show capture sessions*

```
bstnA> show capture sessions

Session Ip               Additional Ip    VLAN File Size  File Count
------- ---------------- ---------------- ---- ---------- ----------
1       192.168.25.21    192.168.25.49    25   16000      3
2       192.168.25.27    192.168.25.27    25   16000      n/a


Session State      File Name
------- ---------- ------------------------------
1       Capturing  ntaps.cap
2       Capturing  fsrvr.cap
```

show fastpath logging

Purpose You can enable one or more network processors to send log messages to a "fastpath" log file. The network-log messages are divided into groups based on the software components that generate them. You can set a logging level for each network-software component, from 1 (terse; shows only the most-urgent messages) to 10 (extremely verbose). You can also filter each logging component so that it logs messages that are relevant to a particular network issue. Use the show fastpath logging command to list all network processors that are configured for logging along with the log levels and filtering for each network component.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show fastpath logging`

Guidelines The output contains two tables. The first shows the network processors that have logging enabled:

Slot is the slot number , and

Processor is the processor that is actively logging its messages. Use the logging fastpath processor command to enable logging on a network processor.

The second table shows all network-log components with a non-default logging level and/or filtering enabled:

Component is an network-log component with non-default settings. A full list of network-log components appears earlier; see Network-Log Components, on page 39-3.

Trace Level is the logging level, where 1 is the most-terse level and 10 is most-verbose. Use the logging fastpath component command to change this.

Filter Type is "inclusive" or "exclusive." This determines whether a log message containing the Filter String is included in or excluded from the "fastpath" log. This is set with the logging fastpath component ... filter command.

Filter String is a match string that the network component applies to each log message. This is also set with the logging fastpath component ... filter command.

Sample `bstnA(cfg)# show fastpath logging`

shows the logging configuration for the network (or "fastpath") processes. See Figure 39.5, below.

Related Commands logging fastpath processor
logging fastpath component
logging fastpath component ... filter

Figure 39.5 *Sample Output: show fastpath logging*

```
bstnA(cfg)# show fastpath logging
   Slot         Processor
-------------------------

    2             1
```

```
    2              2
    2              3


        Component      Trace Level     Filter Type    Filter String
-------------------------------------------------------------------

       NSM_CIFS            6
       NSM_CIFS                       inclusive      192.168.25.15
       NSM_CIFS                       inclusive      172.16.100.183
```

show fastpath resources

Purpose Use the show fastpath resources command for a detailed breakdown of software resources used on a network processor.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show fastpath resources [`*`slot.processor`*` | ip-address `*`ip-address`*`]`**

slot.processor (optional: 2.1-12 on ARX-4000; 1.2-5 on ARX-2000; 1.2 on ARX-500 or ARX-VE) focuses the output on a single network processor (for example, 2.4). If you omit this, the report shows the resources used on all of the ARX's network processors. Use the show processors command for a full list of processors on the ARX.

ip-address (optional) focuses the output on an IP address. Each IP address is handled by a particular network processor; this finds the correct processor and reports on the resources used there.

Guidelines If you enter an IP address, the top of the output shows the network processor that serves the address. You can also use show ip address to find the processor assigned to the address.

The output contains the following tables for each network processor:

- NFS resources,
- CIFS resources,
- System resources,
- Presto resources, and
- IP ports.

This data is difficult to interpret without knowledge of the network software's internal processes. Contact F5 Support for further information.

Samples

```
bstnA(cfg)# show fastpath resources 2.1
```

shows the resources used for a single NSM processor. See Figure 39.6, below.

```
bstnA(cfg)# show fastpath resources ip-address 192.168.25.23
```

shows the resources used on the NSM processor for 192.168.25.23. See Figure 39.7 on page 39-41.

Related Commands show ip address
show processors

Figure 39.6 *Sample Output: show fastpath resources 2.1*

```
bstnA(cfg)# show fastpath resources 2.1

Resources for Processor:  2.1

   NFS Resource                         Current     Maximum     Percent Used
-------------------------------------------------------------------------------
NFS Active Transactions                     0         37888          0%
NFS Shares                                 14         16384          0%
NFS File Handle Cache (kb)                359         30310          1%
NFS Client Connections                     57          2368          2%
NFS Global Service                         29           947          3%
NFS File Servers                            8          3788          0%
NFS Networking                              1          1894          0%
NFS Clients                                 0          4736          0%
XID Range Maps                             30         65519          0%
XID Save Requests                           0          4096          0%
NFS File Cache Resource (kb)               78          6630          1%
NFS File Cache Replication (kb)           427         35993          1%
DNAS Instances                              3           473          1%


  CIFS Resource                         Current     Maximum     Percent Used
-------------------------------------------------------------------------------
CIFS Transaction Blocks                     0         23680          0%
CIFS Shares                                19         16384          0%
CIFS File Handle Cache (kb)                 0         13824          0%
CIFS Connections                            4          9472          0%
CIFS Sessions                               0         11840          0%
CIFS Request ID                             0             0          0%
CIFS DNSS Objects                           4          2048          0%
CIFS Service Objects                        6           512          1%
CIFS Exports Objects                       19         16000          0%
CIFS Tree Connect Contexts                  0         47360          0%
CIFS File Info Blocks                       0         73728          0%
CIFS Find File Info Blocks                  0         94720          0%
CIFS Attachment Points                      6         65536          0%
CIFS TID/UID management                     0           640          0%
CIFS Pathcache Blocks                       0         65536          0%
CIFS BOSTONMED Session Blocks                   0         47360          0%
CIFS Backend File/Dir Opens                 0         30310          0%


 System Resource                        Current     Maximum     Percent Used
-------------------------------------------------------------------------------
File Cache Transaction                      0          3788          0%
NAT Rules                                 217          7577          3%
NAT Actions                                74        242483          0%
NAT HA                                      0        242483          0%
HA Requests                                 0            92          0%
AIPC Client Requests                        0           236          0%
AIPC Events                                 0           236          0%
Syslog Handles                              9            59         15%
RPC Layer                                  70         11840          1%
RON Tunnels                                 0           236          0%
RON Connections                             0            14          0%
Memory Owners                             112          3788          3%
Reactors                                   20            92         22%
Messages                                    0            59          0%
MAC addresses                               6           512          1%
Network Interface Stats                     8           473          2%
Intra-DP Message                            0        121241          0%
```

```
TFTP Client                                0            29          0%
Radix Tables                              25         60620          0%
Connections                              137         11840          1%
User Heap (kb)                          3195       2714621          0%

 Presto Resource                       Current     Maximum     Percent Used
-------------------------------------------------------------------------------
Presto temp routing layer (kb)             3           512          1%
Presto packets (kb)                        0         26343          0%
Presto memory (kb)                     27127        182267         15%
Presto TCP TCB Pool (kb)                  54         16384          0%
Presto UDP Pool (kb)                       3           192          2%
Presto sockets (kb)                        4          1088          0%
Presto IP Fragments (kb)                   0           256          0%

 IP Ports Resource                     Current       Maximum    Percent Used
-------------------------------------------------------------------------------
169.254.80.2 TCP Ports                    12         58368          0%
169.254.80.2 TCP Priv Ports                0           319          0%
169.254.80.2 UDP Ports                     1         58368          0%
169.254.80.2 UDP Priv Ports                0           319          0%
169.254.80.31 TCP Ports                    0         58368          0%
169.254.80.31 TCP Priv Ports               0           319          0%
169.254.80.31 UDP Ports                    0         58368          0%
169.254.80.31 UDP Priv Ports               0           319          0%
192.168.25.5 TCP Ports                     0           512          0%
192.168.25.5 TCP Priv Ports                0           319          0%
192.168.25.5 UDP Ports                     0           512          0%
192.168.25.5 UDP Priv Ports                0           319          0%
10.46.11.253 TCP Ports                     0           512          0%
10.46.11.253 TCP Priv Ports                0           319          0%
10.46.11.253 UDP Ports                     0           512          0%
10.46.11.253 UDP Priv Ports                0           319          0%
192.168.25.31 TCP Ports                    6         58368          0%
192.168.25.31 TCP Priv Ports               0           319          0%
192.168.25.31 UDP Ports                    0         58368          0%
192.168.25.31 UDP Priv Ports               1           319          0%
169.254.80.64 TCP Ports                   28         58368          0%
169.254.80.64 TCP Priv Ports               3           319          1%
169.254.80.64 UDP Ports                   25         58368          0%
169.254.80.64 UDP Priv Ports               0           319          0%
192.168.25.12 TCP Ports                    0          2432          0%
192.168.25.12 TCP Priv Ports               0           319          0%
192.168.25.12 UDP Ports                    0          2432          0%
192.168.25.12 UDP Priv Ports               0           319          0%
192.168.25.15 TCP Ports                    0          2432          0%
192.168.25.15 TCP Priv Ports               0           319          0%
192.168.25.15 UDP Ports                    0          2432          0%
192.168.25.15 UDP Priv Ports               0           319          0%
```

Figure 39.7 *Sample Output: show fastpath resources ip-address 192.168.25.23*

```
bstnA(cfg)# show fastpath resources ip-address 192.168.25.23

Resources for Processor:  2.6

   NFS Resource                        Current     Maximum     Percent Used
-------------------------------------------------------------------------------
NFS Active Transactions                    0         37888          0%
NFS Shares                                14         16384          0%
NFS File Handle Cache (kb)               359         30310          1%
NFS Client Connections                    53          2368          2%
NFS Global Service                        29           947          3%
NFS File Servers                           8          3788          0%
NFS Networking                             1          1894          0%
NFS Clients                                0          4736          0%
XID Range Maps                             8         65519          0%
XID Save Requests                          0          4096          0%
NFS File Cache Resource (kb)              78          6630          1%
NFS File Cache Replication (kb)          427         35993          1%
DNAS Instances                             3           473          1%


  CIFS Resource                        Current     Maximum     Percent Used
-------------------------------------------------------------------------------
CIFS Transaction Blocks                    0         23680          0%
CIFS Shares                               19         16384          0%
CIFS File Handle Cache (kb)                0         13824          0%
CIFS Connections                           4          9472          0%
CIFS Sessions                              0         11840          0%
CIFS Request ID                            0             0          0%
CIFS DNSS Objects                          4          2048          0%
CIFS Service Objects                       6           512          1%
CIFS Exports Objects                      19         16000          0%
CIFS Tree Connect Contexts                 0         47360          0%
CIFS File Info Blocks                      0         73728          0%
CIFS Find File Info Blocks                 0         94720          0%
CIFS Attachment Points                     6         65536          0%
CIFS TID/UID management                    0           640          0%
CIFS Pathcache Blocks                      0         65536          0%
CIFS BOSTONMED Session Blocks                  0         47360          0%
CIFS Backend File/Dir Opens                0         30310          0%


 System Resource                      Current     Maximum     Percent Used
-------------------------------------------------------------------------------
File Cache Transaction                     0          3788          0%
NAT Rules                                217          7577          3%
NAT Actions                               24        242483          0%
NAT HA                                     0        242483          0%
HA Requests                                0            92          0%
AIPC Client Requests                       0           236          0%
AIPC Events                                0           236          0%
Syslog Handles                             9            59         15%
RPC Layer                                 61         11840          1%
RON Tunnels                                0           236          0%
RON Connections                            0            14          0%
Memory Owners                            107          3788          3%
Reactors                                  19            92         21%
Messages                                   0            59          0%
MAC addresses                              5           512          1%
```

```
Network Interface Stats                    7          473          1%
Intra-DP Message                           0       121241          0%
TFTP Client                                0           29          0%
Radix Tables                              21        60620          0%
Connections                              128        11840          1%
User Heap (kb)                          3185      2714621          0%

 Presto Resource                        Current     Maximum     Percent Used
-------------------------------------------------------------------------------
Presto temp routing layer (kb)             2          512          1%
Presto packets (kb)                        0        26343          0%
Presto memory (kb)                     24416       182267         13%
Presto TCP TCB Pool (kb)                  50        16384          0%
Presto UDP Pool (kb)                       3          192          2%
Presto sockets (kb)                        4         1088          0%
Presto IP Fragments (kb)                   0          256          0%

 IP Ports Resource                      Current       Maximum    Percent Used
-------------------------------------------------------------------------------
169.254.80.7 TCP Ports                    12         58368          0%
169.254.80.7 TCP Priv Ports                0           319          0%
169.254.80.7 UDP Ports                     1         58368          0%
169.254.80.7 UDP Priv Ports                0           319          0%
192.168.25.5 TCP Ports                     0           512          0%
192.168.25.5 TCP Priv Ports                0           319          0%
192.168.25.5 UDP Ports                     0           512          0%
192.168.25.5 UDP Priv Ports                0           319          0%
10.46.11.253 TCP Ports                     0           512          0%
10.46.11.253 TCP Priv Ports                0           319          0%
10.46.11.253 UDP Ports                     0           512          0%
10.46.11.253 UDP Priv Ports                0           319          0%
192.168.25.142 TCP Ports                   4         58368          0%
192.168.25.142 TCP Priv Ports              0           319          0%
192.168.25.142 UDP Ports                   0         58368          0%
192.168.25.142 UDP Priv Ports              1           319          0%
169.254.80.69 TCP Ports                   28         58368          0%
169.254.80.69 TCP Priv Ports               3           319          1%
169.254.80.69 UDP Ports                   25         58368          0%
169.254.80.69 UDP Priv Ports               0           319          0%
192.168.25.12 TCP Ports                    0          2432          0%
192.168.25.12 TCP Priv Ports               0           319          0%
192.168.25.12 UDP Ports                    0          2432          0%
192.168.25.12 UDP Priv Ports               0           319          0%
192.168.25.15 TCP Ports                    0          2432          0%
192.168.25.15 TCP Priv Ports               0           319          0%
192.168.25.15 UDP Ports                    0          2432          0%
192.168.25.15 UDP Priv Ports               0           319          0%
```

show filer connections

Purpose Use the show filer connections command to show the current TCP and UDP connections from the network processors to a given back-end filer.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show filer connections {`*`ext-filer-name`*` | ip-addr `*`ip`*`}`**

ext-filer-name (optional, 1-64 characters) identifies the external filer by its configured name. Use show external-filer to display all configured external filers.

ip is the IP address of the filer.

Guidelines The output is a table with one line per connection to the filer.

Proc is the network processor at one end of the connection, in *slot.processor* format.

Proxy Ip is proxy-IP address that the network processor is using. Use show ip proxy-addresses to show all proxy-IP addresses on the switch.

Filer Port shows the transport protocol (TCP or UDP) and the port number being used at the filer.

Conn. is the number of connections from the network processor to the filer.

You can use the drop filer-connections command to drop all connections to a filer. The NSM re-establishes its TCP connections immediately, and CIFS-client applications may do the same; therefore, this output may not change after you run the drop filer-connections command.

The show statistics filer connections command shows statistics on filer connections over time.

Sample `bstnA(cfg)# show filer connections ip-addr 192.168.25.19`

shows connections to the filer at 192.168.25.19. See Figure 39.8 for sample output.

Related Commands
drop filer-connections
show statistics filer connections
show ip proxy-addresses

Figure 39.8 *Sample Output: show filer connections ip-addr*

```
bstnA(cfg)# show filer connections ip-addr 192.168.25.19

Connections to external filer :  192.168.25.19

Proc Proxy Ip        Filer Port Conn.
---- --------------- ---------- -----
2.1  192.168.25.31    UDP/0      1
2.2  192.168.25.32    UDP/0      1
2.3  192.168.25.33    UDP/0      1
2.4  192.168.25.34    UDP/0      1
2.5  192.168.25.141    UDP/0      1
```

```
2.6  192.168.25.142    UDP/0      1
2.7  192.168.25.143    UDP/0      1
2.8  192.168.25.144    UDP/0      1
2.9  192.168.25.145    UDP/0      1
2.10 192.168.25.146    UDP/0      1
2.11 192.168.25.147    UDP/0      1
2.12 192.168.25.148    UDP/0      1
```

show monitor

Purpose Use the show monitor command to show the configuration for any active monitoring session.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show monitor`**

Platforms ARX-2000 and ARX-4000

Guidelines A monitor session, also called port mirroring, duplicates frames from one port to another. You can connect a network analyzer to the destination port to examine the traffic on the source port. Use the monitor command to start or stop a port-monitoring session. Remember to stop a monitoring session if it is not in use; the duplication is a performance strain for both the source and destination ports.

Sample

```
bstnA(cfg)# show monitor

Monitor Session: System
Source : Slot 2 Port 9
Destination :Slot 2 Port 12
Mode: Rx
```

shows the only currently active monitoring session.

Related Commands monitor

show statistics filer connections

Purpose Use the show statistics filer connections command to show filer-connection statistics for a particular filer, or for all filers.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show statistics filer connections
show statistics filer filer connections [processor slot.proc]]
```

filer (optional, 1-64 characters) identifies a particular filer (for example, "nas10"). The show external-filer command lists all filers defined on the ARX. If you omit this, the output contains a table with summary statistics for all connected filers.

slot.proc (optional: 2.1-12 on ARX-4000; 1.2-5 on ARX-2000; 1.2 on ARX-500 or ARX-VE) focuses on one NSM slot and processor. If you use this option, the output focuses on connections from the identified NSM processor.

Guidelines: Summary Output The output from the simplest syntax, show statistics filer connections, shows a table of current connections and maximum connections for all external filers. Each filer appears on one row with the following fields:

Filer is the name of the external filer, as defined on the ARX.

Connections

Current is the current number of connections to the filer. You can use cifs connection-limit to set a maximum number of CIFS connections, and you can use nfs tcp connections to set a ceiling on NFS/TCP connections.

Max is the maximum number of simultaneous connections to the filer since the last reboot, or since the last time someone ran the clear statistics filer connections command.

Guidelines: Detailed Output

If you specify a filer in the command, the output contains two tables with detailed statistics. The CIFS table contains the following fields:

Connection limit is the maximum number of connections allowed to this filer, if any. You can use the cifs connection-limit command to set this limit.

Current data connections is the current number of connections through the data plane. That is, this is the number of connections that are directly related to client connections, and do not come from the control plane (processes on the ACM).

Current control connections counts the current connections from processes on the ACM. These are processes that do not directly relate to any client connections, such as connections from the policy engine or connections that require calculations at the control plane.

Max connections (data+control) is the most simultaneous connections to this filer since the last reboot, or since the last time someone ran clear statistics filer connections.

Time of Max connections is the date that the above maximum occurred.

Current data sessions shows the number of data sessions that are currently connected to the filer. Each data connection can hold multiple data sessions, where each session may correspond to a different client. This is the number of sessions that are directly related to client activity, and do not come from the control plane (processes on the ACM).

Max data sessions is the most simultaneous sessions on this filer since the last reboot, or since the last time someone ran clear statistics filer connections. As above, these are sessions that did not go through the control plane.

Time of Max data sessions is the date that the above maximum occurred.

Max sessions per data connection shows the most CIFS sessions to simultaneously run over a single TCP connection. This applies to data-plane connections between the ARX and an external client or filer.

Time of Max sessions/connections is the date that the above maximum occurred.

The NFS table contains the following fields:

Current connections counts the current NFS connections to the filer.

Max connections counts the maximum NFS connections to the filer since the last reboot, or since the last time someone ran clear statistics filer connections.

Time of Max connections is the date that the above maximum occurred.

Guidelines: Detailed Output for One Processor

If you specify a particular filer and processor, the output shows connection statistics for the traffic between them. These statistics are similar to the ones above, but with a more limited scope:

Data Connection limit is the maximum number of connections allowed to this filer from this processor, if any. You can use the cifs connection-limit command to set the system-wide limit for CIFS connections.

Current data connections is the current number of connections between the processor and the filer. That is, this is the number of connections that are directly related to client connections, and do not come from the control plane (processes on the ACM).

Current control connections counts the current processor/filer connections that terminate at the ACM. The ACM processes do not necessarily relate to any client connections; they are connections from the policy engine or connections that require calculations at the control plane.

Current data sessions shows the number of data sessions that are currently connected to the filer through this processor. Each data connection can hold multiple data sessions, where each session may correspond to a different client. This is the number of sessions that are directly related to client activity, and do not come from the control plane processes on the ACM.

Max sessions per data connection shows the most CIFS sessions to simultaneously run over a single TCP connection. This applies to data-plane connections between the selected processor and filer.

Time of Max sessions/connections is the date that the above maximum occurred.

The NFS table contains a single field:

Current connections counts the current NFS connections between the network processor and the filer.

Guidelines: Other Statistics

You can use the show filer connections command to examine the current connections to a given filer. For client-side statistics, use show statistics global server. For raw, read/write statistics from the NSM fastpath, use show statistics namespace ... fastpath.

Samples

```
bstnA(cfg)# show statistics filer connections
```

sums up all connections to all back-end filers. See Figure 39.9 for sample output.

```
bstnA(cfg)# show statistics filer fs4 connections
```

shows details on the connections and sessions to the fs4 filer. See Figure 39.10 on page 39-49 for sample output.

```
bstnA(cfg)# show statistics filer fs4 connections processor 2.1
```

shows the connections from a particular network processor. See Figure 39.11 on page 39-50 for sample output.

Related Commands

clear statistics filer connections
show filer connections
show statistics global server
show external-filer
show statistics namespace ... fastpath

Figure 39.9 *Sample Output: show statistics filer connections*

```
bstnA(cfg)# show statistics filer connections

                                        Connections
  Filer                                 Current     Max
  ------------------------------------  ----------  ----------
  das1                                  12          12
  fs1                                   9           12
  fs2                                   5           13
  fs3                                   5           19
  fs4                                   11          20
  fs5                                   2           10
  fs6                                   3           17
  fs7                                   2           17
  das2                                  0           0
  das3                                  12          12
  nas1                                  206         210
  das7                                  12          12
  das8                                  12          12
  nas2                                  12          12
  nas3                                  12          12
  nas10                                 4           19
  nas11                                 4           6
  nasE1                                 14          16
  smb1                                  0           0
```

Figure 39.10 *Sample Output: show statistics filer fs4 connections*

```
bstnA(cfg)# show statistics filer fs4 connections

  Filer fs4 (192.168.25.29)              All processors
  -------------------------------------  --------------------

  CIFS:

    Connection limit:                    none

    Current data connections:            5
    Current control connections:         6
    Max connections (data+control):      20
    Time of max connections:             11/16/2012 01:17:26 -0500

    Current data sessions:               5
    Max data sessions:                   5
    Time of max data sessions:           11/16/2012 01:29:15 -0500

    Max sessions per data connection:    1
    Time of max sessions/connection:     11/16/2012 01:28:54 -0500

  NFS:

    Current connections:                 0
    Max connections:                     0
    Time of max connections:             n/a
```

Figure 39.11 *Sample Output: show statistics filer fs4 connections proc...*

```
bstnA(cfg)# show statistics filer fs4 connections processor 2.1

  Filer fs4 (192.168.25.29)               Slot.Proc 2.1
  ---------------------------------------  --------------------

  CIFS:

    Data connection limit:                none

    Current data connections:             0
    Current control connections:          0
    Current data sessions:                0

    Max sessions per data connection:     1
    Time of max sessions/connection:      11/16/2012 01:28:54 -0500

  NFS:

    Current connections:                  0

bstnA#
```

show statistics global server

Purpose Use the show statistics global server command to show the volume of traffic between clients and a global server. The output shows high-level packet counts, and can break the packet counts down to individual RPC calls and/or CIFS commands.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show statistics global server [interface | nfs | cifs]

show statistics global server fqdn [interface |nfs | cifs]

show statistics global server vip vip [interface | nfs | cifs]
```

fqdn (optional, 1-128 characters) is the fully-qualified domain name (for example, "myserver.organization.org") for a global server. Use show global server to see a list of global servers. If you omit this, the output includes all global servers.

vip (optional) identifies a single virtual-IP (VIP) address.

interface | nfs | cifs (optional) focuses the output on a single set of statistics. If you omit this option, the output includes all of the statistics that apply to the chosen ***fqdn*** or ***vip***.

interface shows only a table of high-level packet counts.

nfs displays only a table of NFS-procedure calls.

cifs shows only a table of CIFS commands.

Guidelines The output contains one or more tables for each global server. The following fields identify the global server:

Global Server identifies the global server, and

Virtual Server IP is the Virtual IP (VIP) for the virtual server.

Each global server has up to four tables: one for high-level packet counts, one for NFS RPCs, one for CIFS (SMB) commands, and one for SMB2 commands. The NFS table only appears if the ARX software has collected NFS statistics, and the CIFS and SMB2 tables only appear if the ARX software has collected CIFS or SMB2 statistics. The guidelines below describe each table in more detail.

To clear the counters shown by this command, use clear statistics global server. To examine traffic between the ARX and its filers, use show statistics namespace ... fastpath for raw read/write statistics or show statistics filer for NFS/CIFS call counters.

Guidelines: Interface Statistics

The interface table contains Ingress (incoming) and Egress (outgoing) packet counts for the following packet types:

Octets are the number of bytes received from clients or transmitted to clients.

Unicast Packets,

Non Unicast Packets, and

Dropped Packets are packet counts. The dropped packets in the Egress column are TCP packets that were dropped by the receiving station; the global server retransmits each of these dropped packets.

Guidelines: NFS Statistics

The NFS statistics contain one row per NFS-RPC call. These are calls from NFS clients to the chosen global server. Each row contains the following columns:

NFS RPC is the name of the RPC call.

Count is the number of these RPC calls received from NFS clients.

% Total shows the percentage of these RPC calls compared to all the other NFS calls.

Service Time (uSec) is the average number of microseconds (millionths of one second) from receiving the RPC to sending back a response.

Guidelines: NFS Errors

If there have been any NFS errors, they appear in a separate table under the NFS traffic statistics. The table contains the following counters:

- FormatError is the number of malformed NFS packets received from clients.
- NetworkError counts the send and receive failures.
- ResourceError is the number of times that the ARX software ran out of resources. Contact F5 Support if you see these errors.
- RPCError counts any errors from the RPC layer.
- TimeoutError is the number of timeouts waiting for an NFS response.

Guidelines: CIFS Statistics

The CIFS statistics contain one row per CIFS command. These are commands from CIFS clients to the chosen global server. Each row contains the following columns:

CIFS Command is the name of the CIFS command.

Count is the number of these CIFS commands received from CIFS clients.

% Total shows the percentage of these CIFS commands compared to all the other commands.

Service Time (uSec) is the average number of microseconds (millionths of one second) from receiving the command to sending back a response.

The NTCancel command is designed to count responses to Cancel commands. The CIFS protocol does not require any response to the Cancel command, so this is always 0 (zero).

The NTTransact NotifyChange command counts responses to Notify Change commands (see the documentation for cifs notify-change-mode). These response times can be arbitrarily long, so they are omitted from the count and always appear as 0 (zero).

Guidelines: CIFS Errors

Any CIFS errors appear in a table under the CIFS statistics, with the following fields:

- FormatError is the number of malformed SMB (Server Message Block, an earlier name for CIFS) packets received from clients.
- NetworkError counts the send and receive failures for SMB traffic.
- ResourceError shows the number of times that the ARX software had insufficient resources to complete the transaction. Contact F5 Support if you see these errors.
- SignatureError counts any SMB packets with missing or incorrect signatures. You can use the signatures command to change the SMB-signing policy for this global server's CIFS service.
- SMBError is the number of SMB packets that failed with errors that are not directly related to any of the above errors.
- TimeoutError is the number of timeouts waiting for an SMB response.

Guidelines: SMB2 Statistics

The SMB2 statistics only appear for a service that supports SMB2. You can use the smb2-allowed command to enable SMB2 for the ARX. Like the other tables, this contains one row per SMB2 command. These are commands from SMB2 clients to the chosen global server. Each row contains the following columns:

SMB2 Command is the name of the SMB2 command.

Count is the number of these SMB2 commands received from clients.

% Total shows the percentage of these SMB2 commands compared to all the other commands.

Service Time (uSec) is the average number of microseconds (millionths of one second) from receiving the command to sending back a response.

The Cancel command is designed to count responses to Cancel commands. The CIFS protocol does not require any response to a this command, so this is always 0 (zero).

The ChangeNotify command counts responses to Notify Change commands (see the documentation for cifs notify-change-mode). These response times can be arbitrarily long, so they are omitted from the count and always appear as 0 (zero).

Guidelines: SMB2 Errors

Any SMB2 errors appear in a table under the SMB2 statistics. These have the same fields as the CIFS Errors table described above.

Samples

```
bstnA(cfg)# show statistics global server
```

shows statistics for all global servers. See Figure 39.12 on page 39-54 for sample output.

```
bstnA(cfg)# show statistics global server ac1.medarch.org
```

shows statistics for the global server at "ac1.medarch.org." See Figure 39.13 on page 39-58 for sample output.

```
bstnA(cfg)# show statistics global server vip 192.168.25.12
```

shows the statistics for one virtual server. See Figure 39.14 on page 39-60 for sample output.

Related Commands

clear statistics global server
show statistics namespace ... fastpath
show statistics filer

Figure 39.12 *Sample Output: show statistics global server*

```
bstnA(cfg)# show statistics global server

Global Server:           acopiaFiler
Virtual Server IP:       192.168.25.12

                                  Ingress                  Egress
------------------ ---------------------- -----------------------
Octets                          160601276                78196900
Unicast Packets                    210464                  186034
Non Unicast Packets                     0                       0
Dropped Packets                         0                       0

Last Reset (Local Time): 02/12/2013 08:08:06 -0500

NFS RPC                            Count    % Total Service Time (uSec)
-------------------------- -------------- ---------- ------------------
Null                                   0         0%                 --
Getattr                            40707        42%                222
Setattr                            18460        19%               5497
Lookup                             14465        15%                252
Access                               827         1%                234
Readlink                               3         0%                229
Read                                4186         4%               4695
Write                               8537         9%               9747
Create                              2816         3%               5925
Mkdir                                276         0%               5589
Symlink                                2         0%              14425
Mknod                                  0         0%                 --
Remove                                57         0%               7426
Rmdir                                  0         0%                 --
Rename                                 0         0%                 --
Link                                   0         0%                 --
Readdir                              668         1%                336
Readdirplus                            0         0%                 --
Fsstat                                 4         0%                323
Fsinfo                                 4         0%                291
Pathconf                               0         0%                 --
Commit                              4997         5%              23738
StatFs                                 0         0%                 --
-------------------------- -------------- ---------- ------------------
Total                              96009       100%                3695

NFS Errors                         Count
-------------------------- --------------
FormatError                            0
NetworkError                           0
ResourceError                          0
RPCError                               0
TimeoutError                           0
-------------------------- --------------
Total                                  0


Global Server:           ac1.MEDARCH.ORG
Virtual Server IP:       192.168.25.15

                                  Ingress                  Egress
------------------ ---------------------- -----------------------
Octets                          565533226               379772155
Unicast Packets                    868117                  753566
Non Unicast Packets                     0                       0
```

```
Dropped Packets                          0                     0

Last Reset (Local Time): 02/12/2013 08:08:46 -0500

NFS RPC                              Count    % Total Service Time (uSec)
-------------------------- --------------- ---------- -------------------
Null                                     8         0%                 826
Getattr                              77723        14%                 227
Setattr                              31208         5%                1200
Lookup                               22982         4%                 446
Access                              330088        58%                 246
Readlink                                44         0%                 179
Read                                 38236         7%                4621
Write                                59958        10%                1767
Create                                2868         1%                2721
Mkdir                                  274         0%                5471
Symlink                                 11         0%                8005
Mknod                                    0         0%                  --
Remove                                  98         0%                2735
Rmdir                                    0         0%                  --
Rename                                   3         0%                4730
Link                                     0         0%                  --
Readdir                               1059         0%                1167
Readdirplus                              0         0%                  --
Fsstat                                   7         0%                1594
Fsinfo                                   7         0%                 947
Pathconf                                 0         0%                  --
Commit                                8750         2%                2752
StatFs                                   0         0%                  --
-------------------------- --------------- ---------- -------------------
Total                                573324       100%                 809

NFS Errors                           Count
-------------------------- ---------------
FormatError                              0
NetworkError                             0
ResourceError                            0
RPCError                                 0
TimeoutError                             0
-------------------------- ---------------
Total                                    0

Last Reset (Local Time): 02/12/2013 08:10:23 -0500

CIFS Command                         Count    % Total Service Time (uSec)
-------------------------- --------------- ---------- -------------------
CheckDir                                 0         0%                  --
Close                                  241        14%                 516
Create                                   0         0%                  --
CreateDir                                1         0%                 289
Delete                                   8         0%                1204
DeleteDir                                9         1%                7324
Echo                                     0         0%                  --
FindClose2                               0         0%                  --
Flush                                    0         0%                  --
LockingAndX                              0         0%                  --
LogoffAndX                              62         4%                2973
Negotiate                               42         2%                 249
NTCancel                                 0         0%                  --
NTCreateAndX                           119         7%                4568
NTRename                                 0         0%                  --
NTTransact
```

```
  Create                                  0         0%                  --
  GetUserQuota                            0         0%                  --
  Ioctl                                   0         0%                  --
  NotifyChange                            0         0%                  --
  QuerySecurity                           0         0%                  --
  SetSecurity                             0         0%                  --
Open                                      0         0%                  --
OpenAndX                                124         7%                1944
QueryInfo                                 0         0%                  --
QueryInfoDisk                             6         0%                 145
Read                                      0         0%                  --
ReadAndX                                 16         1%                1081
Rename                                    0         0%                  --
Seek                                      0         0%                  --
SessionSetupAndX                        153         9%                2704
SetInfo                                   0         0%                  --
Transaction                               6         0%               18229
Transaction2
  CreateDir                               0         0%                  --
  FindFirst2                             84         5%                5916
  FindNext2                               0         0%                  --
  GetDFSReferral                          0         0%                  --
  QueryFSInfo                           161         9%                 172
  QueryFileInfo                           5         0%                 341
  QueryPathInfo                          99         6%                3619
  SetFileInfo                            32         2%                 328
  SetPathInfo                             0         0%                  --
TreeConnectAndX                         104         6%                 777
TreeDisconnect                           19         1%                 250
Write                                     0         0%                  --
WriteAndX                               455        26%                 678
UnknownSMB                                0         0%                  --
------------------------- --------------- ---------- -------------------
Total                                  1746       100%                1724

CIFS Errors                           Count
------------------------- ---------------
FormatError                               0
NetworkError                              0
ResourceError                             0
SignatureError                            0
SMBError                                 53
TimeoutError                              0
------------------------- ---------------
Total                                    53


Global Server:           labtests.MEDARCH.ORG
Virtual Server IP:       192.168.25.16

                                    Ingress                  Egress
------------------- ---------------------- ----------------------
Octets                                71304                   87191
Unicast Packets                         406                     358
Non Unicast Packets                       0                       0
Dropped Packets                           0                       0

Last Reset (Local Time): 02/12/2013 08:33:52 -0500

CIFS Command                          Count    % Total Service Time (uSec)
------------------------- --------------- ---------- -------------------
CheckDir                                  0         0%                  --
```

```
Close                                    0         0%                  --
Create                                   0         0%                  --
CreateDir                                0         0%                  --
Delete                                   0         0%                  --
DeleteDir                                0         0%                  --
Echo                                     0         0%                  --
FindClose2                               0         0%                  --
Flush                                    0         0%                  --
LockingAndX                              0         0%                  --
LogoffAndX                               0         0%                  --
Negotiate                                0         0%                  --
NTCancel                                 0         0%                  --
NTCreateAndX                             0         0%                  --
NTRename                                 0         0%                  --
NTTransact
  Create                                 0         0%                  --
  GetUserQuota                           0         0%                  --
  Ioctl                                  0         0%                  --
  NotifyChange                           0         0%                  --
  QuerySecurity                          0         0%                  --
  SetSecurity                            0         0%                  --
Open                                     0         0%                  --
OpenAndX                                 0         0%                  --
QueryInfo                                0         0%                  --
QueryInfoDisk                            0         0%                  --
Read                                     0         0%                  --
ReadAndX                                 0         0%                  --
Rename                                   0         0%                  --
Seek                                     0         0%                  --
SessionSetupAndX                         0         0%                  --
SetInfo                                  0         0%                  --
Transaction                              0         0%                  --
Transaction2
  CreateDir                              0         0%                  --
  FindFirst2                             0         0%                  --
  FindNext2                              0         0%                  --
  GetDFSReferral                         0         0%                  --
  QueryFSInfo                            0         0%                  --
  QueryFileInfo                          0         0%                  --
  QueryPathInfo                          0         0%                  --
  SetFileInfo                            0         0%                  --
  SetPathInfo                            0         0%                  --
TreeConnectAndX                          0         0%                  --
TreeDisconnect                           0         0%                  --
Write                                    0         0%                  --
WriteAndX                                0         0%                  --
UnknownSMB                               0         0%                  --
-------------------------- --------------- ---------- -------------------
Total                                    0         0%                  --

CIFS Errors                           Count
-------------------------- ---------------
FormatError                              0
NetworkError                             0
ResourceError                            0
SignatureError                           0
SMBError                                 0
TimeoutError                             0
-------------------------- ---------------
Total                                    0

SMB2 Command                          Count    % Total Service Time (uSec)
```

```
-------------------------- --------------- ---------- -------------------
Cancel                                   0         0%                  --
ChangeNotify                             0         0%                  --
Close                                   64        23%                2427
Create                                  58        21%                4869
Echo                                     0         0%                  --
Flush                                    0         0%                  --
Ioctl                                   10         4%                 184
Lock                                     0         0%                  --
Logoff                                  12         4%                 390
Negotiate                               25         9%                 219
QueryDirectory                           0         0%                  --
QueryInfo                               12         4%                 256
Read                                     0         0%                  --
SessionSetup                            36        13%                1930
SetInfo                                 23         8%                2506
TreeConnect                             17         6%                 819
TreeDisconnect                          17         6%                 548
Write                                    0         0%                  --
UnknownSMB                               0         0%                  --
-------------------------- --------------- ---------- -------------------
Total                                  274       100%                2201

SMB2 Errors                           Count
-------------------------- ---------------
FormatError                               0
NetworkError                              0
ResourceError                             0
SignatureError                            0
SMBError                                  0
TimeoutError                              0
-------------------------- ---------------
Total                                     0
```

Figure 39.13 *Sample Output: show statistics global server ac1.medarch.org*

```
bstnA(cfg)# show statistics global server ac1.medarch.org

Global Server:           ac1.MEDARCH.ORG
Virtual Server IP:       192.168.25.15

                                  Ingress                  Egress
------------------- ---------------------- ----------------------
Octets                          565533226              379772155
Unicast Packets                    868117                 753566
Non Unicast Packets                     0                      0
Dropped Packets                         0                      0

Last Reset (Local Time): 02/12/2013 08:08:46 -0500

NFS RPC                               Count    % Total Service Time (uSec)
-------------------------- --------------- ---------- -------------------
Null                                      8         0%                 826
Getattr                               77723        14%                 227
Setattr                               31208         5%                1200
Lookup                                22982         4%                 446
Access                               330088        58%                 246
Readlink                                 44         0%                 179
Read                                  38236         7%                4621
Write                                 59958        10%                1767
Create                                 2868         1%                2721
```

```
Mkdir                                274         0%                5471
Symlink                               11         0%                8005
Mknod                                  0         0%                  --
Remove                                98         0%                2735
Rmdir                                  0         0%                  --
Rename                                 3         0%                4730
Link                                   0         0%                  --
Readdir                             1059         0%                1167
Readdirplus                            0         0%                  --
Fsstat                                 7         0%                1594
Fsinfo                                 7         0%                 947
Pathconf                               0         0%                  --
Commit                              8750         2%                2752
StatFs                                 0         0%                  --
--------------------------- --------------- ---------- -------------------
Total                             573324       100%                 809

NFS Errors                          Count
--------------------------- ---------------
FormatError                            0
NetworkError                           0
ResourceError                          0
RPCError                               0
TimeoutError                           0
--------------------------- ---------------
Total                                  0

Last Reset (Local Time): 02/12/2013 08:10:23 -0500

CIFS Command                        Count    % Total Service Time (uSec)
--------------------------- --------------- ---------- -------------------
CheckDir                               0         0%                  --
Close                                241        14%                 516
Create                                 0         0%                  --
CreateDir                              1         0%                 289
Delete                                 8         0%                1204
DeleteDir                              9         1%                7324
Echo                                   0         0%                  --
FindClose2                             0         0%                  --
Flush                                  0         0%                  --
LockingAndX                            0         0%                  --
LogoffAndX                            62         4%                2973
Negotiate                             42         2%                 249
NTCancel                               0         0%                  --
NTCreateAndX                         119         7%                4568
NTRename                               0         0%                  --
NTTransact
  Create                               0         0%                  --
  GetUserQuota                         0         0%                  --
  Ioctl                                0         0%                  --
  NotifyChange                         0         0%                  --
  QuerySecurity                        0         0%                  --
  SetSecurity                          0         0%                  --
Open                                   0         0%                  --
OpenAndX                             124         7%                1944
QueryInfo                              0         0%                  --
QueryInfoDisk                          6         0%                 145
Read                                   0         0%                  --
ReadAndX                              16         1%                1081
Rename                                 0         0%                  --
Seek                                   0         0%                  --
SessionSetupAndX                     153         9%                2704
```

```
SetInfo                                0         0%                --
Transaction                            6         0%             18229
Transaction2
  CreateDir                            0         0%                --
  FindFirst2                          84         5%              5916
  FindNext2                            0         0%                --
  GetDFSReferral                       0         0%                --
  QueryFSInfo                        161         9%               172
  QueryFileInfo                        5         0%               341
  QueryPathInfo                       99         6%              3619
  SetFileInfo                         32         2%               328
  SetPathInfo                          0         0%                --
TreeConnectAndX                      104         6%               777
TreeDisconnect                        19         1%               250
Write                                  0         0%                --
WriteAndX                            455        26%               678
UnknownSMB                             0         0%                --
------------------------- --------------- ---------- ------------------
Total                               1746       100%              1724

CIFS Errors                        Count
-------------------------- ---------------
FormatError                            0
NetworkError                           0
ResourceError                          0
SignatureError                         0
SMBError                              53
TimeoutError                           0
-------------------------- ---------------
Total                                 53
```

Figure 39.14 *Sample Output: show statistics global server vip 192.168.25.12*

```
bstnA(cfg)# show statistics global server vip 192.168.25.12

Global Server:           acopiaFiler
Virtual Server IP:       192.168.25.12

                                 Ingress                  Egress
------------------ ---------------------- ----------------------
Octets                         160601276               78196900
Unicast Packets                   210464                 186034
Non Unicast Packets                    0                      0
Dropped Packets                        0                      0

Last Reset (Local Time): 02/12/2013 08:08:06 -0500

NFS RPC                            Count    % Total Service Time (uSec)
-------------------------- --------------- ---------- ------------------
Null                                   0         0%                --
Getattr                            40707        42%               222
Setattr                            18460        19%              5497
Lookup                             14465        15%               252
Access                               827         1%               234
Readlink                               3         0%               229
Read                                4186         4%              4695
Write                               8537         9%              9747
Create                              2816         3%              5925
Mkdir                                276         0%              5589
Symlink                                2         0%             14425
Mknod                                  0         0%                --
```

```
Remove                                57         0%                7426
Rmdir                                  0         0%                  --
Rename                                 0         0%                  --
Link                                   0         0%                  --
Readdir                              668         1%                 336
Readdirplus                            0         0%                  --
Fsstat                                 4         0%                 323
Fsinfo                                 4         0%                 291
Pathconf                               0         0%                  --
Commit                              4997         5%               23738
StatFs                                 0         0%                  --
--------------------------- -------------- ---------- ------------------
Total                              96009       100%                3695

NFS Errors                         Count
--------------------------- --------------
FormatError                            0
NetworkError                           0
ResourceError                          0
RPCError                               0
TimeoutError                           0
--------------------------- --------------
Total                                  0
```

40

Managed-Volume Troubleshooting Tools

cancel migration

Purpose The cancel migration command stops a file migration that is in progress. You can cancel all migrations in a managed volume or the migration of a single file.

Modes priv-exec

Security Role(s) crypto-officer or storage-engineer

Syntax

```
cancel migration ns volume vol-path
cancel migration file file-name
```

ns (1-30 characters) is the namespace where a migration is in process.

vol-path (1-1024 characters) specifies a managed volume.

file-name (1-1024 characters) is the virtual path of a file that is being migrated. This includes the path name for the managed volume itself (for example, the file name for "mydir/myfile.doc" in the "/bigvol" volume would be "/bigvol/mydir/myfile.doc").

Default(s) None

Guidelines This command removes one file (or all files) from a managed volume's file-migration queue. You can examine a volume's migration queue with the show policy queue command.

If you cancel all migrations in a managed volume, the CLI prompts for confirmation before clearing the queue. Enter **yes** to proceed.

Samples

```
bstnA# cancel migration medarcv volume /rcrds
Are you sure you want to cancel all the queued migrates in volume
'/rcrds'? [yes/no] yes
bstnA#
```

clears the entire migration queue for the "medarcv~/rcrds" volume.

```
bstnA# cancel migration file /claims/2007/masterlist.xsl
```

removes a single file from its volume's migration queue.

Related Commands show policy queue

cancel remove

Purpose The cancel remove command stops a share-removal that is in progress.

Modes priv-exec

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or backup-operator

Syntax **`cancel remove namespace`** ***`ns`*** **`volume`** ***`vol-path`*** **`share`** ***`share-name`***

ns (1-30 characters) is the namespace from which the share is being removed.

vol-path (1-1024 characters) specifies the namespace volume.

share-name (1-64 characters) is the share that is being removed.

Default(s) None

Guidelines You can remove a share with remove-share migrate, remove-share nomigrate, no share, or no filer. The first stage of the share removal is a scan of the filer; this reads all file attributes on the share's back-end directories. During this phase (the longest in the process), you can cancel the share removal. After this phase is over, the CLI does not permit anyone to cancel the operation.

Sample

```
bstnA# cancel remove namespace ns volume /vol share testrun
Storage job cancelled successfully.
bstnA#
```

cancels the removal in progress for the "ns~/vol~testrun" share.

Related Commands

remove-share migrate
remove-share nomigrate
share
filer

clear statistics filer

Purpose Use this command to clear the NFS/CIFS traffic statistics for a back-end filer or share.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
clear statistics filer
clear statistics filer ext-filer-name [share share-name]
```

ext-filer-name (optional, 1-64 characters) identifies the external filer by its configured name. Use show external-filer to display all configured external filers. If you omit this argument, the command clears all traffic statistics on the system. This option is not recommended, as it results in incomplete statistics for some forms of show statistics namespace.

share-name (optional, 1-64 characters) clears traffic statistics for this specific back-end share. This is the name of the share at the back-end filer, not at the namespace. As above, this option is not recommended; it results in incomplete statistics for some forms of show statistics namespace.

Default(s) None

Guidelines The show statistics filer command shows the NFS and/or CIFS traffic between the ARX's software and its back-end filers. Use this command to clear those traffic counters. This also clears the traffic counters for the various forms of show statistics namespace: show statistics namespace ... fastpath, show statistics namespace ... request-detail, show statistics namespace ... response-detail, and show statistics namespace ... summary.

The CLI prompts for confirmation before clearing the statistics; enter **yes** to proceed.

Sample

```
bstnA# clear statistics filer das1
Clear statistics of ALL shares on filer das1.

Are you sure? [yes/no] yes
```

clears the current NFS/CIFS statistics counts for the external filer named "das1."

Related Commands

show statistics filer
show statistics namespace ... fastpath
show statistics namespace ... request-detail
show statistics namespace ... response-detail
show statistics namespace ... summary
show external-filer

clear statistics metadata

Purpose Managed volumes use *metadata* to track the locations of their files on back-end storage. They store the metadata on an external metadata share, and keep statistics on the usage of that share. Use this command to clear these statistics for the entire system, a specific namespace, or particular managed volume.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear statistics metadata [namespace`** ***`ns`*** **`[volume`** ***`vol`*** **`]]`**

ns (optional, 1-30 characters) identifies a namespace.

vol (optional, 1-1024 characters) is a managed volume.

Default(s) None

Guidelines The show statistics metadata command shows usage statistics for the metadata shares behind your managed volumes. Use this command to clear these statistical counters for the whole system, one namespace, or one managed volume.

The CLI prompts for confirmation before clearing the statistics; enter **yes** to proceed.

Sample

```
bstnA# clear statistics metadata
Clear metadata performance statistics for all volumes?

Are you sure? [yes/no]  yes
```

clears all metadata-share statistics for the entire ARX.

Related Commands show statistics metadata

clear statistics metalog

Purpose Volume software stores *metalog* data to be used for recovery from failovers (see redundancy) and crashes. The volume software needs to read and write this metalog data quickly to ensure proper performance; you can use the show statistics metalog command to monitor the read/write statistics for metalog storage. Use this command to clear these statistics for the entire system or a specific volume-group.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear statistics metalog [volume-group`** ***`vol-group-id`*****`]`**

vol-group-id (optional, 1-255 characters) identifies a volume group. You can use the show volume-group command to show all volume groups, which volumes are in each of them, and their ID numbers.

Default(s) None

Guidelines The show statistics metalog command shows usage statistics for the metalog storage behind your managed volumes. These statistics record the latency for metalog reads and writes. Use this command to clear these statistical counters for the whole system or a single volume group.

The CLI prompts for confirmation before clearing the statistics; enter **yes** to proceed.

Sample

```
bstnA# clear statistics metalog volume-group 9
Clear metalog performance statistics for volume group 9?

Are you sure? [yes/no]  yes
```

clears all metalog statistics for a particular volume group.

Related Commands show statistics metalog

clear statistics migration

Purpose Managed volumes keep statistics on all of their file migrations. Use this command to clear these statistics for an entire namespace or particular managed volume.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear statistics migration`** ***`namespace`*** **`[`*****`volume`*****`]`**

namespace (1-30 characters) identifies a namespace.

volume (optional, 1-1024 characters) is a managed volume.

Default(s) None

Guidelines The show statistics migration command shows file-migration statistics for managed volumes. Use this command to clear the statistical counters for one namespace or managed volume.

The CLI prompts for confirmation before clearing the statistics; enter **yes** to proceed.

Sample

```
bstnA# clear statistics migration medarcv /rcrds

Clear data migration statistics? [yes/no] yes
```

clears all file-migration statistics for the “medarcv~/rcrds” volume.

Related Commands show statistics migration

export-mapping

Purpose Some volumes are configured to constrain each directory to a single back-end share.This constraint simplifies volume backups: you can backup the entire volume directly from the filer-share. Use the export-mapping command to create a report that shows the front-end/back-end mapping for these volume configurations.

Modes priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`export-mapping [outputfile`** ***`outputfile`*** **`] [nfs | cifs]`**

outputfile ***outputfile*** (optional; 5-128 characters) specifies a prefix for a customized file name (for example, 'jrandom_exp_map.rpt'), as opposed to the default. Use the show reports command to display the file in the maintenance directory.

nfs | cifs (optional) narrows the report to a single protocol, NFS or CIFS.

Default(s) **outputfile** ***outputfile*** - 'export-mapping.rpt.'

Guidelines The name of the report appears after you issue the command. For long reports, you can use tail to follow the report as it is written. Use show reports export-mapping.rpt to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands. If you want to truncate the report before it finishes, use the truncate-report command.

The report is divided into three tables. The first shows all front-end services, the next shows all back-end filers, and the third table maps all front-end shares to a single back-end path. The single back-end path is the *master directory* for the root of the front-end share; by default, all of the front-end share's subdirectories and files go to this back-end path. If it is the only share in the volume, or if policy is not running in the volume, this contains all of the files and directories in the front-end share. Otherwise, files and directories may be distributed among the volume's shares; use show server-mapping for a map of all back-end shares behind each front-end share.

Samples

```
bstnA# export-mapping
Generated report: export-mapping.rpt on switch bstnA
bstnA# ...
```

A sample appears below in Figure 40.1.

```
bstnA# export-mapping outputfile nfs_only.rpt nfs
Generated report: nfs_only.rpt on switch bstnA
bstnA# ...
```

generates a mapping report for NFS only. See Figure 40.2 on page 40-11 for a sample report.

Related Commands

show global-config **namespace**
restore data
show reports
truncate-report
tail
grep
copy ftp
copy scp
copy {nfs|cifs}
copy tftp
delete

Figure 40.1 *Sample Report: export-mapping*

```
bstnA# show reports export-mapping.rpt
**** Export Mapping Report: Started at Sat Feb 27 01:31:41 2010 ****
**** Software Version: 5.02.000.12541 (Feb 23 2010 20:12:44) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

Proto  Virtual Server                                        IP Address
-----  ------------------------------------------------  ---------------
NFS    acopiaFiler                                         192.168.25.12
NFS    www.wwmed.com                                       192.168.25.10
CIFS   ac1.medarch.org                                     192.168.25.15
CIFS   insur.medarch.org                                   192.168.25.14
NFS    insur.medarch.org                                   192.168.25.14

External Filer                                               IP Address
-----------------------------------------------------------  ---------------
das1                                                       192.168.25.19
fs1                                                        192.168.25.20
fs2                                                        192.168.25.27
fs3                                                        192.168.25.28
fs4                                                        192.168.25.29
fs5                                                        192.168.25.71
das2                                                       192.168.25.22
das3                                                       192.168.25.23
nas1                                                       192.168.25.21
das7                                                       192.168.25.24
das8                                                       192.168.25.25
nas2                                                       192.168.25.44
nas3                                                       192.168.25.47
nas10                                                      192.168.25.49
nasE1                                                      192.168.25.51
smb1                                                       192.168.25.48

Proto  Export                              Source
-----  ----------------------------------  ---------------------------------
NFS    acopiaFiler:/vol                    VIRTUAL PATH
NFS    www.wwmed.com:/acct                 das1:/exports/budgett
CIFS   \\ac1.medarch.org\ARCHIVES          \\fs4\prescriptions\
CIFS   \\ac1.medarch.org\Y2005             \\fs4\prescriptions\2005
CIFS   \\ac1.medarch.org\labs              \\nas10\equipment\
CIFS   \\ac1.medarch.org\bulkstorage       \\fs4\prescriptions\
CIFS   \\ac1.medarch.org\xraysScanners     \\nas10\equipment\
CIFS   \\ac1.medarch.org\chem_results      VIRTUAL PATH
CIFS   \\ac1.medarch.org\CELEBS            \\fs4\prescriptions\VIP_wing
CIFS   \\ac1.medarch.org\Y2004             \\fs4\prescriptions\2004
```

```
CIFS   \\insur.medarch.org\CLAIMS            \\nas1\insurances
CIFS   \\insur.medarch.org\SPECS             \\nas1\insurance\specs
CIFS   \\insur.medarch.org\STATS             \\nas1\insurance\stats
NFS    insur.medarch.org:/claims             nas1:/vol/vol1/NTFS_QTREE/insurances

**** Virtual Servers Processed:                  5
**** External Filers Processed:                 16
**** Exports Processed:                         14

**** Elapsed time:           00:00:00
**** Export Mapping Report: DONE at Sat Feb 27 01:31:41 2010 ****
```

Figure 40.2 *Sample Report: export-mapping ... nfs*

```
bstnA# show reports nfs-only.rpt
**** Export Mapping Report: Started at Sat Feb 27 01:31:42 2010 ****
**** Software Version: 5.02.000.12541 (Feb 23 2010 20:12:44) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

Proto  Virtual Server                                        IP Address
-----  ----------------------------------------------------  ---------------
NFS    acopiaFiler                                           192.168.25.12
NFS    www.wwmed.com                                         192.168.25.10
NFS    insur.medarch.org                                     192.168.25.14

External Filer                                               IP Address
-----------------------------------------------------------  ---------------
das1                                                         192.168.25.19
fs1                                                          192.168.25.20
fs2                                                          192.168.25.27
fs3                                                          192.168.25.28
fs4                                                          192.168.25.29
fs5                                                          192.168.25.71
das2                                                         192.168.25.22
das3                                                         192.168.25.23
nas1                                                         192.168.25.21
das7                                                         192.168.25.24
das8                                                         192.168.25.25
nas2                                                         192.168.25.44
nas3                                                         192.168.25.47
nas10                                                        192.168.25.49
nasE1                                                        192.168.25.51
smb1                                                         192.168.25.48

Proto  Export                                 Source
-----  -------------------------------------  ---------------------------------
NFS    acopiaFiler:/vol                       VIRTUAL PATH
NFS    www.wwmed.com:/acct                    das1:/exports/budgett
NFS    insur.medarch.org:/claims              nas1:/vol/vol1/NTFS_QTREE/insurances

**** Virtual Servers Processed:                  3
**** External Filers Processed:                 16
**** Exports Processed:                          3

**** Elapsed time:           00:00:00
**** Export Mapping Report: DONE at Sat Feb 27 01:31:42 2010 ****
```

find

Purpose Use the find command to locate one file's back-end filer and path.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

`find host` ***`hostname-or-ip`*** **`{nfs| cifs}`** ***`share-name`*** **`path`** ***`path`*** **`[verbose]`**

hostname-or-ip identifies a virtual server by its external DNS hostname or Virtual-IP address. Use show virtual service for a list of all virtual servers and their Virtual IP addresses. (To look up DNS names at an external DNS server, you must first identify the server with the ip name-server command.)

nfs | cifs is a required choice, determining the type of share that follows.

share-name (1-4096 characters) is the name of the front-end share, as seen by clients. The show server-mapping command shows all front-end shares on the ARX, in the left column.

path (1-4096 characters) is the virtual-file path within the front-end share.

verbose (optional) adds the file's NFS filehandles to the output.

`find global-server` ***`fqdn`*** **`{nfs| cifs}`** ***`share-name`*** **`path`** ***`path`*** **`[verbose]`**

fqdn (1-255 characters) identifies a global server by its fully-qualified domain name (FQDN). Use show global service for a list of all global servers and their FQDNs.

The remaining options are described above.

`find namespace` ***`namespace`*** **`path`** ***`path`*** **`[verbose]`**

namespace (1-30 characters) identifies the namespace. Use show namespace for a list of all namespaces.

path (1-4096 characters) is the virtual-file path within the namespace. This is the file path as seen by the client, starting with the volume path.

verbose (optional) adds the file's NFS filehandles to the output.

`find wins` ***`wins-name`*** **`{nfs| cifs}`** ***`share-name`*** **`path`** ***`path`***

wins-name (1-255 characters) identifies a virtual server by its optional wins-name or one of its optional wins-alias names.

The remaining options are described above.

Default(s) None.

Guidelines

This command finds one file's physical location. To find all file locations behind a namespace, use the nsck ... report metadata-only command.

For files in multi-protocol (NFS and CIFS) namespaces, this shows the file's physical location from an NFS point of view and a CIFS point of view.

The output contains several field to map the ARX-volume location of the file to the file's physical location on the filer:

NFS Physical Location is an NFS path to the physical file. This appears in *ip-address*:/*path-to-file* format.

Managed Volume Path indicates that the ***path*** is on a direct volume, and the physical location is on an ARX managed-volume that is standing in as a filer.

CIFS Physical Location is a CIFS path to the physical file. This appears in //*ip-address*/*path-to-file* format. If this is an NFS-only file in a multi-protocol (CIFS and NFS) volume, a message here indicates that the file has an inconsistent name. You can use the nsck ... report inconsistencies command to find all of the volume's inconsistent names.

The **verbose** output shows two NFS filehandles for the file:

Virtual File Handle is the filehandle that the ARX presents to clients.

Physical File Handle is the filehandle that the server gave to the ARX.

The filehandles may be incomplete for the find namespace command. Use one of the other versions (such as find host) for complete NFS filehandles.

Guidelines: File Tracking

The find command identifies a file's location now; for a file's location in the past, you can use *file tracking*. File tracking can be useful for finding the correct filer-backup tape for a lost or compromised file.

The file-tracking feature requires some configuration before you can make any file queries. Specifically, a managed volume requires a snapshot rule to regularly copy its configuration and metadata to a file-history archive. After some copies of metadata have been archived, you can use the show file-history virtual-service command to query file locations at different dates.

Samples

```
bstnA# find namespace wwmed path /acct/index.html

Namespace:             wwmed
Logical path:          /acct/index.html
NFS Physical location: 192.168.25.19:/exports/budget/index.html
```

finds a file in a namespace.

```
bstnA# find host 192.168.25.14 nfs /claims path /stats/piechart.ppt

Namespace:             insur
Logical path:          /claims/stats/piechart.ppt
NFS Physical location: 192.168.25.21:/vol/vol1/NTFS-QTREE/insurance/stats/piechart.ppt
CIFS Physical location: The file/directory on //192.168.25.21/insurance has an inconsistent
name.
```

finds a file behind a virtual server. The virtual server is exporting the file from a multi-protocol namespace, so it shows both NFS and CIFS locations for the file. This file has an inconsistent, NFS-only name.

```
bstnA# find global-server insur.medarch.org nfs /claims path /stats/piechart.ppt

Namespace:              insur
Logical path:           /claims/stats/piechart.ppt
NFS Physical location: 192.168.25.21:/vol/vol1/NTFS-QTREE/insurance/stats/piechart.ppt
CIFS Physical location: The file/directory on //192.168.25.21/insurance has an inconsistent
name.
```

finds the same file behind its global server.

```
bstnA# find global-server insur.medarch.org cifs CLAIMS path /index.html verbose

Namespace:              insur
Logical path:           /claims/index.html
NFS Physical location:  192.168.25.21:/vol/vol1/NTFS-QTREE/insurance/index.html
CIFS Physical location: //192.168.25.21/insurance/index.html

  NFSv3
    Virtual File Handle  (32 bytes):
      0x00000010b0000000100000000000400000000000000000000000000400000000
    Physical File Handle (32 bytes):
      0x01c63600f00da62220000000001ef2a0b0727e2104cf00f1400000008be32b00
```

provides a verbose display for another file behind the same global server.

Related Commands nsck ... report metadata-only

remove service

Purpose Use the remove service command to remove a namespace and all other configuration objects that are exclusively dedicated to the namespace (such as global servers, proxy users, and/or NFS access-lists).

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax `remove service` ***namespace*** `[timeout` ***seconds***`] [sync]`

namespace (1-30 characters) identifies the namespace.

seconds (optional, 300-36,000) is a time limit for each component operation in this service removal, in seconds.

sync (optional) shows the operation's progress at the command line. With this option, the CLI prompt does not return until all components have been removed.

Default(s) None.

Guidelines The CLI prompts for confirmation before removing any configuration objects or metadata; enter **yes** to proceed.

This removes the following objects from the system:

- metadata behind all of the namespace's volumes
- all objects created for CIFS browsing
- all front-end exports that reference the namespace's volumes
- all global and virtual servers that are dedicated to this namespace
- all policy (rules, snapshots, share farms, and filesets)
- all shares and metadata shares
- any NFS access lists that are exclusively devoted to the namespace
- any external filers that are exclusively devoted to the namespace

By default, this command generates a report to show all of the actions it takes to remove the volume(s), in order. The CLI shows the report name after you issue the command, and then returns. You can enter CLI commands as the namespace software removes the objects in the background. Use tail to follow the report as it is written. Use show reports file-name to read the report. You can search through the report with grep. To copy or delete it, use the copy or delete commands. Use the sync option to send the status to the command line instead; the command does not generate a report if you use the sync option.

If you remove a cifs service with any dynamic-dns names, you must remove all of the Dynamic DNS names (with no dynamic-dns) before you run this command.

Use remove namespace to remove a namespace or volume without affecting any global servers, NFS-access-control lists, or external filers. To remove only the policy objects from the namespace or a single volume, use remove namespace ... policy-only. The remove namespace ... volume ... exports-only command finds all of a volumes front-end exports and removes them. To remove a share from a volume, use remove-share migrate or remove-share nomigrate.

Sample

```
bstnA# remove service medco

Remove service 'medco'? [yes/no] yes
Scheduling report: remService_medco_201002240751.rpt
```

removes all configuration objects associated with the “medco” namespace.

Related Commands

remove namespace
remove namespace ... policy-only
remove namespace ... volume ... exports-only
remove-share migrate
remove-share nomigrate

remove-share migrate

Purpose Use the remove-share migrate command to remove a share from a namespace volume and move all of its files to another share or share farm in the same volume. The volume's clients are unaffected by this share removal.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax

```
remove-share migrate namespace volume share dest
    [close-file [exclude fileset]] [async]
```

namespace (1-30 characters) identifies the namespace.

volume (1-1024 characters) is the name of the volume.

share (1-64 characters) is share being removed.

dest (1-64 characters) is the new destination share or share farm for files and master directories.

close-file (optional) causes the operation to close all files open by CIFS clients in order to migrate them. If you omit this option, files opened by CIFS clients cannot migrate off of the selected ***share***. This would leave files stranded on the share and cause the share removal to fail.

> **exclude *fileset*** (optional if you choose **close-file**, 1-64 characters) is a fileset to exclude from automatic closure. If a file in this fileset is open through CIFS, the rule places it on a retry queue instead of automatically closing it. If such a file remains open for the duration of the share removal, it never migrates off of the share and the share-removal fails.

async (optional if you specify a ***dest***) makes this command return immediately, rather than waiting for the share removal to finish.

Default(s) None.

Guidelines This command does not apply to direct volumes (see direct); only metadata-based managed volumes.

Before removing the share, resolve all dependencies on the share. You can use show policy to verify that no policy rules reference the share.

This is not recommended in a volume that supports snapshots. In those volumes, you should use a place-rule to drain the share first, wait until all retained snapshots have "aged out," and then run this command or remove-share nomigrate.

The back-end share must be online for this command to succeed. If the back-end share is offline, you can use remove-share offline to remove the share from the namespace.

You have the option to configure the share to retain copies of its migrated files, in case you need to recover them later: use migrate retain-files on the share before you run this command.

Clients are not impacted by the remove-share command. All files that migrate to the new back-end share remain in the same volume and namespace, so their file associations remain intact. Front-end clients see no change.

Guidelines: Share Removal in Multi-Protocol Namespaces

A directory in a multi-protocol (NFS and CIFS) share cannot migrate if it has an NFS-only name; the volume cannot access the directory through CIFS to replicate its CIFS attributes. You can access the volume through an NFS export and rename the NFS-only directory, or you can temporarily turn off strict-attribute-consistency at the ***dest*** share(s).

◆ Note

The migration may be slower in a multi-protocol volume. If it's possible for a legitimate directory name to collide with a filer-generated name (like "DIR~1") at one of the filers, the volume probes the target filer(s) for conflicting names.

Guidelines: No Migration Needed for Replica-Snap Shares

A volume uses the snapshots on a replica-snap share and ignores all other files and directories. This type of share does not have any files or directories to migrate, so the d***est*** share, **close-file**, and **exclude** options are all ignored.

Guidelines: Reports

This command creates two reports:

- drain_share_rule_for_share_*share-name_time*.rpt, and
- remove.*job-id.share-name.share-id*.rpt

Use show reports to view either of these reports.

Guidelines: Directory Tree Remains After Share Removal

When you remove a share from a namespace, the share's empty directory tree remains on the back-end filer. After the removal is complete, you can access the filer directly and remove the empty directory tree.

Sample

```
prtlndA# remove-share migrate nemed /vol44 rxbills globalmedco
close-file
```

removes the '/vol44~rxbills' share from the 'nemed' namespace, and migrates its files and any primary copies of its directories to the 'globalmedco' share farm. It also closes any files that are currently opened by CIFS clients, so that they will be drained from the share, too.

Related Commands

strict-attribute-consistency
remove-share nomigrate
remove-share offline
no share
wait-for remove

remove-share nomigrate

Purpose Use the remove-share nomigrate command to remove a share from a namespace without migrating any of its files. This causes all of the share's files (if there are any) to be removed from the volume.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax `remove-share nomigrate` ***namespace volume share*** `[`***dest*** `[async]]`

namespace (1 - 30 characters) is the name of the namespace.

volume (1 - 1024 characters) is the name of the volume.

share (1 - 64 characters) is the name of the share.

dest (not allowed for replica-snap shares, but required for all other share types; 1 - 64 characters) is the new destination share or share farm where the original share's primary copies of directories (if identified) are to be migrated. Use show global-config namespace for a list of shares and share farms in the namespace.

async (optional if you specify a ***dest***) makes this command return immediately, rather than waiting for the share removal to finish.

Default(s) None.

Guidelines This command does not apply to direct volumes (see direct); only metadata-based managed volumes.

Before removing the share, resolve all dependencies on the share. You can use show policy to verify that no policy rules reference the share.

This command deletes all of the share's files from the volume's metadata; from the client's perspective, these files disappear. Use this only for shares that have import failures or must keep their files after being removed from the namespace. For a share that is unreachable, use remove-share offline instead. To migrate the files to another share in the same managed volume, use remove-share migrate.

Guidelines: Migrating Master Directories

If the share contains any master directories in the volume (that is, the first-imported instance of a directory), they are copied to the ***dest*** share or share farm. This may fail if the remove process encounters any of the following issues:

- A master directory in a multi-protocol (NFS and CIFS) share cannot migrate if it has an NFS-only name; the volume cannot access the directory through CIFS to replicate its CIFS attributes. You can access the volume through an NFS export and rename the NFS-only directory, or you can temporarily turn off strict-attribute-consistency at the ***dest*** share(s).
- A master directory may not be able to migrate to a ***dest*** that supports CIFS if the destination filer is very low on free space. The CIFS ACL(s) from the master directory could exceed the free space on the destination. You can resolve this issue by choosing a ***dest*** with more free space behind it, increasing disk space on the destination filer(s), or using a place-rule (with a source share) to migrate files from the ***dest*** share to an emptier share in the same managed volume.

A volume uses the snapshots on a replica-snap share and ignores all other files and directories. This type of share does not have any master directories to migrate, so a ***dest*** share is not required.

Guidelines: Report

This command creates a report, "remove.*job-id*.*share-name*.*share-id*.rpt." Use show reports to view its contents.

Sample

```
bstnA# remove-share nomigrate nemed /vol33 oldScripts globalmedco
```

removes the '/vol33~oldScripts' share from the 'nemed' namespace and moves any master directories on that share to the 'globalmedco' share.

Related Commands

remove-share migrate
remove-share offline
no share
nsck ... destage
strict-attribute-consistency
wait-for remove

remove-share offline

Purpose Use the remove-share offline command to remove an offline, unreachable, or otherwise defunct back-end share from a namespace.

Modes priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax `remove-share offline` ***`namespace volume share dest`*** `[async]`

namespace (1 - 30 characters) is the name of the namespace.

volume (1 - 1024 characters) is the name of the volume.

share (1 - 64 characters) is the name of the share.

dest (1 - 64 characters) is the new destination share or share farm where the original share's master directories (if any) are to be migrated. A *master directory* is the first-imported instance of a directory in the managed volume. Use show global-config namespace for a list of shares and share farms in the namespace.

async (optional) makes this command return immediately, rather than waiting for the share removal to finish.

Important

*This operation migrates any master directories from the defunct share to the **dest**. Because the share is unreachable, this operation cannot probe those directories to find their file attributes and/or ACLs. The file attributes for the master directories are therefore set to zero at the **dest**.*

Default(s) None.

Guidelines This command does not apply to direct volumes (see direct); only metadata-based managed volumes.

Only use this command for a back-end share that is unreachable. If the share is still reachable, use remove-share migrate or remove-share nomigrate instead.

Before removing the share, resolve all dependencies on the share. You can use show policy to verify that no policy rules reference the share.

This command deletes all of the share's files from the volume's metadata; from the client's perspective, these files disappear. Given that the back-end filer was unreachable, those files were already inaccessible to clients.

This command creates a report, "remove.*job-id*.*share-name*.*share-id*.rpt." Use show reports to view its contents.

As mentioned above, this command migrates all master directories with 0 (zero) file attributes. Look in the remove report to find these directories: in the section below Prepare Start Time, look for `[AR]` next to each affected directory. You can access the volume through its VIP to manually reset the attributes of these directories.

Sample

```
bstnA# remove-share offline nemed /vol33 rxpresc globalmedco
```

removes the '/vol33~rxpresc' share from the 'nemed' namespace. If the share had any master directories, this command recreates them on the 'globalmedco' share with zeros for their file attributes and/or ACLs.

Related Commands

remove-share migrate
remove-share nomigrate
no share
nsck ... destage
strict-attribute-consistency
wait-for remove

show policy history

Purpose Use the show policy history command to view the policy events that have occurred in a volume, or the detailed events for a specific rule or share farm.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax **`show policy history namespace`** ***`namespace`*** **`volume`** ***`volume`*** **`[rule`** ***`rule-name`*****`]`**

namespace (1-30 characters) is the namespace.

volume (1-1024 characters) identifies the volume.

rule-name (optional, 1-1024 characters) narrows the scope to one rule or share farm in the volume. This shows detailed history for the rule or share farm.

Guidelines The output shows one row per policy event, with the date and time for each. The output contains up to 100 events.

The show policy command shows all namespace policies. Use show policy filesets to show all filesets, and use show schedule to show all schedules. The show policy queue command shows all files currently waiting to be migrated, if any.

Samples `bstnA# show policy history namespace insur volume /claims`

shows the policy history for the "insur~/claims" volume. See Figure 40.3 on page 40-23 for sample output.

`prtlndA# show policy history namespace nemed volume /acctShdw rule farm1`

on a different switch, shows policy history for a share farm. See Figure 40.4 on page 40-24 for sample output.

Related Commands show policy
show schedule
show policy filesets
show policy queue

Figure 40.3 *Sample Output: show policy history*

```
bstnA# show policy history namespace insur volume /claims

Namespace:        insur
Volume:           /claims
Rule
---------------------------------------
/claims

Time                 Message
-------------------  ------------------------------------------------------------
2009-09-23T05:35:16   Rule [insur:/claims:drain_share_rule_for_share_shr1-old]: Created.
2009-09-23T05:35:16   Rule [insur:/claims:drain_share_rule_for_share_shr1-old]: Enabled.
```

```
2009-09-23T05:35:26   Rule [insur:/claims:drain_share_rule_for_share_shr1-old]: Treewalk for share
[insur:/claims:shr1-old] starting on dir [/claims/].
2009-09-23T05:35:26   Rule [insur:/claims:drain_share_rule_for_share_shr1-old]: Starting scan.
2009-09-23T05:35:27   Rule [insur:/claims:drain_share_rule_for_share_shr1-old]: Scan complete.
2009-09-23T05:35:39   Rule [insur:/claims:drain_share_rule_for_share_shr1-old]: Disabled.
2009-09-23T05:55:13   Rule [insur:/claims:insurSV]: Created.
2009-09-23T05:55:14   Rule [insur:/claims:insurSV]: Enabled.
2009-09-23T05:55:14   Rule [insurSV]: Starting synchronization for volume [insur:/claims].
...
```

Figure 40.4 *Sample Output: show policy history ... rule*

```
prtlndA# show policy history namespace nemed volume /acctShdw rule farm1
Namespace:        nemed
Volume:           /acctShdw
Rule
---------------------------------------
farm1

Time                  Message
--------------------  ------------------------------------------------------------
2007-10-02T15:13:31   Enabled.
2007-10-02T15:13:45   The initial scan has been disabled for this rule.
```

show policy queue

Purpose Use the show policy queue command to see which files (if any) are currently queued for migration.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, and operator

Syntax **`show policy queue namespace`** ***`namespace`*** **`volume`** ***`volume`*** **`[auto-close]`**

namespace (1-30 characters) is the namespace.

volume (1-1024 characters) identifies the volume.

auto-close (optional) limits the output to files that have been auto-closed. (A file-placement rule can auto-close any file that is already opened through CIFS, and can hold it closed until the file is migrated. The migrate close-file command enables this feature.)

Guidelines This command is useful for diagnosing migration issues in a particular volume. It shows all the files that are currently waiting to be migrated, three lines per file. These lines contain the following fields:

File, on the top line, is the file being migrated. This is shown from the client's perspective.

Rule is the name of the rule or share farm that is migrating the file.

Target is the name of the target share or share farm. This is the name of the share from the namespace configuration; use the show global-config namespace command to show all shares in a namespace.

Status is the third line. This explains the current status of the file migration.

The output is typically an empty table, as migrations occur very quickly.

If the queue is very full because migrations are repeatedly failing, you can use the policy migrate-delay command to stop the policy engine from infinitely retrying its migrations. This causes the policy engine to remove all old entries from the queue. You can also use policy migrate-retry-delay to change the time between migration retries.

The cancel migration command clears the queue.

Sample `bstnA# show policy queue namespace medarcv volume /rcrds`

shows the policy queue for the "medarcv~/rcrds" volume. See Figure 40.5 for sample output.

Related Commands show policy history
show global-config **`namespace`**
cancel migration

Figure 40.5 *Sample Output: show policy queue*

```
bstnA# show policy queue namespace medarcv volume /rcrds
File
        Rule                  Target
Status
------------------------------------------------------------------------------
/rcrds/recoveryStats/file2.txt (Migrate)
        dailyArchive          bulk
        In Progress (Auto-Close Enabled) (Client Access Denied)
/rcrds/2004/planQ/fsousely_rcrd.dat (Migrate)
        dailyArchive          bulk
        Queued (Auto-Close Enabled)
/rcrds/2004/planQ/hblock_rcrd.dat (Migrate)
        dailyArchive          bulk
        Queued (Auto-Close Enabled)
/rcrds/2004/planQ/ihayes_rcrd.dat (Migrate)
        dailyArchive          bulk
        Queued (Auto-Close Enabled)
/rcrds/2004/planQ/jBradley_record.dat (Migrate)
        dailyArchive          bulk
        Queued (Auto-Close Enabled)
bstnA# ...
```

show statistics filer

Purpose Use the show statistics filer command to show the numbers of NFS RPCs and CIFS SMB commands from various software subsystems to one or more filers.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show statistics filer [nfs | cifs] [summary | request-detail |
response-detail]

show statistics filer ext-filer-name [share {shr | all}] [nfs |
cifs] [summary | request-detail | response-detail]
```

nfs | cifs (optional) limits the output to NFS statistics or CIFS/SMB2 statistics.

summary | request-detail | response-detail is an optional choice, where

summary shows the counts of each NFS and/or CIFS request, the counts of the corresponding responses, and the average round-trip times (RTTs). This is the default output if you select none of these options.

request-detail shows the counts of each NFS and/or CIFS request from each software component.

response-detail shows average RTT for each NFS and/or CIFS response to each software component.

ext-filer-name (optional, 1-64 characters) identifies the external filer by its configured name. Use show external-filer to display all configured external filers.

***shr* | all** is a required choice if you use the optional **share** keyword:

shr (1-64 characters) focuses on traffic to this specific back-end share. This is the name of the share at the back-end filer, not at the namespace.

all shows all shares on the filer. This creates a separate table for each share; if you omit the **share all** option, the output contains a single table showing the sums of all shares.

Guidelines The output from the simplest syntax, show statistics filer, shows 1-3 tables of statistics per external-filer: a table of NFS-traffic statistics, a table for CIFS traffic, and/or a table for SMB2 (an updated version of CIFS) traffic. Each row of the NFS table shows the total packets for a particular NFS RPC. Each row in the CIFS table shows the total packets for a given CIFS SMB command, and each row in the SMB2 table shows the packets for a given SMB2 command. All of these rows also display the average RTT for the RPC or SMB.

To clear the counters shown by this command, use clear statistics filer. To see client-side statistics, use show statistics global server. To sort these statistics by namespace or volume, use show statistics namespace ... summary, show statistics namespace ... request-detail, show statistics namespace ... response-detail, or show statistics namespace ... fastpath.

Guidelines: Output

This output shows a separate set of tables for every selected back-end filer:

Filer is the name of the external filer. As mentioned above, you can use the show external-filer *filer-name* command to see the IP address and/or SPN for *filer-name*.

All shares in filer appears if you did not use the **share** keyword to specify a single share. This indicates that the statistics are at the filer level; they are the sum of the statistics for all of the filers' imported shares. If you select a particular share, the following fields appear instead:

– NFS Export or

– CIFS Share is the name of the export or share at the filer. Both names appear if this is used as a multi-protocol share.

Reset... shows the last time these statistics were cleared. The software resets these statistics whenever someone uses the clear statistics filer command or reloads the chassis.

Guidelines: request-detail and response-detail Output

The request-detail flag focuses on request counts. This option breaks the request-count statistics down into the software subsystems that sent them:

- Data Plane is *fastpath* traffic through the network software. These packets are not handled by any volume-group software on the Control Plane.
- Ctrl Plane (short for Control Plane) is traffic that passes through the volume-group software for further processing.
- Policy shows the traffic that is initiated by the policy engine, such as migrations and shadow copies.
- Probe is the statistics from the health-checking subsystem in the namespace software. This subsystem sends periodic health-check probes to every back-end share that is connected to an ARX volume.

The response-detail flag focuses on average RTTs. This option also breaks the statistics down into the software subsystems that sent the RPCs and/or SMBs.

Guidelines: NFS Errors

NFS errors appear in a separate table under the NFS RPC statistics. The table contains the following counters:

- FormatError is the number of malformed NFS packets received from the filer(s).
- NetworkError counts the send and receive failures.
- ResourceError is the number of times that the ARX software ran out of resources. Contact F5 Support if you see these errors.
- RPCError counts any errors from the RPC layer.
- TimeoutError is the number of timeouts waiting for an NFS response.

Guidelines: CIFS Errors

Any CIFS errors appear in a similar table under the CIFS SMBs, with the following fields:

- FormatError is the number of malformed SMB (CIFS) packets received from the filer(s).
- NetworkError counts the send and receive failures for SMB traffic.
- ResourceError shows the number of times that the ARX software had insufficient resources to complete the transaction. Contact F5 Support if you see these errors.
- SignatureError counts any SMB packets with missing or incorrect signatures. You can use the cifs filer-signatures command to change the SMB-signing policy for the current namespace.
- SMBError is the number of SMB packets that failed with errors that are not directly related to any of the above errors.
- TimeoutError is the number of timeouts waiting for an SMB response.

Guidelines: SMB2 Errors

Any SMB2 errors appear in a table under the SMB2 statistics, with the following fields:

- FormatError is the number of malformed SMB2 packets received from clients.
- NetworkError counts the send and receive failures for SMB2 traffic.
- ResourceError shows the number of times that the ARX software had insufficient resources to complete the transaction. Contact F5 Support if you see these errors.
- SignatureError counts any SMB2 packets with missing or incorrect signatures. You can use the signatures command to change the signing policy for this global server's CIFS service.
- SMBError is the number of SMB2 packets that failed with errors that are not directly related to any of the above errors.
- TimeoutError is the number of timeouts waiting for an SMB2 response.

Samples

```
stkbrgA(cfg)# show statistics filer
```

summarizes all requests and responses to all back-end filers. See Figure 40.6 for sample output.

```
canbyA(cfg)# show statistics filer nfs reponse-detail
```

shows the response times for NFS RPCs to the filers behind "canbyA." See Figure 40.7 on page 40-34 for sample output.

Related Commands

clear statistics filer
show statistics global server
show external-filer
show statistics namespace ... summary
show statistics namespace ... request-detail
show statistics namespace ... response-detail

Figure 40.6 *Sample Output: show statistics filer*

```
stkbrgA(cfg)# show statistics filer

Request/response summary for shares by filer
collected at 02/13/2013 12:00:36 -0500

  Filer: fsvr_10 (192.168.66.134)
  All shares in filer
  Reset at 02/13/2013 11:54:16 -0500

  CIFS SMB                 Req Count      Resp Count   Avg RTT (uSec)
  -----------------  --------------  --------------  --------------
  CheckDir                         0               0              --
  Close                            7               7             456
  Create                           0               0              --
  CreateDir                        0               0              --
  Delete                           1               1             702
  DeleteDir                        0               0              --
  Echo                            19              19            2035
  FindClose2                       0               0              --
  Flush                            0               0              --
  LockingAndX                      0               0              --
  LogoffAndX                       4               4             525
  Negotiate                       13              13             624
  NTCancel                         0               0              --
  NTCreateAndX                     7               7            4764
  NTRename                         0               0              --
  NTTransact:
    Create                         0               0              --
    GetUserQuota                   0               0              --
    Ioctl                          0               0              --
    NotifyChange                   0               0              --
    QuerySecurity                  3               3             369
    SetSecurity                    0               0              --
  Open                             0               0              --
  OpenAndX                         0               0              --
  QueryInfo                        0               0              --
  QueryInfoDisk                    0               0              --
  Read                             0               0              --
  ReadAndX                         2               2            1111
  Rename                           0               0              --
  Seek                             0               0              --
  SessionSetupAndX                15              15            1321
  SetInfo                          0               0              --
  Transaction                      0               0              --
  Transaction2:
    CreateDir                      0               0              --
    FindFirst2                     5               5             646
    FindNext2                      0               0              --
    GetDFSReferral                 0               0              --
    QueryFSInfo                   40              40            2406
    QueryFileInfo                  0               0              --
    QueryPathInfo                  8               8             501
    SetFileInfo                    0               0              --
    SetPathInfo                    0               0              --
  TreeConnectAndX                 15              15             606
  TreeDisconnect                  11              11            2446
  Write                            0               0              --
  WriteAndX                        1               1             773
  UnknownSMB                       0               0              --
  -----------------  --------------  --------------  --------------
      Total/Net Avg:             151             151            1652
```

```
Errors                       Count
-----------------  --------------
FormatError                      0
NetworkError                     0
ResourceError                    0
SignatureError                   0
SMBError                         0
TimeoutError                     0
-----------------  --------------
            Total:               0


Filer: fsvr_11 (192.168.66.135)
All shares in filer
Reset at 02/13/2013 11:54:17 -0500

CIFS SMB                 Req Count      Resp Count   Avg RTT (uSec)
-----------------  --------------  --------------  --------------
CheckDir                         0               0              --
Close                           21              21             368
Create                           0               0              --
CreateDir                        0               0              --
Delete                           2               2             745
DeleteDir                       10              10             558
Echo                            38              38             641
FindClose2                       0               0              --
Flush                            0               0              --
LockingAndX                      0               0              --
LogoffAndX                       8               8             442
Negotiate                       19              19            1638
NTCancel                         0               0              --
NTCreateAndX                    21              21             849
NTRename                         0               0              --
NTTransact:
  Create                         0               0              --
  GetUserQuota                   0               0              --
  Ioctl                          0               0              --
  NotifyChange                   0               0              --
  QuerySecurity                  1               1             359
  SetSecurity                    0               0              --
Open                             0               0              --
OpenAndX                         0               0              --
QueryInfo                        0               0              --
QueryInfoDisk                    0               0              --
Read                             0               0              --
ReadAndX                         4               4             715
Rename                           0               0              --
Seek                             0               0              --
SessionSetupAndX                22              22            3073
SetInfo                          0               0              --
Transaction                      0               0              --
Transaction2:
  CreateDir                      0               0              --
  FindFirst2                    11              11             560
  FindNext2                      0               0              --
  GetDFSReferral                 0               0              --
  QueryFSInfo                   78              78             562
  QueryFileInfo                  0               0              --
  QueryPathInfo                 27              27             434
  SetFileInfo                    0               0              --
  SetPathInfo                    0               0              --
```

```
TreeConnectAndX                 22              22             401
TreeDisconnect                  16              16             460
Write                            0               0              --
WriteAndX                        2               2             698
UnknownSMB                       0               0              --
------------------  --------------  --------------  --------------
    Total/Net Avg:             302             302             800

Errors                       Count
------------------  --------------
FormatError                      0
NetworkError                     0
ResourceError                    0
SignatureError                   0
SMBError                         0
TimeoutError                     0
------------------  --------------
            Total:               0


Filer: fsvr_12 (192.168.66.148)
All shares in filer
Reset at 02/13/2013 11:54:18 -0500

CIFS SMB                Req Count      Resp Count   Avg RTT (uSec)
------------------  --------------  --------------  --------------
CheckDir                         0               0              --
Close                            5               5             843
Create                           0               0              --
CreateDir                        0               0              --
Delete                           1               1             660
DeleteDir                        0               0              --
Echo                            19              19             697
FindClose2                       0               0              --
Flush                            0               0              --
LockingAndX                      0               0              --
LogoffAndX                       0               0              --
Negotiate                        3               3             872
NTCancel                         0               0              --
NTCreateAndX                     5               5            4011
NTRename                         0               0              --
NTTransact:
  Create                         0               0              --
  GetUserQuota                   0               0              --
  Ioctl                          0               0              --
  NotifyChange                   0               0              --
  QuerySecurity                  1               1             336
  SetSecurity                    0               0              --
Open                             0               0              --
OpenAndX                         0               0              --
QueryInfo                        0               0              --
QueryInfoDisk                    0               0              --
Read                             0               0              --
ReadAndX                         2               2             649
Rename                           0               0              --
Seek                             0               0              --
SessionSetupAndX                 3               3            1929
SetInfo                          0               0              --
Transaction                      0               0              --
Transaction2:
  CreateDir                      0               0              --
  FindFirst2                     0               0              --
```

```
  FindNext2                     0               0              --
  GetDFSReferral                0               0              --
  QueryFSInfo                  40              40             622
  QueryFileInfo                 0               0              --
  QueryPathInfo                 0               0              --
  SetFileInfo                   0               0              --
  SetPathInfo                   0               0              --
TreeConnectAndX                 3               3             843
TreeDisconnect                  0               0              --
Write                           0               0              --
WriteAndX                       1               1             668
UnknownSMB                      0               0              --
------------------  --------------  --------------  --------------
    Total/Net Avg:             83              83             919

Errors                       Count
------------------  --------------
FormatError                     0
NetworkError                    0
ResourceError                   0
SignatureError                  0
SMBError                        0
TimeoutError                    0
------------------  --------------
            Total:              0


Filer: fsvr_14 (192.168.66.149)
All shares in filer
Reset at 02/13/2013 11:54:19 -0500

CIFS SMB                Req Count      Resp Count   Avg RTT (uSec)
------------------  --------------  --------------  --------------
CheckDir                        0               0              --
Close                           5               5             348
Create                          0               0              --
CreateDir                       0               0              --
Delete                          1               1             645
DeleteDir                       0               0              --
Echo                           19              19             549
FindClose2                      0               0              --
Flush                           0               0              --
LockingAndX                     0               0              --
LogoffAndX                      0               0              --
Negotiate                       3               3             650
NTCancel                        0               0              --
NTCreateAndX                    5               5             646
NTRename                        0               0              --
NTTransact:
  Create                        0               0              --
  GetUserQuota                  0               0              --
  Ioctl                         0               0              --
  NotifyChange                  0               0              --
  QuerySecurity                 1               1             311
  SetSecurity                   0               0              --
Open                            0               0              --
OpenAndX                        0               0              --
QueryInfo                       0               0              --
QueryInfoDisk                   0               0              --
Read                            0               0              --
ReadAndX                        2               2             509
Rename                          0               0              --
```

```
  Seek                            0               0              --
  SessionSetupAndX                3               3            3958
  SetInfo                         0               0              --
  Transaction                     0               0              --
  Transaction2:
    CreateDir                     0               0              --
    FindFirst2                    0               0              --
    FindNext2                     0               0              --
    GetDFSReferral                0               0              --
    QueryFSInfo                  40              40             581
    QueryFileInfo                 0               0              --
    QueryPathInfo                 0               0              --
    SetFileInfo                   0               0              --
    SetPathInfo                   0               0              --
  TreeConnectAndX                 3               3             940
  TreeDisconnect                  0               0              --
  Write                           0               0              --
  WriteAndX                       1               1            1187
  UnknownSMB                      0               0              --
  ------------------  --------------  --------------  --------------
      Total/Net Avg:             83              83             704

  Errors                       Count
  ------------------  --------------
  FormatError                      0
  NetworkError                     0
  ResourceError                    0
  SignatureError                   0
  SMBError                         0
  TimeoutError                     0
  ------------------  --------------
              Total:               0
```

Figure 40.7 *Sample Output: show statistics filer nfs response-detail*

```
canbyA(cfg)# show statistics filer nfs response-detail

Response-time details for shares by filer
collected at 11/12/2012 04:48:19 -0500

  Filer: solNFS (192.168.121.56)
  All shares in filer
  Reset at 11/12/2012 04:45:00 -0500

                            Average Round-Trip Times per Component (uSec)
  NFS RPC                Data Plane      Ctrl Plane          Policy           Probe
  ------------------  --------------  --------------  --------------  --------------
  Null                           --              --              --             310
  Getattr                       406             411            4472             295
  Setattr                     13602           12242           13554              --
  Lookup                        220             402            5302             359
  Access                        290              --              --              --
  Readlink                      209              --              --              --
  Read                           --              --           11900             646
  Write                          --              --              --          103745
  Create                         --           37850              --         1540799
  Mkdir                          --           47987              --              --
  Symlink                        --              --           69561              --
  Mknod                          --              --              --              --
  Remove                         --           33185              --           44895
  Rmdir                          --           22299              --              --
```

```
Rename                          --           31880             --             --
Link                            --              --             --             --
Readdir                         --             283             --             --
Readdirplus                     --             683             --             --
Fsstat                          --             651             --            276
Fsinfo                          --             458           4090            289
Pathconf                        --              --             --             --
Commit                          --              --           5710             --
StatFs                          --              --             --             --
------------------  -------------  -------------  -------------  -------------
     Net avg time:             325            7669          11499          67870

Errors                  Data Plane     Ctrl Plane         Policy          Probe
------------------  -------------  -------------  -------------  -------------
FormatError                      0              0              0              0
NetworkError                     0              0              0              0
ResourceError                    0              0              0              0
RPCError                         0              0              0              0
TimeoutError                     0              0              0              0
------------------  -------------  -------------  -------------  -------------
            Total:               0              0              0              0


Filer: oldNFS1 (192.168.121.58)
All shares in filer
Reset at 11/12/2012 04:45:01 -0500

                            Average Round-Trip Times per Component (uSec)
NFS RPC                 Data Plane     Ctrl Plane         Policy          Probe
------------------  -------------  -------------  -------------  -------------
Null                            --             --             --            279
Getattr                        287            589            246            264
Setattr                         --           5808          21090             --
Lookup                          --            780            932            266
Access                         272             --             --             --
Readlink                        --             --             --             --
Read                            --             --             --            267
Write                           --             --            925           8520
Create                          --             --           9748           3115
Mkdir                           --          17843           7092             --
Symlink                         --             --             --             --
Mknod                           --             --             --             --
Remove                          --            209             --           2948
Rmdir                           --             --             --             --
Rename                          --           4313          21232             --
Link                            --             --             --             --
Readdir                         --             --             --             --
Readdirplus                     --            842             --             --
Fsstat                          --            339            289            306
Fsinfo                          --            395            359            245
Pathconf                        --             --             --             --
Commit                          --             --         221112             --
StatFs                          --             --             --             --
------------------  -------------  -------------  -------------  -------------
     Net avg time:             286           1028           2445            835

Errors                  Data Plane     Ctrl Plane         Policy          Probe
------------------  -------------  -------------  -------------  -------------
FormatError                      0              0              0              0
NetworkError                     0              0              0              0
ResourceError                    0              0              0              0
RPCError                         0              0              0              0
```

```
TimeoutError                    0              0              0              0
------------------  -------------  -------------  -------------  -------------
           Total:               0              0              0              0


Filer: oldNFS2 (192.168.121.59)
All shares in filer
Reset at 11/12/2012 04:45:03 -0500

                          Average Round-Trip Times per Component (uSec)
NFS RPC                Data Plane     Ctrl Plane         Policy          Probe
------------------  -------------  -------------  -------------  -------------
Null                           --             --             --            213
Getattr                       201            339            529            206
Setattr                        --           2645          17729             --
Lookup                        189            258            594            207
Access                        195             --             --             --
Readlink                       --             --            309             --
Read                           --             --             --            256
Write                          --             --            943           7732
Create                         --             --          17549           2578
Mkdir                          --          57747           7443             --
Symlink                        --             --             --             --
Mknod                          --             --             --             --
Remove                         --           3132             --           3039
Rmdir                          --             --             --             --
Rename                         --             --          78431             --
Link                           --             --             --             --
Readdir                        --            251             --             --
Readdirplus                    --            347             --             --
Fsstat                         --            267            233            238
Fsinfo                         --            340            420            205
Pathconf                       --             --             --             --
Commit                         --             --          92816             --
StatFs                         --             --             --             --
------------------  -------------  -------------  -------------  -------------
     Net avg time:            199            343           2198            734

Errors                 Data Plane     Ctrl Plane         Policy          Probe
------------------  -------------  -------------  -------------  -------------
FormatError                    0              0              0              0
NetworkError                   0              0              0              0
ResourceError                  0              0              0              0
RPCError                       0              0              0              0
TimeoutError                   0              0              0              0
------------------  -------------  -------------  -------------  -------------
           Total:              0              0              0              0
```

show statistics metadata

Purpose Each managed volume stores *metadata* about the files that it keeps on its back-end shares. This metadata is stored on an external metadata share. The latency between the managed-volume software and this external metadata share must be small or the volume's performance suffers. Use the show statistics metadata command to measure the latency between namespace processes and their metadata storage.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show statistics metadata [namespace name [volume vol]]`

name (1-30 characters) specifies one namespace. You can use show namespace to get a list of all namespaces.

vol (optional; 1-1024 characters) specifies one managed volume.

Default(s) None

Guidelines Use this command to measure the performance of managed-volume software with respect to metadata storage. The output contains one table per volume, with each of the following fields:

Namespace,

Volume, and

Location identify a specific metadata share. The name of the back-end filer is its external-filer name as defined on the ARX; you can use show external-filer to find its IP address, SPN, and configuration information.

Last Reset Time shows the last time the volume software started or someone cleared these statistics with clear statistics metadata.

The next group of fields describe statistics for metadata-read operations:

- Read Operations counts the number of reads since the above last-reset time.
- Read Response Time is the average response time for the read operations,
- Read Bytes is the total bytes since the Last Reset Time,
- Read Rate shows the running average in Megabytes per second,
- Read Errors counts the total read errors since the volume group started,
- Read Operations in Progress is the number of read operations currently underway, and
- Last Read Operation is a date stamp for the most-recent read operation.

The above statistics are then repeated for all write operations.

Sample

```
bstnA# show statistics metadata
```

measures and displays the latency between all volumes and their metadata storage. See Figure 40.8 on page 40-38 for sample output.

Related Commands show namespace
clear statistics metadata

Figure 40.8 *Sample Output: show statistics metadata*

```
bstnA> show statistics metadata

  Namespace: wwmed
  Volume:    /acct
  Location:  nas1 /vol/vol2/meta1

      Last Reset Time:              11/16/2012 00:52:17 -0500

      Read Operations:              20
      Read Response Time:           3.20 usec/op
      Read Bytes:                   60416
      Read Rate:                    944.00 MB/s
      Read Errors:                  0
      Read Operations In Progress:  0
      Last Read Operation:          11/16/2012 00:52:17 -0500

      Write Operations:             6955
      Write Response Time:          59.14 usec/op
      Write Bytes:                  6331392
      Write Rate:                   15.39 MB/s
      Write Errors:                 0
      Write Operations In Progress: 0
      Last Write Operation:         11/16/2012 01:36:44 -0500

  Namespace: medarcv
  Volume:    /lab_equipment
  Location:  nas1 /vol/vol2/meta6

      Last Reset Time:              11/16/2012 00:53:08 -0500

      Read Operations:              28
      Read Response Time:           2.79 usec/op
      Read Bytes:                   57344
      Read Rate:                    735.18 MB/s
      Read Errors:                  0
      Read Operations In Progress:  0
      Last Read Operation:          11/16/2012 01:07:33 -0500

      Write Operations:             73
      Write Response Time:          358.95 usec/op
      Write Bytes:                  256000
      Write Rate:                   9.77 MB/s
      Write Errors:                 0
      Write Operations In Progress: 0
      Last Write Operation:         11/16/2012 01:39:03 -0500

  Namespace: medarcv
  Volume:    /rcrds
  Location:  nas1 /vol/vol2/meta3

      Last Reset Time:              11/16/2012 00:52:55 -0500

      Read Operations:              28
      Read Response Time:           2.89 usec/op
      Read Bytes:                   57344
```

```
    Read Rate:                    707.95 MB/s
    Read Errors:                  0
    Read Operations In Progress:  0
    Last Read Operation:          11/16/2012 01:12:32 -0500

    Write Operations:             124
    Write Response Time:          237.94 usec/op
    Write Bytes:                  407552
    Write Rate:                   13.81 MB/s
    Write Errors:                 0
    Write Operations In Progress: 0
    Last Write Operation:         11/16/2012 01:39:02 -0500

Namespace: insur
Volume:    /claims
Location:  nas1 /vol/vol2/meta2

    Last Reset Time:              11/16/2012 01:12:51 -0500

    Read Operations:              26
    Read Response Time:           3.46 usec/op
    Read Bytes:                   77824
    Read Rate:                    864.71 MB/s
    Read Errors:                  0
    Read Operations In Progress:  0
    Last Read Operation:          11/16/2012 01:17:30 -0500

    Write Operations:             86
    Write Response Time:          239.01 usec/op
    Write Bytes:                  330752
    Write Rate:                   16.09 MB/s
    Write Errors:                 0
    Write Operations In Progress: 0
    Last Write Operation:         11/16/2012 01:29:22 -0500

Namespace: labResearch
Volume:    /labMice
Location:  nas1 /vol/vol2/meta8

    Last Reset Time:              11/16/2012 01:17:18 -0500

    Read Operations:              18
    Read Response Time:           4.61 usec/op
    Read Bytes:                   52224
    Read Rate:                    629.20 MB/s
    Read Errors:                  0
    Read Operations In Progress:  0
    Last Read Operation:          11/16/2012 01:17:18 -0500

    Write Operations:             39
    Write Response Time:          252.51 usec/op
    Write Bytes:                  138240
    Write Rate:                   14.04 MB/s
    Write Errors:                 0
    Write Operations In Progress: 0
    Last Write Operation:         11/16/2012 01:18:17 -0500

Namespace: labResearch
Volume:    /labRats
Location:  nas1 /vol/vol2/meta9

    Last Reset Time:              11/16/2012 01:17:25 -0500
```

```
Read Operations:              18
Read Response Time:           3.50 usec/op
Read Bytes:                   52224
Read Rate:                    828.95 MB/s
Read Errors:                  0
Read Operations In Progress:  0
Last Read Operation:          11/16/2012 01:17:25 -0500

Write Operations:             36
Write Response Time:          273.81 usec/op
Write Bytes:                  125952
Write Rate:                   12.78 MB/s
Write Errors:                 0
Write Operations In Progress: 0
Last Write Operation:         11/16/2012 01:18:18 -0500
```

show statistics metalog

Purpose Namespace software records important information in its *metalog*, possibly to be used later for a failure recovery. To ensure proper volume performance, this latency should be low. Use the show statistics metalog command to measure the latency between namespace processes and their metalog storage.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show statistics metalog [volume-group id]`

id (optional, 1-255) identifies a particular volume group. If you omit the number, this command displays all volume groups.

Default(s) None

Guidelines To ensure proper volume performance, this latency and the latency to the backup peer (measurable with show redundancy metalog) should be as low as possible. Use this command to measure the performance of namespace software with respect to metalog storage. The output contains one table per volume group, with each of the following fields:

Last Reset Time shows the last time the volume-group software started, or the last time someone used clear statistics metalog on the volume group.

Read Operations counts the number of reads since the above time.

Read Response Time is the average response time for the read operations,

Read Bytes is the total bytes read since the Last Reset Time,

Read Rate shows the running average round-trip time for read operations, in Megabytes per second,

Read Errors counts the total read errors since the volume group started,

Read Operations in Progress is the number of read operations currently underway, and

Last Read Operation is a date stamp for the most-recent read operation.

The above statistics are then repeated for all write operations.

The ARX-1500 and ARX-2500 store their metalog data on their internal disks. On those platforms, you can use the show metalog usage command to see the metalog-usage statistics.

The ARX-VE stores its metalog data on an external data store. On the ARX-VE, you can use the probe metalog latency command to test the latency between the ARX and this external storage.

The show redundancy metalog command shows the latency for resilvering (copying) metalog data to the backup peer. If that latency is too high, slower volume processing may be noticeable by volume clients.

Samples

```
bstnA# show statistics metalog
```

measures and displays the latency between all volume groups and their metalog storage. See Figure 40.9 on page 40-42 for sample output.

```
bstnA# show statistics metalog volume-group 1
```

focuses on volume group 1. See Figure 40.10 on page 40-44 for sample output.

Related Commands

clear statistics metalog
show redundancy metalog
show metalog usage
probe metalog latency

Figure 40.9 *Sample Output: show statistics metalog*

```
bstnA> show statistics metalog

  Volume Group: 1

      Last Reset Time:              07/07/2011 01:18:13 -0400

      Read Operations:              5
      Read Response Time:           5.80 usec/op
      Read Bytes:                   0
      Read Rate:                    0.00 MB/s
      Read Errors:                  0
      Read Operations In Progress:  0
      Last Read Operation:          07/07/2011 01:18:13 -0400

      Write Operations:             5
      Write Response Time:          60.20 usec/op
      Write Bytes:                  484
      Write Rate:                   1.61 MB/s
      Write Errors:                 0
      Write Operations In Progress: 0
      Last Write Operation:         07/07/2011 01:18:13 -0400

  Volume Group: 2

      Last Reset Time:              07/07/2011 01:18:48 -0400

      Read Operations:              159
      Read Response Time:           5.05 usec/op
      Read Bytes:                   32976
      Read Rate:                    41.07 MB/s
      Read Errors:                  0
      Read Operations In Progress:  0
      Last Read Operation:          07/07/2011 01:54:37 -0400

      Write Operations:             77716
      Write Response Time:          49.55 usec/op
      Write Bytes:                  26504006
      Write Rate:                   6.88 MB/s
      Write Errors:                 0
      Write Operations In Progress: 0
      Last Write Operation:         07/07/2011 01:56:55 -0400
```

```
Volume Group: 9

    Last Reset Time:              07/07/2011 01:58:51 -0400

    Read Operations:              0
    Read Response Time:           0.00 usec/op
    Read Bytes:                   0
    Read Rate:                    0.00 MB/s
    Read Errors:                  0
    Read Operations In Progress:  0
    Last Read Operation:          None

    Write Operations:             0
    Write Response Time:          0.00 usec/op
    Write Bytes:                  0
    Write Rate:                   0.00 MB/s
    Write Errors:                 0
    Write Operations In Progress: 0
    Last Write Operation:         None

Volume Group: 10

    Last Reset Time:              07/07/2011 01:58:51 -0400

    Read Operations:              0
    Read Response Time:           0.00 usec/op
    Read Bytes:                   0
    Read Rate:                    0.00 MB/s
    Read Errors:                  0
    Read Operations In Progress:  0
    Last Read Operation:          None

    Write Operations:             0
    Write Response Time:          0.00 usec/op
    Write Bytes:                  0
    Write Rate:                   0.00 MB/s
    Write Errors:                 0
    Write Operations In Progress: 0
    Last Write Operation:         None
```

Figure 40.10 *Sample Output: show statistics metalog volume-group 1*

```
bstnA> show statistics metalog volume-group 1

  Volume Group: 1

      Last Reset Time:              07/07/2011 01:18:13 -0400

      Read Operations:              5
      Read Response Time:           5.80 usec/op
      Read Bytes:                   0
      Read Rate:                    0.00 MB/s
      Read Errors:                  0
      Read Operations In Progress:  0
      Last Read Operation:          07/07/2011 01:18:13 -0400

      Write Operations:             5
      Write Response Time:          60.20 usec/op
      Write Bytes:                  484
      Write Rate:                   1.61 MB/s
      Write Errors:                 0
      Write Operations In Progress: 0
      Last Write Operation:         07/07/2011 01:18:13 -0400
```

show statistics migration

Purpose Use the show statistics migration command to show the file-migration activity on the current ARX.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show statistics migration [namespace ns [volume vol]]
show statistics migration activity [namespace ns [volume vol]]
show statistics migration history [namespace ns [volume vol]]
```

ns (optional; 1-30 characters) specifies one namespace.

vol (optional; 1-1024 characters) specifies one managed volume.

activity (optional) shows all migrations that are currently in progress for each of the chosen volumes.

history (optional) shows the last 100 migrations that occurred in each of the chosen volumes.

Guidelines The policy engine enforces a managed volume's rules by migrating files from one back-end filer to another. A place-rule causes file migrations, as does the auto-migrate command in a share-farm. Use this command to view in-flight migrations, migration history, and some migration-related statistics.

The output is divided into namespaces, volume groups, and then volumes. All volumes appear in the output, whether or not they have any migration statistics.

Namespace identifies a namespace.

Volume Group is one volume-group in the above namespace.

Volume shows one volume in the above volume group.

Below each volume are the statistics requested, or configuration data if you did not select **activity** or **history**. A volume without any policy associated with it does not have anything below its name.

Guidelines: Default Output The summary view contains a table of namespace-migration statistics followed by a table of per-volume statistics.

Last Reset Time is the last time (in UTC, not local time) that someone ran the clear statistics migration command or rebooted the chassis.

Files Migrated is the total files migrated in the namespace.

Failed Migrations shows how many migrations failed.

Migration Time is the time required to migrate all of the Files Migrated.

Data Migrated counts the total size of files successfully migrated.

Active Migration Count is the number of migrations that are currently underway. You can see the details of these migrations by re-running this command with the **activity** keyword.

Average Data Rate is the average rate of transfer for all migrations behind the namespace.

Guidelines: Default Output (Cont.)

Largest File Size,

Smallest File Size and

Average File Size are high, low, and average samplings of the files migrated in the namespace.

Guidelines: Activity View

If you use the activity option, the output contains a Current Migration Activity table for each volume instead of the migration statistics. This table lists the files that the policy engine is currently migrating, one file per row.

Guidelines: History View

The history option shows two tables per volume with the most-recent migrations (up to 100 total):

Migration History Failure Activity shows each failed migration, one per row. Each row contains four fields.

- Timestamp shows when the migration failed (in UTC, not local time),
- Source is the back-end filer and path where the file started.
- Destination is the back-end filer and path where the policy engine attempted (and failed) to migrate the file.
- Reason is the reason for the failure.

Migration History Success Activity lists all successful migrations. Each row contains the following fields.

- Timestamp shows when the migration occurred (in UTC, not local time),
- Source is the back-end filer and path where the file started.
- Destination is the back-end filer and path where the policy engine migrated the file.

Samples

`bstnA(cfg)# show statistics migration`

shows high-level migration statistics for the "bstnA" chassis. See Figure 40.11 for sample output.

`bstnA(cfg)# show statistics migration history namespace medarcv`

shows the detailed migration history for the "medarcv" namespace. See Figure 40.12 on page 40-48 for sample output.

Related Commands

clear statistics migration
place-rule
auto-migrate

Figure 40.11 *Sample Output: show statistics migration*

```
bstnA(cfg)# show statistics migration

  Namespace:                                    medco
--------------------------------------------------------------------------------
  Volume Group:                                 9
    Volume:                                     /vol
  Namespace:                                    wwmed
--------------------------------------------------------------------------------
  Volume Group:                                 1
    Volume:                                     /acct
      Last Reset Time:                          07/10/2011 01:39:51 -0400
      Files Migrated:                           75
```

```
      Failed Migrations:                          0
      Migration Time:                             30.139 seconds
      Data Migrated:                              120 MB
      Active Migration Count:                     0
      Average Data Rate:                          4.1 MB/s
      Largest File Size:                          14 MB
      Smallest File Size:                         2.6 kB
      Average File Size:                          1.6 MB

  Namespace:                                      medarcv
-------------------------------------------------------------------------------
  Volume Group:                                   2
    Volume:                                       /test_results
    Volume:                                       /acopia$ns3
    Volume:                                       /rcrds
      Last Reset Time:                            07/10/2011 01:40:27 -0400
      Files Migrated:                             366
      Failed Migrations:                          0
      Migration Time:                             7.059 seconds
      Data Migrated:                              22 MB
      Active Migration Count:                     0
      Average Data Rate:                          3.2 MB/s
      Largest File Size:                          131 kB
      Smallest File Size:                         0 B
      Average File Size:                          62 kB

    Volume:                                       /lab_equipment
      Last Reset Time:                            07/10/2011 01:40:39 -0400
      Files Migrated:                             0
      Failed Migrations:                          0
      Migration Time:                             0 seconds
      Data Migrated:                              0 B
      Active Migration Count:                     0
      Average Data Rate:                          0 B/s
      Largest File Size:                          N/A
      Smallest File Size:                         N/A
      Average File Size:                          N/A

  Namespace:                                      insur
-------------------------------------------------------------------------------
  Volume Group:                                   10
    Volume:                                       /claims
      Last Reset Time:                            07/10/2011 01:59:10 -0400
      Files Migrated:                             0
      Failed Migrations:                          0
      Migration Time:                             0 seconds
      Data Migrated:                              0 B
      Active Migration Count:                     0
      Average Data Rate:                          0 B/s
      Largest File Size:                          N/A
      Smallest File Size:                         N/A
      Average File Size:                          N/A

    Volume:                                       /acopia$ns4
```

Figure 40.12 *Sample Output: show statistics migration history...*

```
bstnA(cfg)# show statistics migration history namespace medarcv

  Namespace:                                    medarcv
--------------------------------------------------------------------------------
  Volume Group:                                 2
    Volume:                                     /test_results
    Volume:                                     /acopia$ns3
    Volume:                                     /rcrds

Migration History Failure Activity:
Timestamp            Source                        Destination                          Reason
--------------------------------------------------------------------------------


Migration History Success Activity:
Timestamp             Source                          Destination
--------------------------------------------------------------------------------
07/10/2011 02:31:41 -0400 \\192.168.25.20\histories\VIP_wing\fractures\file8.txt ->
\\192.168.25.27\bulkstorage\VIP_wing\fractures\file8.txt
07/10/2011 02:31:41 -0400 \\192.168.25.29\prescriptions\VIP_wing\fractures\benchNotes.doc ->
\\192.168.25.27\bulkstorage\VIP_wing\fractures\benchNotes.doc
07/10/2011 02:31:41 -0400 \\192.168.25.20\histories\VIP_wing\fractures\labBudget.xls ->
\\192.168.25.27\bulkstorage\VIP_wing\fractures\labBudget.xls
07/10/2011 02:31:41 -0400 \\192.168.25.29\prescriptions\VIP_wing\fractures\examSchedule.doc ->
\\192.168.25.27\bulkstorage\VIP_wing\fractures\examSchedule.doc
07/10/2011 02:31:41 -0400 \\192.168.25.29\prescriptions\VIP_wing\fractures\recoveryStats.xls ->
\\192.168.25.27\bulkstorage\VIP_wing\fractures\recoveryStats.xls
07/10/2011 02:31:41 -0400 \\192.168.25.29\prescriptions\VIP_wing\fractures\paid_holidays.doc ->
\\192.168.25.27\bulkstorage\VIP_wing\fractures\paid_holidays.doc
07/10/2011 02:31:41 -0400 \\192.168.25.20\histories\VIP_wing\fractures\orSched.doc ->
\\192.168.25.27\bulkstorage\VIP_wing\fractures\orSched.doc
07/10/2011 02:31:41 -0400 \\192.168.25.20\histories\VIP_wing\fractures\stdProcedure.doc ->
\\192.168.25.27\bulkstorage\VIP_wing\fractures\stdProcedure.doc
07/10/2011 02:31:41 -0400 \\192.168.25.20\histories\VIP_wing\fractures\memo.doc ->
\\192.168.25.27\bulkstorage\VIP_wing\fractures\memo.doc
...
07/10/2011 02:31:15 -0400 \\192.168.25.29\prescriptions\2001\planA\t_roosevelt.dat ->
\\192.168.25.27\bulkstorage\2001\planA\t_roosevelt.dat
07/10/2011 02:31:15 -0400 \\192.168.25.29\prescriptions\2001\planA\t_jefferson.dat ->
\\192.168.25.27\bulkstorage\2001\planA\t_jefferson.dat
07/10/2011 02:31:15 -0400 \\192.168.25.29\prescriptions\2001\planA\s_adams.dat ->
\\192.168.25.27\bulkstorage\2001\planA\s_adams.dat
07/10/2011 02:31:15 -0400 \\192.168.25.29\prescriptions\2001\planA\j_adams.dat ->
\\192.168.25.27\bulkstorage\2001\planA\j_adams.dat

    Volume:                                     /lab_equipment

Migration History Failure Activity:
Timestamp            Source                        Destination                          Reason
--------------------------------------------------------------------------------


Migration History Success Activity:
Timestamp             Source                          Destination
--------------------------------------------------------------------------------
```

show statistics namespace ... fastpath

Purpose Use the show statistics namespace ... fastpath command to show raw read/write statistics between a namespace and its back-end shares.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show statistics namespace ns [volume vol [share {shr | all}]]
fastpath
```

ns (1-30 characters) identifies the namespace.

vol (optional; 1-1024 characters) specifies one volume in the namespace.

shr **|** **all** is a required choice if you use the optional **share** keyword:

shr (1-64 characters) specifies one share in the volume.

all shows all shares in the volume. This creates a separate table for each share; if you omit the **share all** option, the output contains a single table showing the sums of all shares.

fastpath specifies that you want the raw statistics for the chosen volumes, as opposed to NFS and/or CIFS statistics (see show statistics filer). These statistics are kept by the fastpath processes on the NSM (or Data Plane).

Guidelines This command focuses on the fastpath, or data-plane, communication with back-end filers. These are raw byte counts for reads and writes, and the average packet latency. To see the counts of various CIFS and/or NFS requests to the namespace's filers, you can use the show statistics namespace ... request-detail command. To see the round-trip times (RTTs) for all of these requests, you can use show statistics namespace ... response-detail. For a summary, showing request and response counts together with RTTs, use show statistics namespace ... summary. The output of all three commands includes the counts of various errors. As with this command, you can see these statistics summarized for an entire namespace, focused on a single volume, or focused on one share.

Guidelines: Output This output shows the raw byte counts and throughput between the chosen namespace, volume, or share and its back-end filer(s). Each volume has a separate table of statistics with the following fields:

Namespace and

Volume identify the volume.

All shares in volume appears if you did not use the **share** keyword to specify a single share. This indicates that the statistics are at a volume level; they are the sum of the fastpath (data-plane) statistics for all of the volume's shares. If you select a particular share, the following fields appear instead:

– Share identifies the share, as it is named in the ARX configuration. The show global-config namespace *name* command shows the share names in the ARX configuration.

Guidelines: Output (Cont.)

- Filer is the name of the external filer that hosts the share. Again, this is the name of the filer in the ARX configuration; you can use the show external-filer *filer-name* command to see the IP address and/or SPN for *filer-name*.
- NFS Export or
- CIFS Share is the name of the export or share at the filer. Both names appear if this is used as a multi-protocol share.

Reset... shows the last time these statistics were cleared. The software resets these statistics whenever someone reloads the chassis or uses clear statistics filer on the share.

Bytes Read,

Read Calls,

Bytes Written, and

Write Calls are all measured since the share was last enabled. Use the enable (gbl-ns-vol-shr) command to enable a namespace share.

Other Calls is the sum of other Remote Procedure Calls (RPCs) including statistics operations and file opens.

All Calls is the sum of reads, writes, and other RPCs.

Average Latency is the average packet-round-trip time between the ARX and the share.

Guidelines: Clearing Statistics

You can use the clear statistics filer command to clear the statistics from all variations of show statistics namespace. If you use that command to clear the statistics from one share in a volume but not another, the volume-level statistics in this command are inconsistent; one share may have statistics for 2 months and the other share may have statistics for only 5 minutes. To confirm that the volume has consistent statistics, run this command on each of the volume's shares and verify that the Reset... field has the same date and time for each of them.

Samples

```
bstnA> show statistics namespace wwmed fastpath
```

shows fastpath statistics for the namespace, "wwmed." See Figure 40.13 on page 40-51 for sample output.

```
bstnA> show statistics namespace medarcv volume /rcrds fastpath
```

shows statistics for the volume, "medarcv~/rcrds." See Figure 40.14 on page 40-51 for sample output.

Related Commands

clear statistics filer
enable (gbl-ns-vol-shr)
show global-config **namespace**
show external-filer
show statistics namespace ... request-detail
show statistics namespace ... response-detail
show statistics namespace ... summary

Figure 40.13 *Sample Output: show statistics namespace ... fastpath*

```
bstnA> show statistics namespace wwmed fastpath

Fastpath summary statistics for shares by namespace and volume
collected at 11/16/2012 01:43:36 -0500

  Namespace: wwmed
  Volume: /acct
  All shares in volume

    Bytes read:                258,786,613
    Read calls:                     34,869
    Bytes written:             441,038,367
    Write calls:                    60,269
    Other calls:                   442,726
    All calls:                     537,864
    Average latency:              525 uSec
```

Figure 40.14 *Sample Output: show statistics namespace ... volume ... fastpath*

```
bstnA> show statistics namespace medarcv volume /rcrds fastpath

Fastpath summary statistics for shares by namespace and volume
collected at 11/16/2012 01:43:43 -0500

  Namespace: medarcv
  Volume: /rcrds
  All shares in volume

    Bytes read:                   513,948
    Read calls:                        16
    Bytes written:              1,764,404
    Write calls:                      112
    Other calls:                      152
    All calls:                        280
    Average latency:           169796 uSec
```

show statistics namespace ... request-detail

Purpose Use the show statistics namespace ... request-detail command to show request counts between a namespace and its back-end shares. The request counts are broken down by request type: the output shows separate counters for each CIFS (SMB) procedure call and/or NFS RPC.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show statistics namespace ns [volume vol [share {shr | all}]]
request-detail
```

ns (1-30 characters) identifies the namespace.

vol (optional; 1-1024 characters) specifies one volume in the namespace.

shr **|** **all** is a required choice if you use the optional **share** keyword:

> ***shr*** (1-64 characters) specifies one share in the volume.
>
> **all** shows all shares in the volume. This creates a separate table for each share; if you omit the **share all** option, the output contains a single table showing the sums of all shares.

request-detail causes the output to focus on request counts from ARX software to its back-end shares. This also shows error counts.

Guidelines This command counts the requests from the chosen namespace, volume or share to its back-end filers. These are cross-referenced by request format and by the software component that made the request. To see the round-trip times (RTTs) for all of these requests, you can use show statistics namespace ... response-detail. For a summary, showing request and response counts together with RTTs, use show statistics namespace ... summary. The output of all three commands includes the counts of various errors. To see the raw read/write statistics for the same namespace, volume, or share, you can use the show statistics namespace ... fastpath command. As with this command, you can see these statistics summarized for an entire namespace, focused on a single volume, or focused on one share.

Guidelines: Output The output shows the counts of specific request types from each component of namespace software. Each selected volume has a separate table of statistics with the following fields:

> Namespace and
>
> Volume identify the volume.
>
> All shares in volume appears if you did not use the **share** keyword to specify a single share. This indicates that the statistics are at a volume level; they are the sum of the statistics for all of the volume's shares. If you select a particular share, the following fields appear instead:
>
> – Share identifies the share, as it is named in the ARX configuration. The show global-config namespace *name* command shows the share names in the ARX configuration.

Guidelines: Output (Cont.)

- Filer is the name of the external filer that hosts the share. Again, this is the name of the filer in the ARX configuration; you can use the show external-filer *filer-name* command to see the IP address and/or SPN for *filer-name*.
- NFS Export or
- CIFS Share is the name of the export or share at the filer. Both names appear if this is used as a multi-protocol share.

Reset... shows the last time these statistics were cleared. The software resets these statistics whenever someone reloads the chassis or uses clear statistics filer on the share.

The remainder of the output are up to four tables of request counters: NFS RPCs, NFS Errors, CIFS SMBs, and/or CIFS Errors. These tables all have request (or error) types with four columns:

- Data Plane is *fastpath* traffic from the network software.
- Ctrl Plane (short for Control Plane) is traffic from volume-group software.
- Policy shows the requests from the policy engine, which performs migrations, shadow copies, and other policy-related functions.
- Probe is the requests from the health-checking subsystem in the namespace software. This subsystem sends periodic requests to every back-end share that is connected to an ARX volume.

Guidelines: NFS RPC Table

NFS RPC - Each row of the NFS table shows the total counts for a particular NFS-RPC request. Each column shows the source of the RPC request, as described above.

Guidelines: NFS Error Table

Errors appear in a separate table under the NFS statistics. The table contains the following counters, with the same columns as the RPC table:

- FormatError is the number of malformed NFS packets received from the filer(s).
- NetworkError counts the send and receive failures.
- ResourceError is the number of times that the ARX software ran out of resources. Contact F5 Support if you see these errors.
- RPCError counts any errors from the RPC layer.
- TimeoutError is the number of timeouts waiting for an NFS response.

Guidelines: CIFS SMB Table

CIFS SMB - Each row in the CIFS table shows the total counts for a given CIFS request. This has the same columns as the NFS tables above.

Guidelines: CIFS Error Table

Errors appear in a similar table under the CIFS statistics, with the following fields:

- FormatError is the number of malformed SMB (CIFS) packets received from the filer(s).
- NetworkError counts the send and receive failures for SMB traffic.
- ResourceError shows the number of times that the ARX software had insufficient resources to complete the transaction. Contact F5 Support if you see these errors.
- SignatureError counts any SMB packets with missing or incorrect signatures. You can use the cifs filer-signatures command to change the SMB-signing policy for the current namespace.
- SMBError is the number of SMB packets that failed with errors that are not directly related to any of the above errors.
- TimeoutError is the number of timeouts waiting for an SMB response.

Guidelines: SMB2 Table

SMB2 - Each row in the SMB2 table shows the total counts for a given SMB2 request. You can use the smb2-allowed command to enable SMB2 (an updated version of CIFS) for the ARX. This has the same columns as the NFS and CIFS tables above.

Guidelines: SMB2 Error Table

Errors appear in a similar table under the SMB2 statistics, with the following fields:

- FormatError is the number of malformed SMB2 packets received from clients.
- NetworkError counts the send and receive failures for SMB2 traffic.
- ResourceError shows the number of times that the ARX software had insufficient resources to complete the transaction. Contact F5 Support if you see these errors.
- SignatureError counts any SMB2 packets with missing or incorrect signatures. You can use the cifs filer-signatures command to change the SMB-signing policy for the current namespace.
- SMBError is the number of SMB2 packets that failed with errors that are not directly related to any of the above errors.
- TimeoutError is the number of timeouts waiting for an SMB2 response.

Guidelines: Clearing Statistics

You can use the clear statistics filer command to clear the statistics from all variations of show statistics namespace. If you use that command to clear the statistics from one share in a volume but not another, the volume-level statistics in this command are inconsistent; one share may have statistics for 2 months and the other share may have statistics for only 5 minutes. To confirm that the volume has consistent statistics, run this command on each of the volume's shares and verify that the Reset... field has the same date and time for each of them.

Samples

`stoweA> show statistics namespace lodges request-detail`

shows request statistics for the "lodges" namespace. See Figure 40.15 on page 40-55 for sample output.

`bstnA> show statistics namespace wwmed volume /acct request-detail`

shows request statistics for the "wwmed~/acct" volume. See Figure 40.16 on page 40-56 for sample output.

`bstnA> show statistics namespace wwmed volume /acct share all request-detail`

shows request statistics for the same volume, but for each individual share in the volume instead of the summed statistics. See Figure 40.17 on page 40-57 for sample output.

Related Commands

clear statistics filer
show global-config **namespace**
show external-filer
show statistics namespace ... fastpath
show statistics namespace ... response-detail
show statistics namespace ... summary

Figure 40.15 *Sample Output: show statistics namespace ... request-detail*

```
stoweA> show statistics namespace lodges request-detail

Request-count details for shares by namespace and volume
collected at 02/13/2013 11:18:10 -0500

  Namespace: lodges
  Volume: /skiPatrol
  All shares in volume
  Reset at 02/13/2013 11:15:06 -0500

                                  Request Counts per Component
  CIFS SMB                Data Plane      Ctrl Plane          Policy           Probe
  ------------------  --------------  --------------  --------------  --------------
  CheckDir                         0               0               0               0
  Close                            0              13               0               0
  Create                           0               0               0               0
  CreateDir                        0               0               0               0
  Delete                           0               0               0               0
  DeleteDir                        0               5               0               0
  Echo                             0               0               0              27
  FindClose2                       0               0               0               0
  Flush                            0               0               0               0
  LockingAndX                      0               0               0               0
  LogoffAndX                       0              10               0               0
  Negotiate                        1              14               0               3
  NTCancel                         0               0               0               0
  NTCreateAndX                     0              13               0               0
  NTRename                         0               0               0               0
  NTTransact:
    Create                         0               0               0               0
    GetUserQuota                   0               0               0               0
    Ioctl                          0               0               0               0
    NotifyChange                   0               0               0               0
```

```
    QuerySecurity                  0              4              0              0
    SetSecurity                    0              2              0              0
  Open                             0              0              0              0
  OpenAndX                         0              0              0              0
  QueryInfo                        0              0              0              0
  QueryInfoDisk                    0              0              0              0
  Read                             0              0              0              0
  ReadAndX                         0              0              0              0
  Rename                           0              0              0              0
  Seek                             0              0              0              0
  SessionSetupAndX                 4             14              0              3
  SetInfo                          0              0              0              0
  Transaction                      0              0              0              0
  Transaction2:
    CreateDir                      0              0              0              0
    FindFirst2                     0              6              0              0
    FindNext2                      0              0              0              0
    GetDFSReferral                 0              0              0              0
    QueryFSInfo                    0             68              0             18
    QueryFileInfo                  0              0              0              0
    QueryPathInfo                 32              1              0              0
    SetFileInfo                    0              0              0              0
    SetPathInfo                    0              0              0              0
  TreeConnectAndX                  4             14              0              3
  TreeDisconnect                   4             10              0              0
  Write                            0              0              0              0
  WriteAndX                        0              0              0              0
  UnknownSMB                       0              0              0              0
  ------------------  -------------  -------------  -------------  -------------
              Total:              45            174              0             54

  Errors                  Data Plane      Ctrl Plane          Policy           Probe
  ------------------  -------------  -------------  -------------  -------------
  FormatError                      0              0              0              0
  NetworkError                     0              0              0              0
  ResourceError                    0              0              0              0
  SignatureError                   0              0              0              0
  SMBError                         0              0              0              0
  TimeoutError                     0              0              0              0
  ------------------  -------------  -------------  -------------  -------------
              Total:               0              0              0              0
```

Figure 40.16 *Sample Output: show statistics namespace ... volume ... request-detail*

```
bstnA> show statistics namespace wwmed volume /acct request-detail

Request-count details for shares by namespace and volume
collected at 11/12/2012 05:13:33 -0500

  Namespace: wwmed
  Volume: /acct
  All shares in volume
  Reset at 11/12/2012 01:33:45 -0500

                                  Request Counts per Component
  NFS RPC                 Data Plane      Ctrl Plane          Policy           Probe
  ------------------  -------------  -------------  -------------  -------------
  Null                             0              0              0           5275
  Getattr                     257820           7112              0           2382
  Setattr                      30102           5156              0              0
```

```
  Lookup                        26315           37786               0             526
  Access                       416054               0               0               0
  Readlink                         38               0               0               0
  Read                         124471               0               0             436
  Write                         62062               0               0             348
  Create                            0            3251               0             349
  Mkdir                             0             712               0               0
  Symlink                           0              12               0               0
  Mknod                             0               0               0               0
  Remove                            0              96               0             348
  Rmdir                             0               0               0               0
  Rename                            0               7               0               0
  Link                              0               0               0               0
  Readdir                           0            1387               0               0
  Readdirplus                       0            1862               0               0
  Fsstat                            0            3707               0            1166
  Fsinfo                            0             507               0              36
  Pathconf                          0               0               0               0
  Commit                         9173               0               0               0
  StatFs                            0               0               0               0
  ------------------  --------------  --------------  --------------  --------------
              Total:           926035           61595               0           10866

  Errors                   Data Plane      Ctrl Plane          Policy           Probe
  ------------------  --------------  --------------  --------------  --------------
  FormatError                       0               0               0               0
  NetworkError                      0               0               0               0
  ResourceError                     0               0               0               0
  RPCError                          0               9               0              11
  TimeoutError                      0               9               0              11
  ------------------  --------------  --------------  --------------  --------------
              Total:                0              18               0              22
```

Figure 40.17 *Sample Output: show statistics namespace ... volume ... share all request-detail*

```
bstnA> show statistics namespace wwmed volume /acct share all request-detail

Request-count details for shares by namespace and volume
collected at 11/13/2012 02:24:00 -0500

  Namespace: wwmed
  Volume: /acct
  Share: budget
  Filer: das1
  NFS export: /exports/budget
  Reset at 11/13/2012 01:30:44 -0500

                                  Request Counts per Component
  NFS RPC                  Data Plane      Ctrl Plane          Policy           Probe
  ------------------  --------------  --------------  --------------  --------------
  Null                              0               0               0             318
  Getattr                       24373             758               0             149
  Setattr                       11224            1945               0               0
  Lookup                         4311           12657               0              34
  Access                       183036               0               0               0
  Readlink                          9               0               0               0
  Read                          22658               0               0              28
  Write                         30750               0               0              22
  Create                            0            1053               0              22
  Mkdir                             0             237               0               0
  Symlink                           0               4               0               0
```

```
Mknod                          0              0              0              0
Remove                         0             17              0             22
Rmdir                          0              0              0              0
Rename                         0              3              0              0
Link                           0              0              0              0
Readdir                        0            207              0              0
Readdirplus                    0            766              0              0
Fsstat                         0            203              0             70
Fsinfo                         0             13              0              3
Pathconf                       0              0              0              0
Commit                      3434              0              0              0
StatFs                         0              0              0              0
------------------  -------------  -------------  -------------  -------------
            Total:         279795          17863              0            668

Errors                   Data Plane     Ctrl Plane         Policy          Probe
------------------  -------------  -------------  -------------  -------------
FormatError                    0              0              0              0
NetworkError                   0              0              0              0
ResourceError                  0              0              0              0
RPCError                       0              0              0              0
TimeoutError                   0              0              0              0
------------------  -------------  -------------  -------------  -------------
            Total:             0              0              0              0


Namespace: wwmed
Volume: /acct
Share: bills
Filer: das8
NFS export: /work1/accting
Reset at 11/13/2012 01:30:52 -0500

                                  Request Counts per Component
NFS RPC                  Data Plane     Ctrl Plane         Policy          Probe
------------------  -------------  -------------  -------------  -------------
Null                           0              0              0            318
Getattr                     6408            360              0            149
Setattr                     3109            903              0              0
Lookup                      2558           9818              0             34
Access                     22992              0              0              0
Readlink                       2              0              0              0
Read                        2419              0              0             28
Write                       7349              0              0             22
Create                         0            335              0             22
Mkdir                          0            117              0              0
Symlink                        0              1              0              0
Mknod                          0              0              0              0
Remove                         0             11              0             22
Rmdir                          0              0              0              0
Rename                         0              1              0              0
Link                           0              0              0              0
Readdir                        0             43              0              0
Readdirplus                    0            281              0              0
Fsstat                         0            201              0             70
Fsinfo                         0              9              0              3
Pathconf                       0              0              0              0
Commit                       952              0              0              0
StatFs                         0              0              0              0
------------------  -------------  -------------  -------------  -------------
            Total:         45789          12080              0            668
```

```
Errors                 Data Plane      Ctrl Plane          Policy           Probe
-----------------  --------------  --------------  --------------  --------------
FormatError                    0               0               0               0
NetworkError                   0               0               0               0
ResourceError                  0               0               0               0
RPCError                       0               0               0               0
TimeoutError                   0               0               0               0
-----------------  --------------  --------------  --------------  --------------
            Total:             0               0               0               0


Namespace: wwmed
Volume: /acct
Share: bills2
Filer: das3
NFS export: /exports/acct2
Reset at 11/13/2012 01:30:53 -0500

                                 Request Counts per Component
NFS RPC                Data Plane      Ctrl Plane          Policy           Probe
-----------------  --------------  --------------  --------------  --------------
Null                           0               0               0             318
Getattr                    28490             766               0             149
Setattr                    14343            2051               0               0
Lookup                      7135           13496               0              34
Access                    122947               0               0               0
Readlink                       7               0               0               0
Read                        3443               0               0              28
Write                      15646               0               0              22
Create                         0            1480               0              22
Mkdir                          0             270               0               0
Symlink                        0               6               0               0
Mknod                          0               0               0               0
Remove                         0               3               0              22
Rmdir                          0               0               0               0
Rename                         0               3               0               0
Link                           0               0               0               0
Readdir                        0             161               0               0
Readdirplus                    0             695               0               0
Fsstat                         0             201               0              70
Fsinfo                         0               9               0               3
Pathconf                       0               0               0               0
Commit                      4277               0               0               0
StatFs                         0               0               0               0
-----------------  --------------  --------------  --------------  --------------
            Total:        196288           19141               0             668

Errors                 Data Plane      Ctrl Plane          Policy           Probe
-----------------  --------------  --------------  --------------  --------------
FormatError                    0               0               0               0
NetworkError                   0               0               0               0
ResourceError                  0               0               0               0
RPCError                       0               0               0               0
TimeoutError                   0               0               0               0
-----------------  --------------  --------------  --------------  --------------
            Total:             0               0               0               0


Namespace: wwmed
Volume: /acct
Share: it5
Filer: das7
```

```
NFS export: /lhome/it5
Reset at 11/13/2012 01:30:54 -0500

                                 Request Counts per Component
NFS RPC                Data Plane      Ctrl Plane          Policy           Probe
-----------------  --------------  --------------  --------------  --------------
Null                            0               0               0             318
Getattr                      1474              29               0             149
Setattr                       276              76               0               0
Lookup                        137            1518               0              34
Access                       1114               0               0               0
Readlink                        0               0               0               0
Read                         6606               0               0              28
Write                        6525               0               0              22
Create                          0               0               0              22
Mkdir                           0               0               0               0
Symlink                         0               0               0               0
Mknod                           0               0               0               0
Remove                          0              65               0              22
Rmdir                           0               0               0               0
Rename                          0               0               0               0
Link                            0               0               0               0
Readdir                         0              38               0               0
Readdirplus                     0             119               0               0
Fsstat                          0             201               0              70
Fsinfo                          0               9               0               3
Pathconf                        0               0               0               0
Commit                        126               0               0               0
StatFs                          0               0               0               0
-----------------  --------------  --------------  --------------  --------------
            Total:          16258            2055               0             668

Errors                 Data Plane      Ctrl Plane          Policy           Probe
-----------------  --------------  --------------  --------------  --------------
FormatError                     0               0               0               0
NetworkError                    0               0               0               0
ResourceError                   0               0               0               0
RPCError                        0               0               0               0
TimeoutError                    0               0               0               0
-----------------  --------------  --------------  --------------  --------------
            Total:              0               0               0               0
```

show statistics namespace ... response-detail

Purpose Use the show statistics namespace ... response-detail command to show latency statistics between a namespace and its back-end shares.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show statistics namespace ns [volume vol [share {shr | all}]]
response-detail
```

ns (1-30 characters) identifies the namespace.

vol (optional; 1-1024 characters) specifies one volume in the namespace.

shr | **all** is a required choice if you use the optional **share** keyword:

shr (1-64 characters) specifies one share in the volume.

all shows all shares in the volume. This creates a separate table for each share; if you omit the **share all** option, the output contains a single table showing the sums of all shares.

response-detail shows the round-trip times between various ARX software and back-end shares. This also shows error counts.

Guidelines This command displays round-trip times (RTTs) between the chosen namespace, volume or share and its back-end filers. These are cross-referenced by request format and by the software component that made the request. To see the counts for all of these requests, you can use show statistics namespace ... request-detail. For a summary, showing request and response counts together with RTTs, use show statistics namespace ... summary. The output of all three commands includes the counts of various errors. To see the raw read/write statistics for the same namespace, volume, or share, you can use the show statistics namespace ... fastpath command. As with this command, you can see these statistics summarized for an entire namespace, focused on a single volume, or focused on one share.

Guidelines: Output The output shows the round-trip times (RTTs) for specific requests from each component of namespace software. Each selected volume has a separate table of RTT statistics with the following fields:

Namespace and

Volume identify the volume.

All shares in volume appears if you did not use the **share** keyword to specify a single share. This indicates that the statistics are at a volume level; they are the sum of the RTT statistics for all of the volume's shares. If you select a particular share, the following fields appear instead:

– Share identifies the share, as it is named in the ARX configuration. The show global-config namespace *name* command shows the share names in the ARX configuration.

Guidelines: Output (Cont.)

- Filer is the name of the external filer that hosts the share. Again, this is the name of the filer in the ARX configuration; you can use the show external-filer *filer-name* command to see the IP address and/or SPN for *filer-name*.
- NFS Export or
- CIFS Share is the name of the export or share at the filer. Both names appear if this is used as a multi-protocol share.

Reset... shows the last time these statistics were cleared. The software resets these statistics whenever someone reloads the chassis or uses clear statistics filer on the share.

The remainder of the output are up to four tables: average times for NFS RPCs, counts of NFS Errors, times for CIFS SMBs, and/or counts of CIFS Errors. Each table shows request (or error) types with four columns:

- Data Plane is *fastpath* traffic from the network software.
- Ctrl Plane (short for Control Plane) is traffic from the volume-group software.
- Policy shows the RTT times from the policy engine, which performs migrations, shadow copies, and other policy-related functions.
- Probe is the RTT times from the health-checking subsystem in the namespace software. This subsystem sends periodic requests to every back-end share that is connected to an ARX volume.

Guidelines: NFS RPC Table

NFS RPC - Each row of the NFS table shows the average RTTs for a particular NFS-RPC request. Each column shows the source of the RPC, as described above.

Guidelines: NFS Error Table

Errors appear in a separate table under the NFS statistics. The table contains the following counters, with the same columns as the RPC table:

- FormatError is the number of malformed NFS packets received from the filer(s).
- NetworkError counts the send and receive failures.
- ResourceError is the number of times that the ARX software ran out of resources. Contact F5 Support if you see these errors.
- RPCError counts any errors from the RPC layer.
- TimeoutError is the number of timeouts waiting for an NFS response.

Guidelines: CIFS SMB Table

CIFS SMB - Each row in the CIFS table shows the average RTTs for a given CIFS request. The columns show the source of each SMB, as above.

The NTCancel command is designed for responses to Cancel commands. The CIFS protocol does not require any response to the Cancel command, so this is always "--."

The NTTransact NotifyChange command is for responses to Notify Change commands (see the documentation for cifs notify-change-mode). These response times can be arbitrarily long, so they are omitted and always appear as "--."

Guidelines: CIFS Error Table

Errors appear in a similar table under the CIFS statistics, with the following fields:

- FormatError is the number of malformed SMB (CIFS) packets received from the filer(s).
- NetworkError counts the send and receive failures for SMB traffic.
- ResourceError shows the number of times that the ARX software had insufficient resources to complete the transaction. Contact F5 Support if you see these errors.
- SignatureError counts any SMB packets with missing or incorrect signatures. You can use the cifs filer-signatures command to change the SMB-signing policy for the current namespace.
- SMBError is the number of SMB packets that failed with errors that are not directly related to any of the above errors.
- TimeoutError is the number of timeouts waiting for an SMB response.

Guidelines: SMB2 Table

SMB2 - Each row in the SMB2 table shows the total counts for a given SMB2 request. You can use the smb2-allowed command to enable SMB2 (an updated version of CIFS) for the ARX. This has the same columns as the NFS and CIFS tables above.

The Cancel command is designed for responses to Cancel commands. The CIFS protocol does not require any response to the Cancel command, so this is always "--."

The ChangeNotify command is for responses to Notify Change commands (see the documentation for cifs notify-change-mode). These response times can be arbitrarily long, so they are omitted and always appear as "--."

Guidelines: SMB2 Error Table

Errors appear in a similar table under the SMB2 statistics, with the following fields:

- FormatError is the number of malformed SMB2 packets received from clients.
- NetworkError counts the send and receive failures for SMB2 traffic.
- ResourceError shows the number of times that the ARX software had insufficient resources to complete the transaction. Contact F5 Support if you see these errors.
- SignatureError counts any SMB2 packets with missing or incorrect signatures. You can use the cifs filer-signatures command to change the SMB-signing policy for the current namespace.
- SMBError is the number of SMB2 packets that failed with errors that are not directly related to any of the above errors.
- TimeoutError is the number of timeouts waiting for an SMB2 response.

Guidelines: Clearing Statistics

You can use the clear statistics filer command to clear the statistics from all variations of show statistics namespace. If you use that command to clear the statistics from one share in a volume but not another, the volume-level statistics in this command are inconsistent; one share may have statistics for 2 months and the other share may have statistics for only 5 minutes. To confirm that the volume has consistent statistics, run this command on each of the volume's shares and verify that the Reset... field has the same date and time for each of them.

Sample `bstnA> show statistics namespace medarcv volume /rcrds response-detail`
shows RTT statistics for the volume, "medarcv~/rcrds." See Figure 40.18 on page 40-64 for sample output.

Related Commands
clear statistics filer
show global-config **namespace**
show external-filer
show statistics namespace ... fastpath
show statistics namespace ... request-detail
show statistics namespace ... summary

Figure 40.18 *Sample Output: show statistics namespace ... volume ... response-detail*

```
bstnA> show statistics namespace medarcv volume /rcrds response-detail

Response-time details for shares by namespace and volume
collected at 02/12/2013 09:03:23 -0500

  Namespace: medarcv
  Volume: /rcrds
  All shares in volume
  Reset at 02/12/2013 08:09:00 -0500

                              Average Round-Trip Times per Component (uSec)
  CIFS SMB                 Data Plane      Ctrl Plane          Policy           Probe
  ------------------  --------------  --------------  --------------  --------------
  CheckDir                         --              --              --              --
  Close                           232             249              --             314
  Create                           --              --              --              --
  CreateDir                        --              --              --              --
  Delete                           --             620              --             601
  DeleteDir                        --             432              --              --
  Echo                             --              --              --             455
  FindClose2                       --             245              --              --
  Flush                            --              --              --              --
  LockingAndX                      --              --              --              --
  LogoffAndX                      899             540              --              --
  Negotiate                       427             494              --            1114
  NTCancel                         --              --              --              --
  NTCreateAndX                    487             363              --            1752
  NTRename                         --              --              --              --
  NTTransact:
    Create                         --             441              --              --
    GetUserQuota                   --              --              --              --
    Ioctl                          --              --              --              --
    NotifyChange                   --              --              --              --
    QuerySecurity                  --             260              --              --
    SetSecurity                    --             240              --              --
  Open                             --              --              --              --
  OpenAndX                        454              --              --              --
  QueryInfo                        --              --              --              --
  QueryInfoDisk                    --              --              --              --
  Read                             --              --              --              --
  ReadAndX                        650              --              --             481
  Rename                           --              --              --              --
  Seek                             --              --              --              --
  SessionSetupAndX                903            1083              --            1556
  SetInfo                          --              --              --              --
```

```
Transaction                     --              --              --              --
Transaction2:
  CreateDir                     --              --              --              --
  FindFirst2                    --             386              --              --
  FindNext2                     --              --              --              --
  GetDFSReferral                --              --              --              --
  QueryFSInfo                   --             338              --             505
  QueryFileInfo                328             282              --              --
  QueryPathInfo                254             595              --              --
  SetFileInfo                  200             263              --              --
  SetPathInfo                   --              --              --              --
TreeConnectAndX                362             391              --             457
TreeDisconnect                 434             321              --              --
Write                           --              --              --              --
WriteAndX                      866              --              --             669
UnknownSMB                      --              --              --              --
------------------  --------------  --------------  --------------  --------------
     Net avg time:             616             342              --             642

Errors                   Data Plane      Ctrl Plane          Policy           Probe
------------------  --------------  --------------  --------------  --------------
FormatError                      0               0               0               0
NetworkError                     0               0               0               0
ResourceError                    0               0               0               0
SignatureError                   0               0               0               0
SMBError                         0               0               0               0
TimeoutError                     0               0               0               0
------------------  --------------  --------------  --------------  --------------
            Total:               0               0               0               0

                            Average Round-Trip Times per Component (uSec)
SMB2                     Data Plane      Ctrl Plane          Policy           Probe
------------------  --------------  --------------  --------------  --------------
Cancel                          --              --              --              --
ChangeNotify                    --              --              --              --
Close                           --             448              --             647
Create                          --             464              --             723
Echo                            --              --              --             466
Flush                           --              --              --              --
Ioctl                           --              --              --              --
Lock                            --              --              --              --
Logoff                          --              --              --              --
Negotiate                       --              --              --              --
QueryDirectory                  --             297              --              --
QueryInfo                       --             527              --             722
Read                            --              --              --             592
SessionSetup                    --             702              --            1542
SetInfo                         --             296              --             988
TreeConnect                     --             320              --             468
TreeDisconnect                  --              --              --              --
Write                           --              --              --             974
UnknownSMB                      --              --              --              --
------------------  --------------  --------------  --------------  --------------
     Net avg time:              --             451              --             660

Errors                   Data Plane      Ctrl Plane          Policy           Probe
------------------  --------------  --------------  --------------  --------------
FormatError                      0               0               0               0
NetworkError                     0               0               0               0
ResourceError                    0               0               0               0
SignatureError                   0               0               0               0
SMBError                         0               0               0               0
```

```
  TimeoutError                    0              0              0              0
  ------------------  --------------  -------------  -------------  -------------
              Total:              0              0              0              0
```

show statistics namespace ... summary

Purpose Use the show statistics namespace ... summary command to show request counts, response counts, and average round-trip times (RTTs) between a namespace and its back-end-filer shares.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show statistics namespace` ***name*** `[volume` ***vol*** `[share {`***shr*** `| all}]] [summary]`

name (1-30 characters) identifies the namespace.

vol (optional; 1-1024 characters) specifies one volume in the namespace.

shr **| all** is a required choice if you use the optional **share** keyword:

> ***shr*** (1-64 characters) specifies one share in the volume.
>
> **all** shows all shares in the volume. This creates a separate table for each share; if you omit the **share all** option, the output contains a single table showing the sums of all shares.

summary (optional) shows request and response counts per ARX request type, and shows the average RTT for each request type. This also shows error counts.

Guidelines This command shows the request counts, response counts, and round-trip times (RTTs) from the chosen namespace, volume or share to its back-end filers. These are cross-referenced by request format and by the software component that made the request. To show the detailed request counts, you can use show statistics namespace ... request-detail. To see the detailed RTTs for all of these requests, you can use show statistics namespace ... response-detail. The output of all three commands includes the counts of various errors. To see the raw read/write statistics for the same namespace, volume, or share, you can use the show statistics namespace ... fastpath command. As with this command, you can see these statistics summarized for an entire namespace, focused on a single volume, or focused on one share.

Guidelines: Output The output shows the request counts, response counts, and round-trip times (RTTs) for specific NFS and/or CIFS requests. Each selected volume has a separate table of summary statistics with the following fields:

> Namespace and
>
> Volume identify the volume.
>
> All shares in volume appears if you did not use the **share** keyword to specify a single share. This indicates that the statistics are at a volume level; they are the sum of the statistics for all of the volume's shares. If you select a particular share, the following fields appear instead:
>
> – Share identifies the share, as it is named in the ARX configuration. The show global-config namespace *name* command shows the share names in the ARX configuration.

Guidelines: Output (Cont.)

- Filer is the name of the external filer that hosts the share. Again, this is the name of the filer in the ARX configuration; you can use the show external-filer *filer-name* command to see the IP address and/or SPN for *filer-name*.
- NFS Export or
- CIFS Share is the name of the export or share at the filer. Both names appear if this is used as a multi-protocol share.

Reset... shows the last time these statistics were cleared. The software resets these statistics whenever someone reloads the chassis or uses clear statistics filer on the share.

The remainder of the output contains up to four tables: summaries for NFS RPCs, counts of NFS Errors, summaries for CIFS SMBs, and/or counts of CIFS Errors.

Guidelines: NFS RPC Table

NFS RPC - Each row of the NFS table shows the summaries for a particular NFS-RPC request. This has the following columns:

- Req Count is the total number of requests from this namespace, including requests from the control plane, data plane, policy engine, and probe subsystem.
- Resp Count is the total number of responses from this namespace's filers. These counts should match the request counts.
- Avg RTT (uSec) is average round-trip time (RTT) for the request/response pairs.

Guidelines: NFS Error Table

Errors appear in a separate table under the NFS statistics. The table contains the following counters:

- FormatError is the number of malformed NFS packets received from the filer(s).
- NetworkError counts the send and receive failures.
- ResourceError is the number of times that the ARX software ran out of resources. Contact F5 Support if you see these errors.
- RPCError counts any errors from the RPC layer.
- TimeoutError is the number of timeouts waiting for an NFS response.

Guidelines: CIFS SMB Table

CIFS SMB - Each row in the CIFS table shows the summaries for a given CIFS request. This has the same columns as the NFS RPC table.

The NTCancel command is designed to count responses to Cancel commands. The CIFS protocol does not require any response to the Cancel command, so these counts are always 0 (zero) and the average RTT is always "--."

The NTTransact NotifyChange command counts responses to Notify Change commands (see the documentation for cifs notify-change-mode). These response times can be arbitrarily long, so they always have 0 (zero) counts and the RTT is always "--."

Guidelines: CIFS Error Table

Errors appear in a similar table under the CIFS statistics, with the following fields:

- FormatError is the number of malformed SMB (CIFS) packets received from the filer(s).
- NetworkError counts the send and receive failures for SMB traffic.
- ResourceError shows the number of times that the ARX software had insufficient resources to complete the transaction. Contact F5 Support if you see these errors.
- SignatureError counts any SMB packets with missing or incorrect signatures. You can use the cifs filer-signatures command to change the SMB-signing policy for the current namespace.
- SMBError is the number of SMB packets that failed with errors that are not directly related to any of the above errors.
- TimeoutError is the number of timeouts waiting for an SMB response.

Guidelines: SMB2 Table

SMB2 - Each row in the SMB2 table shows the total counts for a given SMB2 request. You can use the smb2-allowed command to enable SMB2 (an updated version of CIFS) for the ARX. This has the same columns as the NFS and CIFS tables above.

The Cancel command is designed to count responses to Cancel commands. The CIFS protocol does not require any response to the Cancel command, so these counts are always 0 (zero) and the average RTT is always "--."

The ChangeNotify command counts responses to Notify Change commands (see the documentation for cifs notify-change-mode). These response times can be arbitrarily long, so they always have 0 (zero) counts and the RTT is always "--."

Guidelines: SMB2 Error Table

Errors appear in a similar table under the SMB2 statistics. This has the same columns as the CIFS Error table described above.

Guidelines: Clearing Statistics

You can use the clear statistics filer command to clear the statistics from all variations of show statistics namespace. If you use that command to clear the statistics from one share in a volume but not another, the volume-level statistics in this command are inconsistent; one share may have statistics for 2 months and the other share may have statistics for only 5 minutes. To confirm that the volume has consistent statistics, run this command on each of the volume's shares and verify that the Reset... field has the same date and time for each of them.

Sample

```
stoweA> show statistics namespace lodges summary
```

shows summary statistics for the namespace, "lodges." See Figure 40.19 for sample output.

Related Commands

clear statistics filer
show global-config **namespace**
show external-filer
show statistics namespace ... fastpath
show statistics namespace ... request-detail
show statistics namespace ... response-detail

Figure 40.19 *Sample Output: show statistics namespace ... summary*

```
stoweA> show statistics namespace lodges summary

Request/response summary for shares by namespace and volume
collected at 02/13/2013 11:18:10 -0500

  Namespace: lodges
  Volume: /skiPatrol
  All shares in volume
  Reset at 02/13/2013 11:15:06 -0500

  CIFS SMB                 Req Count      Resp Count   Avg RTT (uSec)
  -----------------  --------------  --------------  --------------
  CheckDir                         0               0              --
  Close                           13              13             310
  Create                           0               0              --
  CreateDir                        0               0              --
  Delete                           0               0              --
  DeleteDir                        5               5             361
  Echo                            27              27             315
  FindClose2                       0               0              --
  Flush                            0               0              --
  LockingAndX                      0               0              --
  LogoffAndX                      10              10             539
  Negotiate                       18              18             517
  NTCancel                         0               0              --
  NTCreateAndX                    13              13             767
  NTRename                         0               0              --
  NTTransact:
    Create                         0               0              --
    GetUserQuota                   0               0              --
    Ioctl                          0               0              --
    NotifyChange                   0               0              --
    QuerySecurity                  4               4             215
    SetSecurity                    2               2             249
  Open                             0               0              --
  OpenAndX                         0               0              --
  QueryInfo                        0               0              --
  QueryInfoDisk                    0               0              --
  Read                             0               0              --
  ReadAndX                         0               0              --
  Rename                           0               0              --
  Seek                             0               0              --
  SessionSetupAndX                21              21            1054
  SetInfo                          0               0              --
  Transaction                      0               0              --
  Transaction2:
    CreateDir                      0               0              --
    FindFirst2                     6               6             494
    FindNext2                      0               0              --
    GetDFSReferral                 0               0              --
    QueryFSInfo                   86              86             339
    QueryFileInfo                  0               0              --
    QueryPathInfo                 33              33             318
    SetFileInfo                    0               0              --
    SetPathInfo                    0               0              --
  TreeConnectAndX                 21              21             448
  TreeDisconnect                  14              14             356
  Write                            0               0              --
  WriteAndX                        0               0              --
  UnknownSMB                       0               0              --
  ------------------  --------------  --------------  --------------
```

```
    Total/Net Avg:             273             273             438

Errors                       Count
-----------------  --------------
FormatError                      0
NetworkError                     0
ResourceError                    0
SignatureError                   0
SMBError                         0
TimeoutError                     0
-----------------  --------------
            Total:               0
```

strict-attribute-consistency

Purpose

This command applies only to a share in a multi-protocol (NFS and CIFS) volume with one or more NFS-only directories. An *NFS-only* directory has a name that cannot be expressed in both NFS and CIFS, so its back-end filer has created a filer-generated name (*FGN*) for CIFS clients. The CIFS attributes for NFS-only directories are obscured because it is not possible to deterministically find the directory's CIFS-side FGN. Directory duplication, called *striping*, is required for many standard operations in a managed volume, including file migration (place-rule, auto-migrate, balance) and share removal (remove-share migrate, remove-share nomigrate). The volume can never stripe an NFS-only directory if strict-attribute consistency is required. The no strict-attribute-consistency command allows striping of NFS-only directories; it permits the volume to ignore any CIFS attributes that it cannot find. The volume replaces any undiscovered CIFS attribute with a 0 (zero).

Use the affirmative form, strict-attribute-consistency, for multi-protocol volumes where NFS-only striping is not required.

Modes

gbl-ns-vol-shr

Security Role(s)

storage-engineer or crypto-officer

Syntax

```
no strict-attribute-consistency
strict-attribute-consistency
```

Default(s)

```
strict-attribute-consistency
```

Guidelines

This command is not allowed in a CIFS-only or NFS-only namespace. Strict-attribute consistency, the default, is recommended for most situations. Use the no form of this command to permit migrations of one or more NFS-only directories.

Important

If you migrate to a share with no strict-attribute-consistency, *NFS-only directories lose CIFS attributes at the destination share.*

You may want to re-enable strict-attribute-consistency on the share after completing a one-time migration (such as a remove-share migrate).

Guidelines (Cont.)

The best practice in this case is to rename all NFS-only directories so that they are accessible to CIFS clients. If all directories are renamed, you can keep the strict-attribute-consistency setting for all shares in the volume and support all forms of migration and striping.

The following characters are supported for NFS directory names, but not for CIFS:

`any control character, /, \, :, *, ?, >, <, ”, or |`

Any of these names result on an FGN at the back-end filer, which makes the directory “NFS-only.”

Additionally, NFS supports case collisions and CIFS does not. A *case collision* is two files or directories whose names differ only in case: for example, “oneDir” and “ONEDIR.”

Note that NFS clients can create NFS-only directories while the volume is in service. If strict-attribute-consistency is enabled, any such directory-creation (or rename) disables all striping for the directory. It also makes the directory and its contents inaccessible to CIFS clients, regardless of the strict-attribute-consistency setting.

To rename the directory, access it through one of the volume’s nfs exports.

Samples

`bstnA(gbl-ns-vol-shr[insur~/claims~precs])# no strict-attribute-consistency`

allows the volume to drop any CIFS-directory attributes that it cannot find when it stripes NFS-only directories in the ‘precs’ share.

`bstnA(gbl-ns-vol-shr[wwmed~/acct~bills])# strict-attribute-consistency`

forces the volume to reconcile all CIFS attributes before striping any directory in the ‘bills’ share. NFS-only directories cannot be striped, but no CIFS attributes are ever discarded.

Related Commands

place-rule
auto-migrate
balance
remove-share migrate
remove-share nomigrate

wait-for remove

Purpose Use the wait-for remove command to wait until a volume-removal or share-removal operation has completed.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`wait-for remove`** ***`namespace vol-path`*** **`[`*****`share`***`] [timeout`** ***`timeout`***`]`

namespace (1-30 characters) is the name of the namespace.

vol-path (1-1024 characters) identifies the volume.

share (optional, 1-64 characters) is the share that is being removed.

timeout (optional, 1-2096) is the timeout value in seconds.

Default(s) ***timeout*** - none, wait indefinitely

Guidelines This command is useful in CLI scripts.

When removing a volume with remove namespace ... volume, or a share with remove-share migrate or remove-share nomigrate, you can use the wait-for remove command to wait for the operation to complete. This can be useful for CLI scripts, which you can copy onto the switch (with copy ftp, copy scp, copy {nfs|cifs}, or copy tftp), and run.

If you set a ***timeout*** and it expires before the volume or share is removed, the command exits with a warning. To interrupt the wait-for remove command, press **<Ctrl-C>**.

Samples

```
bstnA> wait-for remove ns21 /vol4
```

waits indefinitely for the '/vol4' volume to be removed from the 'ns21' namespace. The CLI prompt does not return until the volume-removal operation is complete.

```
bstnA> wait-for remove ns66 /vol31 rxbills timeout 60
```

waits up to 60 seconds for the '/vol31~rxbills' share to be removed from the 'ns66' namespace.

Related Commands

remove namespace **`... volume`**
remove-share migrate
remove-share nomigrate

41

CIFS-Service Troubleshooting Tools

cifs promote-subshares

Purpose A managed volume duplicates all CIFS subshares between its back-end filers. That is, if a CIFS share exists below one of its back-end shares, that lower share (called a subshare) and its ACL is duplicated to all of its peer subshares in the managed volume. In former software releases, the managed volume used a special name for the replicated subshares, with the format "_acopia_*subshare-name_id*$" (for example, "acopia_CELEBS_3$"). You can use this command to promote all such subshare names to a native name, such as "CELEBS." You can invoke this command on one share at a time.

Mode priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax **`cifs promote-subshares namespace`** ***`ns`*** **`volume`** ***`vol`*** **`share`** ***`shr`*** **`[tentative]`**

ns (1-30 characters) identifies the share's namespace. Use the show namespace command for a list of all namespaces.

vol (optional, 1-1024 characters) identifies the managed volume with CIFS subshares. The show global-config namespace *ns-name* command lists all volumes (and other parameters) in a particular namespace.

shr (optional, 1-64 characters) is the share with ARX-generated subshares.

tentative (optional) creates a report showing all of the ARX-generated subshares at the back-end share, and showing the native share names that the managed volume would give to each subshare. You can use this option to confirm that the volume would rename the subshares as needed, then you can re-run the command without the option to actually perform the renames.

Default(s) None

Guidelines When you use the command without the tentative option, the CLI prompts for confirmation before it performs any subshare renames. Enter **yes** to proceed with the rename operations.

This command generates a report to show the results (or tentative results) of the operation. The report name has a prefix of "cifsPromoteSubshares" and a time stamp; the exact name of the report appears after you enter the command.

You can use the tentative option to generate the report, which reveals any subshares with names that were generated unnecessarily. (Some subshare names are generated to avoid a collision with another share name on the same back-end filer; you cannot use native names for those shares.) If you find any such subshare names, you can re-run the command without the tentative option to change them to their native names.

If the volume has no modify set, preventing it from changing any configuration on its back-end filers, this operation always runs as though the tentative flag is raised. The command generates a report but does not change any subshare names.

Use the filer-subshares command to enable subshare support in the managed volume.

Samples

```
bstnA# cifs promote-subshares namespace medarcv volume /rcrds share
bulk tentative

% INFO: Promote subshares operation completed. See report
'cifsPromoteSubshares_201009020052.rpt' for more detail.
```

runs a tentative command to promote all subshares in the "bulk" share. The report shown in the prompt lists all of that share's subshares (if there are any) and the new names that the command would give them. A sample report, showing no subshares that require promotion, is in Figure 41.1 on page 41-4.

Had the report shown any subshares that required promotion, you could promote them to native names by re-running the command without the tentative argument.

```
bstnA# cifs promote-subshares namespace ns4 volume /vol share shr3

 This operation renames subshares on filer fs14[172.16.195.44]
 corresponding to ARX-generated subshares under share
 bulk in use by front-end passthrough exports for volume /vol.
 Continue? [yes/no] yes

% INFO: Promote subshares operation completed. See report
'cifsPromoteSubshares_201003161659.rpt' for more detail.
```

promotes all subshares in the "shr3" share. The report lists all of that share's subshares and their new names.

Related Commands filer-subshares

Figure 41.1 *Sample Report: cifsPromoteSubshares...*

```
**** CIFS Promote Subshares Report: Started at Thu Sep  2 00:52:45 2010 ****
**** Software Version: 5.02.000.12617 (Aug 26 2010 22:44:45) [nbuilds]
**** Hardware Platform: ARX-4000
**** Report Destination:

Operation Parameters
--------------------------------------------------------------------------------
Namespace: medarcv
Volume: /rcrds (modifiable, no access-based-enum)

Share Information
    Share: bulk
    Filer Name: fs2
    Filer IP: 192.168.25.27
    Filer Port: 445
    Filer Share: bulkstorage

Options
    Tentative. Report, but do not perform the operation.

ABE - access-based-enumeration.
ACL - security descriptor access control list.

Operation Execution Detail
--------------------------------------------------------------------------------


Summary:
```

```
--------------------------------------------------------------------------------

No share promote operations were executed.


**** Total processed:               1
**** Elapsed time:           00:00:00
**** Promote Subshares Report: DONE at Thu Sep  2 00:52:45 2010 ****
```

clear statistics cifs authentication

Purpose Use this command to clear the authentication statistics for a front-end CIFS service.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear statistics cifs authentication {`*`fqdn`*` | all}`**

fqdn (1-128 characters) is the fully-qualified domain name (for example, "myserver.organization.org") for a CIFS service. Use show cifs-service to see a list of CIFS services.

all clears statistics for all CIFS services. The CLI requires confirmation before doing this.

Default(s) None

Guidelines The show statistics cifs authentication command shows counters and statistics for all CIFS authentications. Use this command to clear the traffic counters for one (or all) CIFS services.

If you clear all statistics with this command, the CLI prompts for confirmation; enter **yes** to proceed.

Sample

```
bstnA# clear statistics cifs authentication all

Confirmation of this command causes CIFS authentication statistics to
be cleared for all CIFS services.

Proceed [yes/no] yes
```

clears all CIFS-authentication statistics from the system.

Related Commands show statistics cifs authentication
show cifs-service

clear statistics cifs path-cache

Purpose Each NSM processor keeps a cache of file and directory paths for CIFS managed volumes. The NSM processors learn these paths from the managed-volume software, then store them in their path cache. The system keeps running statistics on the usage of the path cache. This command clears the path-cache statistics for a CIFS-supporting namespace or volume.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear statistics cifs path-cache`** ***`namespace`*** **`[`*** ***`volume`*** **`[`*** ***`slot.processor`*** **`]]`**

namespace (1-30 characters) is the name of a namespace. Use the show namespace command for a list of all namespaces.

volume (optional, 1-1024 characters) focuses on the path cache for a single volume. The show global-config namespace ns-name command lists all volumes (and other parameters) in a particular namespace. If you omit this, the statistics are cleared for all volumes in the ***namespace***.

slot.processor (optional: for example, 1.4) focuses on a single NSM processor's cache. Each processor keeps an individual cache for each volume.

Default(s) None

Guidelines Use the cifs path-cache command to enable the CIFS path-cache for a volume. The path cache is only used for SMB traffic, not SMB2 (see smb2-allowed).

This command clears statistics only. It does not clear any path-cache entries.

The show statistics cifs path-cache command shows counters and statistics for the path cache. Use this command to clear the counters for CIFS-supporting volumes. The CLI prompts for confirmation before clearing the statistics; enter **yes** to proceed.

Samples

```
bstnA# clear statistics cifs path-cache insur /rcrds 3.3

Clear path-cache statistics? [yes/no] yes
```

clears the path-cache statistics that processor 3.3 kept for the "insur~/rcrds" volume.

```
bstnA# clear statistics cifs path-cache medarcv

Clear path-cache statistics? [yes/no] yes
```

clears all path-cache statistics on all volumes in the "medarcv" namespace. This clears the statistics at all NSM processors.

Related Commands

cifs path-cache
show statistics cifs path-cache
show cifs-service path-cache

clear statistics cifs symlinks

Purpose CIFS clients can de reference NFS symbolic links (or symlinks) in a multi-protocol volume. Use this command to clear the statistics for these de referencing operations.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear statistics cifs symlinks [volume-group id]`**

volume-group *id* (optional) narrows the scope to a single volume group. A *volume group* is a failure domain for a group of volumes in the same namespace. If you omit this option, the command clears statistics from all volumes on the system.

Default(s) None

Guidelines The show statistics cifs symlinks command shows counters and statistics for all symlink de referencing operations. The statistics only include symlink operations on behalf of CIFS clients; NFS clients perform symlink de referencing at the client machine. Use this command to clear the symlink counters for one (or all) volume groups.

The CLI prompts for confirmation before clearing any statistics; enter **yes** to proceed.

Sample

```
bstnA# clear statistics cifs symlinks

Clear statistics for CIFS symlink resolution for all Volume Groups?

Are you sure? [yes/no] yes
```

clears all CIFS-symlink statistics from the system.

Related Commands show statistics cifs symlinks

clear statistics cifs work-queues

Purpose CIFS volumes use internal *work queues* to process CIFS commands, CIFS authentication requests, and all the component tasks required to accomplish them. The volume software keeps statistics for the amount of processing time used by each work item in these queues. Use this command to clear those work-queue statistics.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax

```
clear statistics cifs work-queues volume-group vg-id
clear statistics cifs work-queues instance instance-id
```

volume-group *vg-id* (optional, 1-255) chooses a volume group. A *volume group* is a failure domain for a group of volumes in the same namespace. You can use show volume-group for a list of all volume groups on the system, and to see which volumes are assigned to each volume group.

instance *instance-id* (optional) chooses a volume group by its namespace-instance ID. Instance IDs often appear in syslog messages, which you can view with show logs syslog. You can can also see instance IDs with show namespace all, which shows full details on all namespaces in the system.

Default(s) None

Guidelines The show statistics cifs work-queues command shows counters and statistics for all CIFS work queues. These statistics include time required to process each work item, number of work items completed and discarded from each queue, average processing time for each CIFS operation, and average round-trip times between the volume(s) and their back-end CIFS servers. The back-end-server statistics can only be cleared by the clear statistics filer command; this command clears all the remaining statistics.

The CLI prompts for confirmation before clearing any statistics; enter **yes** to proceed.

Sample

```
bstnA# clear statistics cifs work-queues volume-group 2
Statistics for filer shares are not cleared by this command.
It may take up to 30 seconds after the command for other statistics to
be cleared.

Clear CIFS work-queue statistics? [yes/no]  yes
```

clears all CIFS-work-queue statistics for volume group 2.

Related Commands show statistics cifs work-queues
clear statistics filer

clear statistics domain-controller

Purpose The ARX keeps statistics for its NTLM/NTLMv2 authentications against each of its domain controllers (DCs). Use this command to clear the statistics for a particular DC, or for all DCs at once.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear statistics domain-controller {`*`ip-address`*` | all}`**

ip-address (for example, 1.2.3.4) identifies the DC for which you want to clear statistics. The show active-directory command shows all DCs that are known to the ARX.

all clears NTLM statistics for all DCs.

Default(s) None

Guidelines The show statistics domain-controller command shows counters and statistics for NTLM and NTLMv2 authentications from all CIFS front-end services with constrained delegation. Use this command to clear the NTLM counters against a particular DC.

The CLI prompts for confirmation before clearing any statistics; enter **yes** to proceed.

Sample

```
bstnA# clear statistics domain-controller 192.168.25.102
Clear domain-controller statistics? [yes/no] yes
```

clears all NTLM/NTLMv2 statistics for the DC at 192.168.25.102.

Related Commands show statistics domain-controller
show active-directory

clear statistics domain-controller load-balancing

Purpose The ARX balances its Kerberos requests between multiple domain controllers (DCs) in the same Windows Domain, and keeps count of the number of requests to each DC. The ARX keeps these counters until the next reboot. Use this command to clear the counters immediately.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`clear statistics domain-controller load-balancing`**

Default(s) None

Guidelines The show statistics domain-controller load-balancing command shows counters for all Kerberos requests to all DCs. Use this command to clear those counters.

The CLI prompts for confirmation before clearing any statistics; enter **yes** to proceed.

Sample

```
bstnA# clear statistics domain-controller load-balancing
Clear load-balancing statistics? [yes/no] yes
```

clears all Kerberos-request counters for all DCs.

Related Commands show statistics domain-controller load-balancing
show active-directory

clear subshare-cache

Purpose A managed volume that supports filer-subshares keeps all of its subshare information in cache memory. This decreases the number of RPC calls between the ARX and its back-end filers during sync subshare operations. On the advice of F5 personnel, you can use this command to clear this cache.

Mode priv-exec

Security Role(s) storage-engineer or crypto-officer

Syntax `clear subshare-cache [ext-filer-name]`

ext-filer-name (optional, 1-64 characters) identifies an external filer with outdated subshare information. This is the name of the filer in the ARX configuration. For a list of configured external filers, use show external-filer.

Default(s) None

Guidelines Use the filer-subshares command to enable filer subshares for a volume. Then use export (gbl-cifs) ... filer-subshare to export a single subshare, or sync subshares from-namespace to export all of the subshares from the back-end filers.

A *filer subshare* is any CIFS share that is inside an imported CIFS share. A client who connects to a front-end subshare, if the subshares are configured as described above, is passed through the managed volume directly to a corresponding subshare on a back-end filer. The filer can then enforce its subshare ACL, as opposed to the top-level ACL of the imported share. Each front-end subshare (visible to your CIFS clients) maps to one or more back-end subshares. The state of each filer share and share ACL resides in a memory cache; you can use show subshare-cache to see the contents o f the cache.

The sync subshares from-namespace and sync subshares from-service commands clear this cache automatically before they begin, but only for the filers affected by the command. You can use this command to clear the entire cache, or the cache for one filer where you are sure that subshare information has changed. This slows the performance of future sync subshare commands, but ensures that the cache is updated with the latest subshare information on the back-end filer(s).

Sample

```
bstnA# clear subshare-cache fs2
```

clears the subshare-cache for the external filer named “fs2.”

Related Commands

filer-subshares
sync subshares from-namespace
sync subshares from-service
show subshare-cache

close cifs file

Purpose You can use the close cifs file command to close a file that is being held open by a CIFS client.

Modes priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`close cifs file`** ***`fqdn slot.processor`*** **`fid`** ***`file-id`***

fqdn (1-128 characters) is the fully-qualified domain name for one CIFS service (for example, "www.organization.org"). This identifies the host for the open file.

slot.processor (for example, 2.4) is the NSM slot and processor that is hosting the CIFS session, from the output of show cifs-service open-files.

file-id (0-65535) identifies the file to close. You can also find this in the output of show cifs-service open-files.

Default(s) None

Guidelines Some Windows applications open a file, read it, immediately close it, and then present its contents to the user as though the file was still open. Notepad and WordPad are two commonly-used applications that use this practice. These applications do not actually hold the files open, so other clients can still access the file for writing. Other applications, such as Microsoft Word, hold the file open until the user closes it through the application interface. This prevents access from any competing applications or clients. This command closes any such file from the CLI.

Use show cifs-service open-files to see all files that are currently held open by CIFS clients. This is the same "open files" listing seen through MMC.

Use show cifs-service user-sessions to find clients with open CIFS sessions. You can also use drop cifs-service user-session to disconnect a client session.

An authorized Windows client can perform this operation from an MMC interface (or a similar Windows-management client), assuming the CIFS service has browsing enabled and the client belongs to a properly-enabled windows-mgmt-auth group.

Sample

```
bstnA# close cifs file ac1.medarch.org 2.6 fid 1241
slot:2.6 CIFS open file with FID:1241  closed.
```

closes an open CIFS file that is hosted by processor 2.6, through a CIFS service at "ac1.medarch.org."

Related Commands

cifs
show cifs-service open-files
show cifs-service user-sessions
browsing
windows-mgmt-auth

drop cifs-service user-session

Purpose Use the drop cifs-service user-session command to drop a client connection to a CIFS service.

Modes priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer, or crypto-officer

Syntax **`drop cifs-service user-session`** ***`fqdn slot.processor`*** **`ipaddress`** ***`client-ip`***

fqdn (1-128 characters) is the fully-qualified domain name for one CIFS service (for example, "www.organization.org"). This identifies the host for the session.

slot.processor (for example, 1.4) is the NSM slot and processor that is hosting the session, from the output of show cifs-service user-sessions.

client-ip (an IP address in *a.b.c.d* format) identifies the client to drop.

Default(s) None

Guidelines This is useful for malfunctioning CIFS clients that fail to tear down their connections, or for testing.

Use show cifs-service user-sessions to find clients with open CIFS sessions. You can also use show cifs-service open-files to see all files that are currently open through CIFS (and holding an exclusive-write lock).

An authorized Windows client can perform this operation from an MMC (or similar) interface, assuming the CIFS service has browsing enabled and the client belongs to a properly-enabled windows-mgmt-auth group.

Sample

```
bstnA# drop cifs-service user-session ac1.medarch.org 2.5 ipaddress
172.16.100.214
```

disconnects a CIFS client at 172.16.100.214.

Related Commands

cifs
show cifs-service user-sessions
show cifs-service open-files
browsing
windows-mgmt-auth

show cifs-service client-activity

Purpose Use the show cifs-service client-activity command to show details about one client's connection to a CIFS service.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show cifs-service client-activity fqdn ip-address
  [connection-id [open-files | pending-transactions]]
```

fqdn (1-128 characters) is the fully-qualified domain name for one CIFS service (for example, "www.organization.org").

ip-address is the IP address of a single client.

connection-id (optional) identifies a single client connection. You can get this ID from the summary output, then rerun the command with the ID to show details about the client connection.

open-files (optional, if you chose a ***connection-id***) shows only the Open Files table for the given client connection.

pending-transactions (optional, if you chose a ***connection-id***) shows only the pending transactions for the given client connection, if any.

Guidelines: Summary Output

The simplest syntax outputs the following fields, table, and summary:

Client IP Address is the IP address of the client, specified on the command line.

SlotProc is the NSM processor that is hosting the CIFS session, in *slot.processor* format.

Then a table shows all of the connections from the above client. Each connection appears on two rows. These rows contain the following fields:

- Conn is an integer ID to identify the front-end CIFS connection. You can rerun this command with this specific ID to see details about the connection, such as the specific files that the client is holding open.
- Port is the port on the NSM to which the client is connected.
- Age is the number of seconds that the client connection has been up.
- Protocol shows which CIFS protocol the client is using, SMB or SMB2. To enable SMB2 for the ARX, use smb2-allowed.
- Sessions is the number of sessions from the client to the CIFS service (at the specified NSM processor).
- Tree Cons is the number of different front-end shares to which client is connected. These may be different volumes in the same namespace and/or different share names for the same front-end share (see export (gbl-cifs)).
- Open Files is the number of files open in this client session. Details on the open files appear in a table below.
- CredMax is the maximum number of credits allowed in the currently-active connection. This is only for an SMB2 connection. SMB2 clients can use *pipelining* to place multiple SMB2 requests on the connection at the same time instead of serially. Each SMB2 request consumes one credit. This is the maximum simultaneous SMB2 requests allowed in the current connection.
- CredUsed is the number of SMB2 credits the client has in use (out of the above maximum). This is the current number of outstanding SMB2 requests for this client.
- Username appears on its own line. This is the name provided by the CIFS client.

Total number of connections is the sum of the above client connections.

Guidelines: Detailed Output for One Client Connection

If you specify a ***connection-id*** in the command, the output also includes two tables. You can use the **open-files** or **pending-transactions** keyword to include only one of these two tables.

Guidelines: Open Files Table

The Open Files table contains three rows per open file. The top row shows various ID's used for the client's front-end and back-end connections:

FE Conn identifies the front-end connection to the client. This is the ***connection-id*** entered in the command.

BE Conn identifies the back-end connection to a file server.

FE UID is the client's User ID. The CIFS service provides the UID to the client after a successful authentication. This is only valid for the duration of the CIFS-client session.

BE UID is the User ID used for the back-end connection. The CIFS service provides the UID to the client after a successful authentication to the back-end server. This is only valid for the duration of the CIFS session.

FE TID is the client's Tree-connection ID. This identifies the client's CIFS connection to a particular resource, such as a directory tree offered by the front-end service. The TID changes for each CIFS session (whenever the connection is broken and re-established).

BE TID is the Tree-connection ID for the back-end connection.

FE FID is the front-end File ID for this file. This changes with each CIFS-client session.

BE FID is the back-end File ID for this file.

In cases where the back-end connection is to an internal process instead of a filer, the BE IDs are 0 (zero).

The second row is the FE Path to the open file. This is the virtual path, as seen by the front-end client.

The final row is the BE Path to the open file. This is the physical path on the back-end server.

Guidelines: Pending Transactions Table

Some CIFS transactions may wait in a queue before they get a response, such as messages from the NSM processes to back-end filers or to ACM processes. These are called *pending transactions*. This table lists all pending transactions, two rows per transaction. The top row shows CIFS IDs from the front-end (FE) perspective, along with the state of the associated back-end transaction:

FE Conn is the client ID from the front-end CIFS service. The service assigns this to the client when it establishes a TCP connection. The CIFS service uses this internally to identify the client session.

FE UID is the client's User ID. The CIFS service provides the UID to the client after a successful authentication. As above, this only valid for the duration of the CIFS session.

FE MID is the Multiplex ID. The client software sets this. Whenever the client has its own pending transaction (waiting for response from the ARX's CIFS service), it can send a packet with a new MID to differentiate the new transaction.

FE PID is the process ID, provided by the client.

FE TID is the client's Tree-connection ID. This identifies the client's CIFS connection to a particular resource, such as a directory tree offered by the front-end service. The TID changes for each CIFS session (whenever the connection is broken and re-established).

FE FID is the File ID for this file. This also changes with each CIFS session.

State explains the current state of the transaction. This is the state of the CIFS connection between the NSM process and the remote process; unlike the preceding fields on the top row, this concerns the back-end (BE) CIFS session. The remote process is either a filer (referenced as the "backend" or the "filer" in the state text) or a namespace process (called "dnas" in the state text).

The bottom row shows the same IDs for the back-end (BE) CIFS session with either a back-end filer or a namespace process on the ACM. If the latter, there is no authentication, tree connection, or file involved, so the UID, TID, and FID are all 0 (zero). The final field, Command, is the CIFS command that is pending.

Guidelines: Table Summaries

Total Open Files and

Total Pending Transactions appear beneath the above tables.

Guidelines: Manipulating the Client Session

Use the show cifs-service user-sessions command for a list of all client connections to a CIFS service. You can use drop cifs-service user-session to disconnect a client session. An authorized Windows client can show and drop CIFS sessions from a Windows-management interface like MMC, assuming the CIFS service has browsing enabled and the client belongs to a properly-enabled windows-mgmt-auth group.

For statistics on CIFS authentication, use show statistics cifs authentication.

Samples

```
bstnA> show cifs-service client-activity ac1.medarch.org 172.16.100.20
```

lists all client sessions with the "ac1.medarch.org" CIFS service. See Figure 41.2 for sample output.

```
bstnA> show cifs-service client-activity ac1.medarch.org 172.16.100.20
24
```

lists one client session with the "ac1.medarch.org" CIFS service. This uses the client-connection ID from the previous output, 24. See Figure 41.3 on page 41-19 for sample output.

Related Commands

cifs
show cifs-service user-sessions
drop cifs-service user-session
show cifs-service open-files
browsing
windows-mgmt-auth
show statistics cifs authentication

Figure 41.2 *Sample Output: show cifs-service client-activity*

```
bstnA> show cifs-service client-activity ac1.medarch.org 172.16.100.20

Client IP Address: 172.16.100.20
SlotProc:          2.7

Conn      Port   Age        Protocol    Sessions    Tree Conn   Open Files   CredMax   CredUsed
Username
-------------------------------------------------------------------------------------------------
22        1170   00:00:01     SMB         1           1           0            8         8
Administrator@MEDARCH.ORG

24        1178   00:00:00     SMB         1           1           0            8         8
administrator@MEDARCH.ORG


Total number of connections: 2
```

Figure 41.3 *Sample Output: show cifs-service client-activity (Detailed)*

```
bstnA> show cifs-service client-activity ac1.medarch.org 172.16.100.20 24

Client IP Address: 172.16.100.20
SlotProc:          2.7

Conn      Port   Age        Protocol    Sessions    Tree Conn   Open Files   CredMax   CredUsed
Username
-------------------------------------------------------------------------------------------------
24        2179   00:00:22     SMB         1           1           8            8         8
juser@MEDARCH.ORG


 Open Files
 ----------
   FE Conn  BE Conn  FE UID  BE UID  FE TID  BE TID  FE FID  BE FID
   FE Path
   BE Path
   ---------------------------------------------------------------

   24       31       4       2050    4       2050    533     16386
   \rcrds\VIP_wing\holdback\examSchedule.doc
   \holdback\examSchedule.doc

   24       31       4       2050    4       2050    536     16387
   \rcrds\VIP_wing\holdback\stdProcedure.doc
   \holdback\stdProcedure.doc

   24       31       4       2050    4       2050    538     16388
   \rcrds\VIP_wing\holdback\recoveryStats.xls
   \holdback\recoveryStats.xls
```

```
   24       30       4       2049    4       2049    539     16388
   \rcrds\VIP_wing\holdback\memo.doc
   \holdback\memo.doc

   24       30       4       2049    4       2049    541     16389
   \rcrds\VIP_wing\holdback\paid_holidays.doc
   \holdback\paid_holidays.doc

   24       30       4       2049    4       2049    543     16390
   \rcrds\VIP_wing\holdback\labBudget.xls
   \holdback\labBudget.xls

   24       30       4       2049    4       2049    545     16391
   \rcrds\VIP_wing\holdback\benchNotes.doc
   \holdback\benchNotes.doc

   24       30       4       2049    4       2049    547     16392
   \rcrds\VIP_wing\holdback\orSched.doc
   \holdback\orSched.doc


Total number of open files:           8
Total number of pending Transactions: 0

Key
----------------------------
Conn - Connection Identifier
  FE - Front End
  BE - Back End
 UID - User Identifier
 TID - Tree Identifier
 FID - File Identifier
```

show cifs-service exports

Purpose Use the show cifs-service exports command to list CIFS exports and their client-connection statistics.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show cifs-service exports {`*`fqdn`* `[`*`slot.processor`*`] | all}`

fqdn (1-128 characters) specifies a particular CIFS service (for example, "www.company.com").

slot.processor (optional: for example, 2.3) focuses on the connections to a particular NSM processor.

all shows the connections statistics for all CIFS services.

Guidelines The output is a table of CIFS shares that are exported by each NSM processor. Each row in the table describes a single export:

Proc is the NSM that is exporting the share. Use show processors for a full list of all processors on the ARX.

Export is the name of the CIFS share as seen by clients. This is created as part of the export (gbl-cifs) command.

Namespace identifies the namespace behind this CIFS share. This is also chosen with the export (gbl-cifs) command.

Virtual Path is the path to the CIFS share from the root of the volume. Again, this is from the perspective of the front-end client, and it is established by the export (gbl-cifs) command.

Tree Connects is the heading for a series of connection statistics:

- Curr is the number of CIFS clients currently connected to the share and processor.
- Peak is the highest number of simultaneous connections.
- Total is the sum of all CIFS connections to the share.

To see the client sessions that are connected to the CIFS service, use show cifs-service user-sessions. To disconnect a client session, use drop cifs-service user-session. The show cifs-service open-files command shows all files that are currently open through CIFS (with an exclusive-write lock). A share farm cannot auto-migrate a file in this state; use close cifs file to close one from the CLI.

Samples `bstnA> show cifs-service exports all`

lists all exports from all CIFS services. See Figure 41.4 on page 41-22 for sample output.

`bstnA> show cifs-service exports ac1.medarch.org 2.8`

lists all exports from a particular CIFS service and NSM processor. See Figure 41.5 on page 41-23 for sample output.

Related Commands

- cifs
- show cifs-service
- show cifs-service user-sessions
- show cifs-service open-files

Figure 41.4 *Sample Output: show cifs-service exports all*

```
bstnA> show cifs-service exports all


Global Server: ac1.MEDARCH.ORG            [192.168.25.15]

                              Namespace                                 Tree Connects
Proc  Export                  Virtual Path                          Curr   Peak   Total
----  ----------------------  ------------------------------------  ----   ----   -----
2.1   acopia#ns3_lab_equipment$   medarcv                             0      0      0
                              /lab_equipment
2.1   acopia#ns3_rcrds$        medarcv                               0      0      0
                              /rcrds
2.1   acopia#ns4_claims$       insur                                 0      0      0
                              /claims
2.1   ARCHIVES                medarcv                               0      0      0
                              /rcrds
2.1   bulkstorage             medarcv                               0      0      0
                              /rcrds
2.1   CELEBS                  medarcv                               0      0      0
                              /rcrds/VIP_wing
2.1   chem_results            medarcv                               0      0      0
                              /test_results
2.1   CLAIMS                  insur                                 0      0      0
                              /claims
2.1   E$                      medarcv                               0      0      0
                              /acopia$ns3
2.1   F$                      insur                                 0      0      0
                              /acopia$ns4
2.1   labs                    medarcv                               0      0      0
                              /lab_equipment
2.1   MP3S                    medarcv                               0      0      0
                              /rcrds/2011/mp3downloads
2.1   SPECS                   insur                                 0      0      0
                              /claims/specs
2.1   STATS                   insur                                 0      0      0
                              /claims/stats

...

2.12  xraysScanners          medarcv                               0      0      0
                              /lab_equipment
2.12  Y2004                   medarcv                               0      1      1
                              /rcrds/2004
2.12  Y2005                   medarcv                               0      0      0
                              /rcrds/2005
2.12  Y2008                   medarcv                               0      0      0
                              /rcrds/2008
2.12  Y2010                   medarcv                               0      0      0
                              /rcrds/2010
2.12  Z$                      medarcv                               0      0      0
                              /test_results
                                                                    ----   ----   -----
                                                          Totals:     0     15     40
```

Figure 41.5 *Sample Output: show cifs-service exports ac1.medarch.org 2.8*

```
bstnA> show cifs-service exports ac1.medarch.org 2.8


Global Server: ac1.MEDARCH.ORG            [192.168.25.15]

                              Namespace                                 Tree Connects
Proc  Export                  Virtual Path                          Curr   Peak   Total
----  ----------------------  ------------------------------------  ----   ----   -----
2.8   acopia#ns3_lab_equipment$   medarcv                             0      0      0
                              /lab_equipment
2.8   acopia#ns3_rcrds$       medarcv                               0      0      0
                              /rcrds
2.8   acopia#ns4_claims$      insur                                 0      0      0
                              /claims
2.8   ARCHIVES                medarcv                               0      0      0
                              /rcrds
2.8   bulkstorage             medarcv                               0      0      0
                              /rcrds
2.8   CELEBS                  medarcv                               0      0      0
                              /rcrds/VIP_wing
2.8   chem_results            medarcv                               0      0      0
                              /test_results
2.8   CLAIMS                  insur                                 0      0      0
                              /claims
2.8   E$                      medarcv                               0      0      0
                              /acopia$ns3
2.8   F$                      insur                                 0      0      0
                              /acopia$ns4
2.8   labs                    medarcv                               0      0      0
                              /lab_equipment
2.8   MP3S                    medarcv                               0      0      0
                              /rcrds/2011/mp3downloads
2.8   SPECS                   insur                                 0      0      0
                              /claims/specs
2.8   STATS                   insur                                 0      0      0
                              /claims/stats
2.8   xraysScanners           medarcv                               0      0      0
                              /lab_equipment
2.8   Y2004                   medarcv                               1      1      3
                              /rcrds/2004
2.8   Y2005                   medarcv                               0      0      0
                              /rcrds/2005
2.8   Y2010                   medarcv                               0      0      0
                              /rcrds/2010
2.8   Z$                      medarcv                               0      0      0
                              /test_results
                                                                    ----   ----   -----
                                                          Totals:     1      1      3
```

show cifs-service kerberos-tickets

Purpose If a CIFS service authenticates its clients with Kerberos, it caches all of the *tickets* granted to its clients. The tickets have expiration times, and are cached by the CIFS service until they expire. Use the show cifs-service kerberos-tickets command to list the Kerberos tickets in the cache.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show cifs-service kerberos-tickets {`***`fqdn`*** `| all} [user` ***`username`***`]`

fqdn | **all** is a required choice:

fqdn (1-128 characters) is the fully-qualified domain name for one CIFS service (for example, "www.organization.org").

all shows the tickets granted by all CIFS-services.

username (optional: 1-128 characters) is used in a case-blind search: this shows all principals whose names start with this string. For example, "myuser" matches "myuser", "MYUSER@myco.com," and "myusername@myorg.org."

Guidelines This command shows the Kerberos tickets granted by CIFS services. Kerberos authentication must be enabled for the service to grant any such tickets: use cifs authentication kerberos (in gbl-ns mode) to enable Kerberos for a namespace, and domain-join to join the CIFS service to a Windows domain so that Kerberos works.

The output has a separate table of principals (clients) and their tickets for each CIFS service. It contains the following fields:

Service identifies the CIFS service by name. This is the FQDN of the CIFS service's global server, used in the cifs command.

Principal is name of the client who requested the ticket(s).

Start Time(UTC) is the date and time when the CIFS service granted the ticket. As noted, this is not in local time.

Expiry Time(UTC) is the date and time when the ticket is due to expire next. As above, this is not in local time.

Service Principal is name of the server or Ticket-Granting Ticket that granted this ticket to the principal.

Renew Till only appears for renewable Ticket-Granting Tickets.

Total number of principals displayed,

Total number of ticket-granting-tickets ..., and

Total number of service tickets ... appear at the end of each CIFS-service section.

An SNMP trap appears if the ticket cache begins to fill up. Use snmp-server traps and snmp-server trusthost to set up SNMP traps, and/or use email-event to deliver the traps via E-mail. The show health output also shows this condition if it arises.

To see the client sessions that are connected to the CIFS service, use show cifs-service user-sessions. Use the show statistics cifs authentication command for statistics on all CIFS-service authentications: Kerberos, NTLMv2, and NTLM.

Sample `bstnA> show cifs-service kerberos-tickets all`
lists all Kerberos tickets granted by all CIFS services. See Figure 41.6 on page 41-25 for sample output.

Related Commands
cifs
cifs authentication **kerberos**
domain-join
show cifs-service user-sessions
show statistics cifs authentication

Figure 41.6 *Sample Output: show cifs-service kerberos-tickets all*

```
bstnA> show cifs-service kerberos-tickets all


Service:        ac1.MEDARCH.ORG

  Start Time(UTC)    Expiry Time(UTC)   Service Principal
  ------------------------------------------------------------------------

Principal:  ac1$@MEDARCH.ORG

  01/20/2011 06:25:52 AM 01/20/2011 04:25:52 PM krbtgt/MEDARCH.ORG@MEDARCH.ORG
        Renew Till:   01/21/2011 06:25:52 AM

Principal:  juser@MEDARCH.ORG

  01/20/2011 06:38:16 AM 01/20/2011 04:25:52 PM HOST/ac1.MEDARCH.ORG@MEDARCH.ORG
  01/20/2011 06:38:16 AM 01/20/2011 04:25:52 PM CIFS/VM-SWP2008S-02@MEDARCH.ORG
  01/20/2011 06:38:21 AM 01/20/2011 04:25:52 PM CIFS/vm-swp2003s2-04@MEDARCH.ORG

Principal:  ac1$@MEDARCH.ORG

  01/20/2011 06:25:52 AM 01/20/2011 04:25:52 PM krbtgt/MEDARCH.ORG@MEDARCH.ORG
        Renew Till:   01/21/2011 06:25:52 AM

Principal:  lfine_md@MEDARCH.ORG

  01/20/2011 06:26:10 AM 01/20/2011 04:25:52 PM HOST/ac1.MEDARCH.ORG@MEDARCH.ORG
  01/20/2011 06:26:10 AM 01/20/2011 04:25:52 PM CIFS/VM-SWP2008S-02@MEDARCH.ORG
  01/20/2011 06:26:10 AM 01/20/2011 04:25:52 PM CIFS/vm-swp2003s2-04@MEDARCH.ORG

Principal:  ac1$@MEDARCH.ORG

  01/20/2011 06:25:52 AM 01/20/2011 04:25:52 PM krbtgt/MEDARCH.ORG@MEDARCH.ORG
        Renew Till:   01/21/2011 06:25:52 AM

Principal:  Administrator@MEDARCH.ORG

  01/20/2011 06:38:22 AM 01/20/2011 04:25:52 PM HOST/ac1.MEDARCH.ORG@MEDARCH.ORG
  01/20/2011 06:38:22 AM 01/20/2011 04:25:52 PM CIFS/VM-SWP2008S-02@MEDARCH.ORG
  01/20/2011 06:38:22 AM 01/20/2011 04:25:52 PM CIFS/vm-swp2003s2-04@MEDARCH.ORG

Principal:  ac1$@MEDARCH.ORG

  01/20/2011 06:25:52 AM 01/20/2011 04:25:52 PM krbtgt/MEDARCH.ORG@MEDARCH.ORG
        Renew Till:   01/21/2011 06:25:52 AM

Principal:  choward_md@MEDARCH.ORG
```

```
  01/20/2011 06:26:06 AM 01/20/2011 04:25:52 PM HOST/ac1.MEDARCH.ORG@MEDARCH.ORG
  01/20/2011 06:26:06 AM 01/20/2011 04:25:52 PM CIFS/VM-SWP2008S-02@MEDARCH.ORG
  01/20/2011 06:26:06 AM 01/20/2011 04:25:52 PM CIFS/vm-swp2003s2-04@MEDARCH.ORG

Principal:  ac1$@MEDARCH.ORG

  01/20/2011 06:25:52 AM 01/20/2011 04:25:52 PM krbtgt/MEDARCH.ORG@MEDARCH.ORG
        Renew Till:   01/21/2011 06:25:52 AM

Principal:  lfine_md@MEDARCH.ORG

  01/20/2011 06:26:07 AM 01/20/2011 04:25:52 PM HOST/ac1.MEDARCH.ORG@MEDARCH.ORG
  01/20/2011 06:26:07 AM 01/20/2011 04:25:52 PM CIFS/VM-SWP2008S-02@MEDARCH.ORG
  01/20/2011 06:26:07 AM 01/20/2011 04:25:52 PM CIFS/vm-swp2003s2-04@MEDARCH.ORG

Principal:  ac1$@MEDARCH.ORG

  01/20/2011 06:25:52 AM 01/20/2011 04:25:52 PM krbtgt/MEDARCH.ORG@MEDARCH.ORG
        Renew Till:   01/21/2011 06:25:52 AM

Principal:  choward_md@MEDARCH.ORG

  01/20/2011 06:38:16 AM 01/20/2011 04:25:52 PM HOST/ac1.MEDARCH.ORG@MEDARCH.ORG
  01/20/2011 06:38:16 AM 01/20/2011 04:25:52 PM CIFS/VM-SWP2008S-02@MEDARCH.ORG
  01/20/2011 06:38:16 AM 01/20/2011 04:25:52 PM CIFS/vm-swp2003s2-04@MEDARCH.ORG

  Total number of principals displayed is 12
  Total number of ticket-granting-tickets cached for this selection is 6
  Total number of service tickets cached for this selection is 18
```

show cifs-service open-files

Purpose Use the show cifs-service open-files command to list the files that CIFS clients are holding open.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show cifs-service open-files {`*`fqdn`* **`[`***`slot.processor`***`] | all}`**

fqdn (1-128 characters) is the fully-qualified domain name for one global server (for example, “www.organization.org”).

slot.processor (optional: for example, 1.3) focuses on the open files served by one NSM slot and processor.

all shows the open files for all CIFS-service offerings.

Guidelines This command shows all of the open files in a CIFS front-end service. A separate table describes each open file with the following fields:

Proc is the NSM processor that is serving the file, in *slot.processor* format.

User IP is the IP address of the client that has the file open.

User Name identifies the client with a username and domain name.

Namespace indicates the ARX namespace where the file resides.

Virtual IP is the IP address that the client is using to access the file. This is established by the virtual server command.

Virtual Path is the pathname of the open file from the root of the ARX volume.

Virtual Share is the name of the CIFS share from the client perspective. This is established by the export (gbl-cifs) command, or by a Windows-management application like MMC (if browsing is enabled for the CIFS service).

Virtual FID is the CIFS file ID for the file, as seen by the client application. Each NSM processor assigns its own set of file IDs, so the same file ID may be reused by multiple processors.

Filer IP is the IP address of the back-end filer that hosts the open file.

Filer Share is the share name at the back-end filer.

Path on Filer is the pathname of the open file from the root of the Filer Share shown above.

Open Mode shows the read/write mode used to open the file. This is either “Read+Write” or “Read.”

Locks is the number of *range locks* held by the client, if any. These are locks for ranges of bytes in the file.

To close an open CIFS file, use close cifs file. To see the client sessions that are connected to the CIFS service, use show cifs-service user-sessions. An authorized Windows client can perform all of these operations from an MMC (or similar) interface, assuming the CIFS service has browsing enabled and the client belongs to a properly-enabled windows-mgmt-auth group.

Samples

`bstnA> **show cifs-service open-files all**`

lists all open files. See Figure 41.7.

`bstnA> **show cifs-service open-files ac1.medarch.org 2.9**`

focuses on files served by a particular processor, 2.9. See Figure 41.8 on page 41-29.

Related Commands

cifs
close cifs file
show cifs-service user-sessions
browsing
windows-mgmt-auth

Figure 41.7 *Sample Output: show cifs-service open-files ac1.medarch.org*

```
bstnA> show cifs-service open-files all
Cifs Open Files
---------------

Proc:  2.7
    User IP:                172.16.100.20
    User Name:              juser@MEDARCH.ORG
    Namespace:              medarcv
    Virtual IP:             192.168.25.15
    Virtual Path:           \rcrds\VIP_wing\fractures\examSchedule.doc
    Virtual Share:          CELEBS
    Virtual FID:            589
    Filer IP:               192.168.25.29
    Filer Share:            CELEBS$
    Path on Filer:          \fractures\examSchedule.doc
    Open Mode:              Read+Write
    Locks:                  0

...
Proc:  2.8
    User IP:                172.16.108.112
    User Name:              Administrator@MEDARCH.ORG
    Namespace:              medarcv
    Virtual IP:             192.168.25.15
    Virtual Path:           \rcrds\2004\recoveryStats\recoveryStats.xls
    Virtual Share:          Y2004
    Virtual FID:            547
    Filer IP:               192.168.25.20
    Filer Share:            Y2004
    Path on Filer:          \recoveryStats\recoveryStats.xls
    Open Mode:              Read+Write
    Locks:                  0

Proc:  2.9
    User IP:                172.16.100.68
    User Name:              lfine_md@MEDARCH.ORG
    Namespace:              medarcv
    Virtual IP:             192.168.25.15
    Virtual Path:           \rcrds\2010\holdback\examSchedule.doc
    Virtual Share:          Y2010
    Virtual FID:            536
    Filer IP:               192.168.25.29
    Filer Share:            Y2010
```

```
    Path on Filer:            \holdback\examSchedule.doc
    Open Mode:                Read+Write
    Locks:                    0

...
Proc:  2.10
    User IP:                  172.16.100.209
    User Name:                choward_md@MEDARCH.ORG
    Namespace:                medarcv
    Virtual IP:               192.168.25.15
    Virtual Path:             \rcrds\holdback\memo.doc
    Virtual Share:            ARCHIVES
    Virtual FID:              535
    Filer IP:                 192.168.25.20
    Filer Share:              histories
    Path on Filer:            \holdback\memo.doc
    Open Mode:                Read+Write
    Locks:                    0

Total number of files displayed is 25
```

Figure 41.8 *Sample Output: show cifs-service open-files ac1.medarch.org 2.9*

```
bstnA> show cifs-service open-files ac1.medarch.org 2.9
Cifs Open Files
---------------

Proc:  2.9
    User IP:                  172.16.100.68
    User Name:                lfine_md@MEDARCH.ORG
    Namespace:                medarcv
    Virtual IP:               192.168.25.15
    Virtual Path:             \rcrds\2010\holdback\examSchedule.doc
    Virtual Share:            Y2010
    Virtual FID:              536
    Filer IP:                 192.168.25.29
    Filer Share:              Y2010
    Path on Filer:            \holdback\examSchedule.doc
    Open Mode:                Read+Write
    Locks:                    0

Proc:  2.9
    User IP:                  172.16.100.68
    User Name:                lfine_md@MEDARCH.ORG
    Namespace:                medarcv
    Virtual IP:               192.168.25.15
    Virtual Path:             \rcrds\2010\holdback\stdProcedure.doc
    Virtual Share:            Y2010
    Virtual FID:              537
    Filer IP:                 192.168.25.20
    Filer Share:              Y2010
    Path on Filer:            \holdback\stdProcedure.doc
    Open Mode:                Read+Write
    Locks:                    0

Total number of files displayed is 2
```

show cifs-service path-cache

Purpose Each NSM processor can keep a cache of file paths in memory. Whenever a CIFS client requests a file or directory path, the NSM processor queries namespace software (on the ACM) for the virtual path and records the answer in its cache. The next request for the same path goes to the cache instead of the namespace software. This decreases the number of repetitive path queries to namespace software, thereby increasing performance. Use the show cifs-service path-cache command to show all file/directory paths currently in the path cache.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show cifs-service path-cache`** ***`namespace`*** **`[`*** ***`volume`*** **`[`*** ***`slot.processor`*** **`]] [detailed]`**

namespace (1-30 characters) is the name of a namespace. Use the show namespace command for a list of all namespaces.

volume (optional, 1-1024 characters) focuses on the path cache for a single volume. The show global-config namespace ns-name command lists all volumes (and other parameters) in a particular namespace.

slot.processor (optional: for example, 2.2) focuses on a single NSM processor's cache.

detailed (optional) shows the additional front-end export, "/acopia#," if it is exported by this CIFS service. This export is automatically generated by a CIFS service with browsing enabled.

Guidelines Use the cifs path-cache command to enable the CIFS path cache for a volume.

At the top of the output is one line per selected volume, indicating whether or not the volume has CIFS path caching enabled. Below these lines is a table with all of the currently-cached file paths. Each file (or directory) path is described with three rows of information.

The first row shows the volume and share (and, possibly subshare) with the virtual path, and the NSM processor that has the path in its cache:

Volume is the name of the volume with this file path.

ShareName is the name of the share as it appears in the volume configuration. (The name of the filer share appears in the next row.)

SubShareName, if applicable, is the name of the subshare where this path resides. A subshare is any share under the imported share; see the documentation for the filer-subshares command.

slot.proc is the NSM processor that is serving the file, in *slot.processor* format.

Guidelines (Cont.) The second row contains more information:

\\Filer\ShareName identifies the back-end share for the cached path. The external-filer name is the one configured on the ARX; use show external-filer for a full list of all configured filers. The share name is the one configured on the filer itself.

Age is the age of the cache entry in seconds. After 120 seconds of non-use, the entry is invalidated; the next search for the path causes a new namespace query. This aging process keeps the cache from occupying excessive memory.

State is the current state of this cache entry. This can be "up" (the path is valid), "init," "pending" (the NSM processor is waiting for a response from the namespace software), "down" (the path information is known to be stale), "none," or "unknown."

The final row identifies the virtual path:

Path shows the path from the client perspective. This always starts from the root of the volume.

Use the show statistics cifs path-cache command to show the total path-cache usage since the last ARX reboot.

Samples

```
bstnA> show cifs-service path-cache medarcv
```

lists all cached paths for the "medarcv" namespace. See Figure 41.9 on page 41-31.

```
bstnA> show cifs-service path-cache insur /claims
```

shows the path cache for a specific volume, "insur~/claims." See Figure 41.10 on page 41-32.

Related Commands

cifs path-cache
show statistics cifs path-cache

Figure 41.9 *Sample Output: show cifs-service path-cache medarcv*

```
bstnA> show cifs-service path-cache medarcv

 CIFS service path cache on volume '/rcrds' is enabled
 CIFS service path cache on volume '/test_results' is enabled
 CIFS service path cache on volume '/lab_equipment' is enabled

Volume                    ShareName                 SubShareName          slot.proc
\\Filer\ShareName                   Age                                 State
Path
-----------------------------------------------------------------


/rcrds                    rx                        CELEBS                  2.7
\\fs4\prescriptions                 0:00:35                             up
/rcrds/VIP_wing

/rcrds                    rx                        Y2004                   2.8
\\fs4\prescriptions                 0:00:13                             up
/rcrds/2004/recoveryStats
```

Figure 41.10 *Sample Output: show cifs-service path-cache insur /claims*

```
bstnA> show cifs-service path-cache insur /claims

 CIFS service path cache on volume '/claims' is enabled

Volume                     ShareName                 SubShareName          slot.proc
\\Filer\ShareName                   Age                                  State
Path
-------------------------------------------------------------
```

show cifs-service transactions

Purpose Use the show cifs-service transactions command to list the active CIFS transactions from a particular client-IP address.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show cifs-service transactions`** ***`ip-addr`***

ip-addr (optional) is the source-IP address of a client machine.

Guidelines This command shows the current CIFS transactions between a client station and the ARX.

CIFS Transactions for the client: is the IP address of the client, entered in the command.

- via Processor shows the NSM processor that is hosting the CIFS session, in *slot.processor* format.

Src Port is the transport protocol (TCP or UDP) and port of the client, in *transport/port* format.

RPC is the name of the Remote Procedure Call invoked by the client transaction.

Status describes the current state of the transaction:

- Fwd Filer means that the NSM is forwarding the client connection to a back-end filer.
- Fwd Volume Group Proc indicates that the NSM is forwarding the client connection to managed-volume software on the control plane.
- Fwd Fastpath Proc means that the namespace software is forwarding the client application back to an NSM processor (the fastpath).

To see the client sessions that are connected to a particular CIFS service, use show cifs-service user-sessions. Use show cifs-service open-files to see all files that are currently open through CIFS.

Sample `bstnA> show cifs-service transactions 172.16.100.20`

shows all currently-active transactions between the PC at 172.16.100.20 and any CIFS services on the switch. See Figure 41.11 for sample output.

Related Commands cifs
show cifs-service user-sessions
show cifs-service open-files

Figure 41.11 *Sample Output: show cifs-service transactions*

```
bstnA> show cifs-service transactions 172.16.100.20
CIFS Transactions for the client:  172.16.100.20 via Processor 2.7

Src Port     RPC         Status
------------ ----------- ---------
TCP/61275    trans2      Unknown
TCP/61275    ntcreate    Unknown
```

show cifs-service user-sessions

Purpose Use the show cifs-service user-sessions command to list the client connections to a CIFS service.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax

```
show cifs-service user-sessions all [summary]
show cifs-service user-sessions all [namespace ns [volume vol]]
show cifs-service user-sessions fqdn [slot.processor]
show cifs-service user-sessions fqdn [namespace ns [volume
vol]]
```

all shows all client sessions with all CIFS services.

summary (optional) shows client-authentication counters for all CIFS services.

namespace (1-30 characters) focuses on client sessions with one of the service's namespaces. Use the show namespace command for a list of all namespaces.

volume (optional, 1-1024 characters) focuses on the client sessions with a single volume. The show global-config namespace *ns-name* command lists all volumes (and other parameters) in a particular namespace.

fqdn (1-128 characters) is the fully-qualified domain name for one CIFS service (for example, "www.organization.org"). This option focuses on the client connections to a single CIFS service.

slot.processor (optional: 2.1-12 on ARX-4000; 1.2-5 on ARX-2000; 1.2 on ARX-500 or ARX-VE) focuses on the sessions with one NSM slot and processor.

Guidelines For each selected cifs service, this command shows a table of current client sessions. The table contains one client session per row.

Proc is the NSM processor that is hosting the CIFS session, in *slot.processor* format.

IP Address is the source IP of the client.

Proto shows which CIFS protocol was used, SMB or SMB2. To enable SMB2 for the ARX, use smb2-allowed.

Auth is the authentication protocol used for the client connection. This is "Kerberos," "NTLMv2," "NTLM," or "Anon." The "Anon" means anonymous access; clients can log into the CIFS service (below) anonymously, but then have extremely limited access to the service's storage and other resources. Anonymous access to the IPC$ share is only supported if the CIFS service is backed by a namespace with cifs anonymous-access.

Sign shows whether or not the client connection is using SMB signing, a CIFS security feature. You can use the signatures command to determine (or change) SMB-signing support at the CIFS service.

Age is the time that the client connection has been up.

Username shows the Windows credentials used by the client.

Total number of users displayed appears at the bottom of the output.

Guidelines (Cont.)

To disconnect a CIFS client, use the drop cifs-service user-session command. Use show cifs-service open-files to see the open files in the service, or close cifs file to close one from the command line.

Guidelines: Summary Output

The **all summary** options show counters for each type of client authentication. Each CIFS service appears in a row with the following counters:

Kerberos shows the number of Kerberos authentications to the service.

NTLMv2 and

NTLM are the number of clients that authenticate with NTLMv2 or NTLM.

Anon shows the number anonymous accesses to this service. Clients can log into the CIFS service anonymously, but then have extremely limited access to the service's storage and other resources. Anonymous access to the IPC$ share is only supported if the CIFS service is backed by a namespace with cifs anonymous-access.

Samples

```
bstnA> show cifs-service user-sessions all
```

lists all client sessions with all CIFS services on the switch. See Figure 41.12 for sample output.

```
bstnA> show cifs-service user-sessions labtests.medarch.org
```

lists all client sessions with one CIFS service on the switch. This service supports SMB2. See Figure 41.13 on page 41-36 for sample output.

```
bstnA> show cifs-service user-sessions all summary
```

shows authentication counters for all CIFS services. See Figure 41.14 on page 41-36 for sample output.

```
bstnA> show cifs-service user-sessions ac1.medarch.org
```

lists all client sessions with the CIFS service at "ac1.medarch.org." See Figure 41.15 on page 41-36 for sample output.

```
bstnA> show cifs-service user-sessions ac1.medarch.org namespace
medarcv volume /rcrds
```

lists all client sessions with a particular volume behind the CIFS service at "ac1.medarch.org." See Figure 41.16 on page 41-37 for sample output.

```
bstnA> show cifs-service user-sessions ac1.medarch.org 2.7
```

narrows the list of client sessions to those that are hosted by a single processor, 2.7. For sample output, see Figure 41.17 on page 41-37.

Related Commands

cifs
show cifs-service open-files
browsing
windows-mgmt-auth
signatures
show statistics cifs authentication

Figure 41.12 *Sample Output: show cifs-service user-sessions all*

```
bstnA> show cifs-service user-sessions all


CIFS Service ac1.MEDARCH.ORG

Proc  IP Address       Proto  Auth      Sign  Age            Username
----  ---------------  -----  --------  ----  -------------  ----------------------------
2.7   172.16.100.20    SMB    NTLM      Yes   00:00:31       juser@MEDARCH.ORG
2.7   172.16.100.20    SMB    Kerberos  Yes   00:00:32       juser@MEDARCH.ORG
2.8   172.16.108.112   SMB    Kerberos  Yes   00:00:17       Administrator@MEDARCH.ORG
2.9   172.16.100.68    SMB    Kerberos  Yes   00:00:38       lfine_md@MEDARCH.ORG
2.10  172.16.100.209   SMB    Kerberos  Yes   00:00:31       choward_md@MEDARCH.ORG


Total number of users displayed is 5
```

Figure 41.13 *Sample Output: show cifs-service user-sessions (SMB2)*

```
bstnA> show cifs-service user-sessions labtests.medarch.org


CIFS Service labtests.MEDARCH.ORG

Proc  IP Address       Proto  Auth      Sign  Age            Username
----  ---------------  -----  --------  ----  -------------  ----------------------------
2.7   172.16.108.17    SMB2   NTLMv2    No    00:00:18       MEDARCH\Administrator
2.9   172.16.100.30    SMB2   NTLMv2    No    00:00:24       MEDARCH\Administrator
2.10  172.16.100.31    SMB2   NTLMv2    No    00:00:19       MEDARCH\Administrator


Total number of users displayed is 3
```

Figure 41.14 *Sample Output: show cifs-service user-sessions all summary*

```
bstnA> show cifs-service user-sessions all summary

Service                        Kerberos    NTLMv2      NTLM      Anon
-------------------------------------------------------------------
ac1.MEDARCH.ORG                       4         0         1         0
labtests.MEDARCH.ORG                  0         0         0         0

Total                                 4         0         1         0
```

Figure 41.15 *Sample Output: show cifs-service user-sessions ac1.medarch.org*

```
bstnA> show cifs-service user-sessions ac1.medarch.org


CIFS Service ac1.MEDARCH.ORG

Proc  IP Address       Proto  Auth      Sign  Age            Username
----  ---------------  -----  --------  ----  -------------  ----------------------------
2.7   172.16.100.20    SMB    NTLM      Yes   00:00:00       administrator@MEDARCH.ORG
2.7   172.16.100.20    SMB    Kerberos  Yes   00:00:00       Administrator@MEDARCH.ORG
2.8   172.16.108.112   SMB    Kerberos  Yes   00:00:06       HP1200B-RED$@MEDARCH.ORG
2.9   172.16.100.68    SMB    Kerberos  Yes   00:00:04       lfine_md@MEDARCH.ORG
```

```
2.9   172.16.100.69    SMB    Kerberos  Yes   00:00:06       HP1200B-BLUE$@MEDARCH.ORG
2.10  172.16.100.209   SMB    Kerberos  Yes   00:00:00       Administrator@MEDARCH.ORG


Total number of users displayed is 6
```

Figure 41.16 *Sample Output: show cifs-service user-sessions ... namespace ... volume ...*

```
bstnA> show cifs-service user-sessions ac1.medarch.org namespace medarcv volume /rcrds


CIFS Service ac1.MEDARCH.ORG

Proc  IP Address       Proto  Auth      Sign  Age            Username
----  ---------------  -----  --------  ----  -------------  ----------------------------
2.7   172.16.100.20    SMB    NTLM      Yes   00:00:02       juser@MEDARCH.ORG
2.7   172.16.100.20    SMB    Kerberos  Yes   00:00:03       juser@MEDARCH.ORG
2.9   172.16.100.68    SMB    Kerberos  Yes   00:00:10       lfine_md@MEDARCH.ORG
2.10  172.16.100.209   SMB    Kerberos  Yes   00:00:03       choward_md@MEDARCH.ORG


Total number of users displayed is 4
```

Figure 41.17 *Sample Output: show cifs-service user-sessions ... 2.7*

```
bstnA> show cifs-service user-sessions ac1.medarch.org 2.7


CIFS Service ac1.MEDARCH.ORG

Proc  IP Address       Proto  Auth      Sign  Age            Username
----  ---------------  -----  --------  ----  -------------  ----------------------------
2.7   172.16.100.20    SMB    NTLM      Yes   00:00:22       juser@MEDARCH.ORG
2.7   172.16.100.20    SMB    Kerberos  Yes   00:00:23       juser@MEDARCH.ORG


Total number of users displayed is 2
```

show fastpath cifs-signatures

Purpose Use the show fastpath cifs-signatures command to show counters and statistics for SMB signing. S*MB signing* is a CIFS security feature for client/server communication, where all packets contain a digital signature that the sender creates and the receiver verifies. This command shows counters for various SMB-signing activities, both on the client side and the filer side.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show fastpath cifs-signatures [slot.processor]`

slot.processor (optional: for example, 1.3) focuses on the signing statistics for one NSM slot and processor. If you omit this, the output contains statistics from all network processors.

Guidelines The default output shows the SMB-signing counters for all network processors. Each processor's counters appear in a separate table, with one row per counter. Each row shows the counter for the client side of the ARX and the filer side.

Outbound Unsigned SMBs is the number of unsigned CIFS packets sent from the ARX. The first number is the number of unsigned packets sent to clients, and the second is the number of unsigned packets sent to back-end filers.

Outbound Signed SMBs is the number of signed CIFS packets sent from the ARX to its clients, followed by the number of signed packets sent to back-end filers.

Inbound Unsigned SMBs counts the unsigned CIFS packets received by the ARX from clients and filers, respectively.

Inbound Verified SMBs is the number of signed CIFS packets received and verified by the ARX.

Inbound SMB Verify Errors counts all inbound CIFS packets that were rejected because their SMB signatures failed verification.

You can use the cifs filer-signatures command to enable, require, or disable SMB signing between a namespace and its back-end filers. The signatures command enables, requires, or disables SMB signing between a CIFS service and its clients.

Samples

```
bstnA(cfg)# show fastpath cifs-signatures
```

shows the SMB-signing counters for all network processors. See Figure 41.18 on page 41-39 for sample output.

```
bstnA(cfg)# show fastpath cifs-signatures 2.8
```

focuses on the SMB-signing counters for a particular network processor. See Figure 41.19 on page 41-41 for sample output.

Related Commands cifs filer-signatures
signatures

Figure 41.18 *Sample Output: show fastpath cifs-signatures*

```
bstnA# show fastpath cifs-signatures

CIFS SMB Signature Statistics for Processor:  2.1

Signature Counts                                Client                  Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                             52                      10
Outbound Signed SMBs                                0                       0
Inbound Unsigned SMBs                              52                      11
Inbound Verified SMBs                               0                       0
Inbound SMB Verify Errors                           0                       0


CIFS SMB Signature Statistics for Processor:  2.2

Signature Counts                                Client                  Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                              0                       0
Outbound Signed SMBs                                0                       0
Inbound Unsigned SMBs                               0                       0
Inbound Verified SMBs                               0                       0
Inbound SMB Verify Errors                           0                       0


CIFS SMB Signature Statistics for Processor:  2.3

Signature Counts                                Client                  Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                             75                     217
Outbound Signed SMBs                              203                       0
Inbound Unsigned SMBs                              84                     225
Inbound Verified SMBs                             195                       0
Inbound SMB Verify Errors                           0                       0


CIFS SMB Signature Statistics for Processor:  2.4

Signature Counts                                Client                  Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                           1773                      48
Outbound Signed SMBs                                0                       0
Inbound Unsigned SMBs                            1773                      51
Inbound Verified SMBs                               0                       0
Inbound SMB Verify Errors                           0                       0


CIFS SMB Signature Statistics for Processor:  2.5

Signature Counts                                Client                  Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                            168                       0
Outbound Signed SMBs                                0                       0
Inbound Unsigned SMBs                             168                       0
Inbound Verified SMBs                               0                       0
Inbound SMB Verify Errors                           0                       0


CIFS SMB Signature Statistics for Processor:  2.6

Signature Counts                                Client                  Filer
-------------------------------------------------------------------------------
```

```
Outbound Unsigned SMBs                          1234                       0
Outbound Signed SMBs                               0                       0
Inbound Unsigned SMBs                           1234                       0
Inbound Verified SMBs                              0                       0
Inbound SMB Verify Errors                          0                       0


CIFS SMB Signature Statistics for Processor:  2.7

Signature Counts                                Client                   Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                            14                     107
Outbound Signed SMBs                             186                       0
Inbound Unsigned SMBs                             18                     116
Inbound Verified SMBs                            182                       0
Inbound SMB Verify Errors                          0                       0


CIFS SMB Signature Statistics for Processor:  2.8

Signature Counts                                Client                   Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                             2                      68
Outbound Signed SMBs                             112                       0
Inbound Unsigned SMBs                              4                      75
Inbound Verified SMBs                            110                       0
Inbound SMB Verify Errors                          0                       0


CIFS SMB Signature Statistics for Processor:  2.9

Signature Counts                                Client                   Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                             8                      61
Outbound Signed SMBs                             109                       0
Inbound Unsigned SMBs                             11                      67
Inbound Verified SMBs                            106                       0
Inbound SMB Verify Errors                          0                       0


CIFS SMB Signature Statistics for Processor:  2.10

Signature Counts                                Client                   Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                             7                      30
Outbound Signed SMBs                              73                       0
Inbound Unsigned SMBs                              9                      34
Inbound Verified SMBs                             71                       0
Inbound SMB Verify Errors                          0                       0


CIFS SMB Signature Statistics for Processor:  2.11

Signature Counts                                Client                   Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                             0                       0
Outbound Signed SMBs                               0                       0
Inbound Unsigned SMBs                              0                       0
Inbound Verified SMBs                              0                       0
Inbound SMB Verify Errors                          0                       0
```

```
CIFS SMB Signature Statistics for Processor:  2.12

Signature Counts                                Client                  Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                             9                      10
Outbound Signed SMBs                              19                       0
Inbound Unsigned SMBs                             11                      12
Inbound Verified SMBs                             17                       0
Inbound SMB Verify Errors                          0                       0
```

Figure 41.19 *Sample Output: show fastpath cifs-signatures 2.8*

```
bstnA# show fastpath cifs-signatures 2.8

CIFS SMB Signature Statistics for Processor:  2.8

Signature Counts                                Client                  Filer
-------------------------------------------------------------------------------
Outbound Unsigned SMBs                             2                      68
Outbound Signed SMBs                             112                       0
Inbound Unsigned SMBs                              4                      75
Inbound Verified SMBs                            110                       0
Inbound SMB Verify Errors                          0                       0
```

show statistics cifs authentication

Purpose Use the show statistics cifs authentication command to show counters and statistics for CIFS-authentication.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show statistics cifs authentication {fqdn | all} [verbose]`

fqdn | **all** is a required choice:

fqdn (1-128 characters) is the fully-qualified domain name for one CIFS service (for example, "www.mystate.gov").

all shows all authentication statistics from all CIFS-services.

verbose (optional) adds an additional table to the output. The table shows details for all CIFS-authentication failures.

Guidelines The default output shows the CIFS-service FQDN, its Windows Domain, and two tables.

Guidelines: Client Authentication Table Client Authentication, the first table, shows Kerberos, NTLMv2, and NTLM counters for front-end CIFS clients. These count the authentications between clients and the CIFS service.

The NTLM and NTLMv2 counts include Netlogon authentications and authentications through a separately-installed Secure Agent. A CIFS service that uses constrained delegation uses Netlogon, but a CIFS service that use unconstrained delegation uses the Secure Agent. You choose between constrained and unconstrained delegation when you run domain-join for the service.

Client authentication success and

Client authentication failure count the number of successful and failed authentication attempts for NTLM, NTLMv2, and Kerberos.

Principals in same realm is a Kerberos-only counter. This is the number of principals (users and hosts) who were in the same Kerberos realm (Windows domain) as the CIFS service. The CIFS-service realm/domain is established in the global-server configuration: use show global server fqdn to see it and windows-domain (gbl-gs) change it.

Principals in trusted realm is another Kerberos-only counter. These are the number of principals in a trusted realm. Trust relationships in realms are established by the local Active-Directory forest, which you configure on the ARX with active-directory-forest and its sub commands. Use show active-directory to show the Active-Directory forest on this ARX.

Principals in realm unknown counts the Kerberos-using principals who were outside any realm (domain) in the Active-Directory forest. This may indicate one or more realms/domains are missing from the ARX's Active-Directory-forest configuration.

Guidelines (Cont.)

Errors contacting Secure Agent are NTLM and/or NTLMv2 errors for a CIFS service that uses unconstrained delegation or is not joined to its domain (see the domain-join command). These may indicate a connection problem with the Secure Agent's DC host. Use show ntlm-auth-server to show the host(s) for the Secure Agent, and use show exports host dc-ip-address connectivity to check connectivity to the DC.

SMB signing incompatibility counts the number of times that a client and the CIFS service could not successfully negotiate SMB signing. In this case, one end of the connection requires SMB signing and the other end of the connection refuses it. You can use the signatures command to set the SMB signing policy for the CIFS service.

Guidelines: Filer Authentication Table

Filer Authentication shows statistics for back-end Windows authentication. These count the authentications between the CIFS service and the back-end CIFS servers. Some of these statistics have different meanings depending on the delegation setting for the CIFS service; when the service joins its Windows domain (see domain-join), it is set for either constrained delegation or unconstrained delegation. Constrained delegation is more secure, and therefore recommended.

Control plane authentication success and

Control plane authentication failure count all authentication attempts that originate from the control plane. The control plane is where volumes, global-servers, and the policy engine run. These count the client requests that require processing at the control plane.

Fast path authentication success and

Fast path authentication failure are the authentications that do not require processing at the control plane; they are processed at the NSM only.

Cross-Realm TGTs granted applies only to Kerberos.

- For a CIFS service that uses constrained delegation, the CIFS service follows the Kerberos S4U protocol and obtains service tickets on behalf of its clients. This counts the number of cross-realm service tickets that the service obtained.
- For unconstrained delegation, the CIFS client presents a Ticket-Granting Ticket (TGT) to the CIFS service, which then uses that TGT to get a Service Ticket for a back-end filer. This counts the number of cross-realm TGTs that the DCs granted.

Cross-Realm TGTs denied are also Kerberos-only statistics, with different meanings for CIFS services with constrained or unconstrained delegation:

- For constrained delegation, this is the number of S4U service tickets that have been denied, where the client was in a different realm (or domain) from the CIFS service. A constrained CIFS service uses the Kerberos Service-for-User (S4U) protocol extension, where the CIFS service gets service tickets to its back-end filers on behalf of each of its Kerberos clients. If these tickets are denied, verify that the CIFS service still has an account configuration at a local DC, and that the account's password key has not expired. If the password key has expired, you must use domain-join to rejoin the CIFS service to its domain.
- For unconstrained delegation, the CIFS client presents a TGT to the CIFS service, which then uses that TGT to get a Service Ticket for a back-end filer. This counts the number of cross-realm TGTs that the DCs denied.

Guidelines: Filer Authentication Table (Cont.)

Cross-Forest TGTs granted, and

Cross-Forest TGTs denied are also Kerberos-only statistics. These only apply to a service with unconstrained delegation. The CIFS client presents a Ticket-Granting Ticket (TGT) to the unconstrained CIFS service, which then uses that TGT to get a Service Ticket for a back-end filer. These fields count the number of cross-forest TGTs that the ARX requested. A cross-forest TGT grants access from one forest to another. Use show active-directory to view all realms (or Windows domains), all forests, their connections to one another, and the Domain Controller(s) for each.

Service tickets granted, and

Service tickets denied are Kerberos-only statistics that apply to both constrained and unconstrained delegation. These are the results from attempting to attain Service Tickets on behalf of clients.

Control plane expired TGTs count the problems that the control plane has encountered with TGTs.

Fast path expired TGTs are TGT errors for NSM-only transactions.

Errors contacting Secure Agent may indicate a connection problem with the Secure Agent's DC host. These counters only apply to a CIFS service that uses unconstrained delegation (or is not joined to its domain). This is an NTLM/NTLMv2 counter. Use show ntlm-auth-server to show the host(s) for the Secure Agent, and use show exports host dc-ip-address connectivity to check connectivity to the DC.

SMB signing incompatibility counts the number of times that a filer and the namespace behind the CIFS service could not successfully negotiate SMB signing. In this case, one end of the connection requires SMB signing and the other end of the connection refuses it. You can use the cifs filer-signatures command to set the SMB signing policy for a namespace.

S4U-to-self tickets granted and

S4U-to-self tickets denied count the successful and failed attempts for the CIFS service to get a service ticket for itself. These only apply to a CIFS service configured for constrained delegation. Such a CIFS service needs a service ticket to itself for each of its Kerberos clients. A constrained CIFS service uses the Kerberos Service-for-User (S4U) protocol extension, which necessitates this ticket. This is the total number of such tickets that have been denied, including the cross-realm tickets shown in the Cross-Realm TGTs denied field above. If these S4U tickets are denied, verify that the CIFS service still has an account configuration at a local DC, and that the account's password key has not expired. If the password key has expired, you must use domain-join to rejoin the CIFS service to its domain.

Guidelines: Verbose Output

If you use the verbose keyword, additional tables appear to display CIFS-authentication failures. The tables display detailed reasons for the last 20 failures.

The Client Authentication section has a table entitled Authentication Failure Reason Table. Each failure appears on two lines with the following fields:

Error Code is an internal code name for the error (for example, KRB5KRB_AP_ERR_MODIFIED).

Error Description is the text for the error (for example, "Message stream modified").

Count is the number of times the error has occurred. Only unique errors appear in this table; this counter increments each time the CIFS service gets each error.

Last Time is the date and time the error was last received.

Last Client IP is the IP address of the client to get the error.

The Filer Authentication section has a similar table. Instead of showing a filer IP in all cases, this attempts to show the Principal involved with each back-end authentication error.

These tables contain up to 20 unique authentication errors. After a service reaches 20 errors, any new errors are dropped until you clear the statistics. The clear statistics cifs authentication command clears all CIFS-authentication statistics, including these tables.

Sample

`bstnA(cfg)# show statistics cifs authentication ac1.medarch.org`

shows CIFS-authentication statistics for the CIFS service at "ac1.medarch.org." See Figure 41.20 on page 41-45 for sample output.

Related Commands

clear statistics cifs authentication
show cifs-service kerberos-tickets
show statistics domain-controller
show statistics cifs work-queues

Figure 41.20 *Sample Output: show statistics cifs authentication ac1.medarch.org*

```
bstnA# show statistics cifs authentication ac1.medarch.org

CIFS service: ac1.MEDARCH.ORG
in Windows domain: MEDARCH.ORG

 Client Authentication
 =====================

  Counter                                    Kerberos      NTLMv2        NTLM
  ------------------------------------------------------------------------------
  Client authentication success                    30           0          15
  Client authentication failure                     0           0           4
  Principals in same realm                         30         n/a         n/a
  Principals in trusted realm                       0         n/a         n/a
  Principals realm unknown                          0         n/a         n/a
  Errors contacting Secure Agent                  n/a           0           0
  SMB signing incompatibility                       0           0           0


 Filer Authentication
 ====================
```

Counter	Kerberos	NTLMv2	NTLM
Control plane authentication success	39	0	0
Control plane authentication failure	0	0	0
Fast path authentication success	31	0	0
Fast path authentication failure	0	0	0
Cross-Realm TGTs granted	0	n/a	n/a
Cross-Realm TGTs denied	0	n/a	n/a
Cross-Forest TGTs granted	0	n/a	n/a
Cross-Forest TGTs denied	0	n/a	n/a
Service tickets granted	49	n/a	n/a
Service tickets denied	0	n/a	n/a
Control plane expired TGTs	0	n/a	n/a
Fast path expired TGTs	0	n/a	n/a
Errors contacting Secure Agent	n/a	0	0
SMB signing incompatibility	0	0	0
S4U-to-self tickets granted	20	n/a	n/a
S4U-to-self tickets denied	0	n/a	n/a

show statistics cifs fastpath

Purpose Use the show statistics cifs fastpath command to show counters and statistics for CIFS servers. The NSM processors keep these statistics.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show statistics cifs fastpath [all | slot.processor]`

all (optional) shows statistics from all NSM processors.

slot.processor (optional: 2.1-12 on ARX-4000; 1.2-5 on ARX-2000; 1.2 on ARX-500 or ARX-VE) focuses on the statistics for one NSM slot and processor. If you omit this, the output contains statistics from all NSM processors; the effect is the same as using the **all** keyword.

Guidelines The output is a table with one row per NSM processor. Each row contains the following fields:

Proc identifies the NSM processor. This is in *slot.processor* format.

Transactions Handled is the sum of front-end and back-end CIFS transactions handled by the processor since the last reboot.

FrontEnd Connections shows the current number of connections from clients.

BackEnd Connections is the current number of connections made to filers on behalf of CIFS clients.

File Info shows the current number of files, pipes, and directories held open on this processor.

File Handles is the total file handles associated with the above files, directories, and/or pipes. This number may be lower than the above number because multiple clients may be accessing the same files.

Samples `bstnA(cfg)# show statistics cifs fastpath`
shows CIFS statistics for all CIFS services. See Figure 41.21 for sample output.

`bstnA(cfg)# show statistics cifs fastpath 2.4`
focuses on a single NSM processor. See Figure 41.22 for sample output.

Related Commands

Figure 41.21 *Sample Output: show statistics cifs fastpath*

```
bstnA# show statistics cifs fastpath


Proc Transactions FrontEnd    BackEnd     File File
     Handled      Connections Connections Info Handles
---- ------------ ----------- ----------- ---- -------
2.1  0            0           0           0   0
2.2  0            0           0           0   0
2.3  219          0           0           0   0
2.4  974          6           0           0   0
```

```
2.5  91          1          0          0   0
2.6  0           0          0          0   0
2.7  0           0          0          0   0
2.8  0           0          0          0   0
2.9  0           0          0          0   0
2.10  0           0          0          0   0
2.11  0           0          0          0   0
2.12  24          0          0          0   0
```

Figure 41.22 *Sample Output: show statistics cifs fastpath 2.4*

```
bstnA# show statistics cifs fastpath 2.4


Proc Transactions FrontEnd    BackEnd     File File
     Handled      Connections Connections Info Handles
---- ------------ ----------- ----------- ---- -------
2.4  974          6           0           0   0
```

show statistics cifs path-cache

Purpose Each NSM processor can keep a cache of file/directory paths to improve CIFS performance. (This applies to SMB, not SMB2.) Use the show statistics cifs path-cache command to show counters and statistics for this cache.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show statistics cifs path-cache`** ***`namespace`*** **`[`*** ***`volume`*** **`[`*** ***`slot.processor`*** **`]]`**

namespace (1-30 characters) is the name of a namespace. Use the show namespace command for a list of all namespaces.

volume (optional, 1-1024 characters) focuses on the path cache for a single volume. The show global-config namespace *ns-name* command lists all volumes (and other parameters) in a particular namespace.

slot.processor (optional: for example, 2.4) focuses on a single NSM processor's cache.

Guidelines Use the cifs path-cache command to enable the CIFS path-cache for a volume.

The output is one or more tables, where each table contains statistics for each volume on each NSM processor. Each NSM processor keeps a separate cache for every managed volume it serves. The header for each table identifies its managed volume and its NSM processor (in *slot.processor* format). The table body contains the following counters:

Total Hits is the number of path queries from CIFS clients where the path was found in the cache.

Total Misses counts the path queries where the answer was not in the cache. These result in a "lookup" query to the namespace software; the path goes into the cache (for future requests) after a successful lookup.

DNAS Lookup Fails shows the number of path lookups that failed. (The namespace software is called *DNAS*.) A lookup can fail when a file or directory's path is in flux due to migrations, filer error, or some other issue. Many of these issues are counted in the fields below.

The next block of fields show the number of path-cache entries that have been declared "invalid," and removed from the cache.

DNAS Invalidates shows the number of paths that were removed from the cache due to an "invalidate" message from the namespace software. The namespace software sends an "invalidate" message when a file or directory's path is changing: during a migration (caused by rules like place-rule, and auto-migrate), during a directory rename, while a managed-volume share is being removed (caused by, for example, remove-share migrate), while the volume is being disabled (no enable (gbl-ns, gbl-ns-vol)), or during a share import (enable (gbl-ns-vol-shr)).

Client Invalidates is the number of paths that have been invalidated (removed from the cache) due to some client action, such as a rename or delete.

Guidelines (Cont.)

Filer Reply Invalidates is the number of paths that have been invalidated by a filer's unexpected "not found" response. This is typically a case where a filer application, such as anti-virus software, has moved a file. If the volume has auto sync files enabled, it responds to this by probing the filer and updating its metadata with the correct path.

Age Invalidates counts the number of paths that were removed from the cache due to a 120-second timeout. Each path stays in the cache for at least 120 seconds before being removed, to avoid wasting valuable memory. The clock restarts every time a client requests the path.

Insert Invalidates shows the number of paths that were removed from the cache to make room for a new path. This only occurs when the path cache exceeds its maximum size; the NSM processor removes the oldest entries first.

Mgmt Invalidates counts the paths that were invalidated due to a CLI or GUI command that destages a volume (nsck ... destage) or rebuilds it (nsck ... rebuild). This invalidates all paths in the volume.

Coll Invalidates is the number of paths that were declared invalid when a case collision was discovered. A case collision is a path on another share that is exactly the same except in letter case (for example, "/myDir/yourDir/yourFile.doc" has a case collision with "/myDir/yourDir/YOURFILE.DOC").

The final block of fields summarizes the path-cache statistics:

Current Entries shows the number of paths in the cache now.

Max Entries shows the maximum size of the cache since the last reboot.

Total Entries is the sum of all path-cache entries since the last reboot.

Use the clear statistics cifs path-cache command to clear the above statistics.

To see the current contents of the path cache, use the show cifs-service path-cache command.

Samples

```
bstnA(cfg)# show statistics cifs path-cache insur
```

shows path-cache statistics for the "insur" namespace. See Figure 41.23 for sample output.

```
bstnA(cfg)# show statistics cifs path-cache medarcv /rcrds
```

focuses on a single volume. See Figure 41.24 on page 41-51 for sample output.

```
bstnA(cfg)# show statistics cifs path-cache medarcv /rcrds 2.6
```

focuses on a single volume and processor. See Figure 41.25 on page 41-52 for sample output.

Related Commands

cifs path-cache
show cifs-service path-cache
clear statistics cifs path-cache

Figure 41.23 *Sample Output: show statistics cifs path-cache insur*

```
bstnA# show statistics cifs path-cache insur
Volume  /claims     Slot.Proc   2.12
---------------------------------------
    Total Hits                      42
    Total Misses                    10

    DNAS Lookup Fails               0

    DNAS Invalidates                0
    Client Invalidates              0
    Filer Reply Invalidates         0
    Age Invalidates                 10
    Insert Invalidates              0
    Mgmt Invalidates                0
    Coll Invalidates                0

    Current Entries                 0
    Max Entries                     8
    Total Entries                   10
```

Figure 41.24 *Sample Output: show statistics cifs path-cache medarcv /rcrds*

```
bstnA# show statistics cifs path-cache medarcv /rcrds
Volume  /rcrds     Slot.Proc   2.2
---------------------------------------
    Total Hits                      34
    Total Misses                    10

    DNAS Lookup Fails               6

    DNAS Invalidates                0
    Client Invalidates              0
    Filer Reply Invalidates         0
    Age Invalidates                 4
    Insert Invalidates              0
    Mgmt Invalidates                0
    Coll Invalidates                0

    Current Entries                 0
    Max Entries                     3
    Total Entries                   10

Volume  /rcrds     Slot.Proc   2.3
---------------------------------------
    Total Hits                      1
    Total Misses                    1

    DNAS Lookup Fails               0

    DNAS Invalidates                0
    Client Invalidates              0
    Filer Reply Invalidates         0
    Age Invalidates                 1
    Insert Invalidates              0
    Mgmt Invalidates                0
    Coll Invalidates                0

    Current Entries                 0
    Max Entries                     1
    Total Entries                   1
```

```
Volume  /rcrds     Slot.Proc   2.4
---------------------------------------
    Total Hits                      44
    Total Misses                    16

    DNAS Lookup Fails               9

    DNAS Invalidates                0
    Client Invalidates              0
    Filer Reply Invalidates         0
    Age Invalidates                 7
    Insert Invalidates              0
    Mgmt Invalidates                0
    Coll Invalidates                0

    Current Entries                 0
    Max Entries                     4
    Total Entries                   16

Volume  /rcrds     Slot.Proc   2.6
---------------------------------------
    Total Hits                      698
    Total Misses                    268

    DNAS Lookup Fails               253

    DNAS Invalidates                0
    Client Invalidates              0
    Filer Reply Invalidates         0
    Age Invalidates                 15
    Insert Invalidates              0
    Mgmt Invalidates                0
    Coll Invalidates                0

    Current Entries                 0
    Max Entries                     9
    Total Entries                   268
```

Figure 41.25 *Sample Output: show statistics cifs path-cache medarcv /rcrds 2.6*

```
bstnA# show statistics cifs path-cache medarcv /rcrds 2.6
Volume  /rcrds     Slot.Proc   2.6
---------------------------------------
    Total Hits                      698
    Total Misses                    268

    DNAS Lookup Fails               253

    DNAS Invalidates                0
    Client Invalidates              0
    Filer Reply Invalidates         0
    Age Invalidates                 15
    Insert Invalidates              0
    Mgmt Invalidates                0
    Coll Invalidates                0

    Current Entries                 0
    Max Entries                     9
    Total Entries                   268
```

show statistics cifs symlinks

Purpose In a multi-protocol (CIFS and NFS) volume, CIFS clients can follow the symbolic links (or symlinks) created by NFS clients. Use the show statistics cifs symlinks command to show counters and statistics for symlink usage by CIFS clients.

Mode (any)

Security Role(s) network-technician, network-engineer, storage-engineer, backup-operator or crypto-officer

Syntax `show statistics cifs symlinks [volume-group id]`

volume-group *id* (optional) narrows the scope to a single volume group. A *volume group* is a failure domain for a group of volumes in the same namespace. If you omit this option, the output shows statistics for all of the volumes on the system.

Guidelines Multi-protocol volumes keep these statistics for de referencing symlinks for its CIFS clients.

To disable all CIFS access to symlinks in a given volume, you can use the cifs deny-symlinks command.

You can use the clear statistics cifs symlinks command to clear these statistics. The statistics also clear after every chassis reboot.

Guidelines: Output The output is a group of tables, one per volume group. Each table contains the following statistics:

Symlink requests dereferenced from backend is the number of symlink reads that required a query to the back-end filer. These are symlinks requested by CIFS clients. The volume software queries the back-end filer once and then caches the symlink target for future CIFS-client access.

Symlink requests dereferenced from cache counts the symlinks that were accessed from the internal cache, without requiring any back-end access.

Total symlink requests dereferenced is the sum of the above two counters. This is the total number of times that CIFS clients accessed symlinks in the given volume group.

Entries in symlink cache shows the current number of symlinks stored in the internal cache. These are symlinks accessed by one or more CIFS clients in the past 2 minutes.

Symlink cache size indicates the amount of memory currently used by the internal symlink cache. This is rounded to the nearest Kilobyte.

Cache hit rate is the percentage of symlinks resolved by accessing the internal cache. The remaining symlinks were resolved by querying a back-end filer.

Guidelines: Output (Cont.)

The remaining statistics count all failed attempts to de reference symlinks for CIFS clients:

Failed symlink requests dereferenced to a dangling path counts all attempts to access a dangling symlink. A *dangling* symlink is one that points to a non-existent file or directory.

Failed symlink requests dereferenced to an absolute path is the number of attempts to access an absolute symlink. An *absolute* symlink is one starts with a slash (/) or back slash (\); for example, "/vol/vol3/flightRecords/2009." A CIFS-client machine invariably has a different root path than the one in the symlink (such as "e:\flightRecords\2009"), so it cannot possibly interpret an absolute symlink.

Failed symlink requests due to filer timeout is the number of back-end queries that failed due to a back-end-filer timeout.

Failed symlink requests due to other filer error counts the back-end queries that failed due to any non-timeout error by the filer.

Samples

`bstnA(cfg)# show statistics cifs symlinks`

shows CIFS symlink statistics for all volume groups on the ARX. See Figure 41.26 for sample output.

`bstnA(cfg)# show statistics cifs symlinks volume-group 4`

focuses on a single volume group. See Figure 41.27 on page 41-55 for sample output.

Related Commands

clear statistics cifs symlinks
cifs deny-symlinks

Figure 41.26 *Sample Output: show statistics cifs symlinks*

```
bstnA# show statistics cifs symlinks

CIFS Symlink Resolution Statistics
==================================

Statistics for all volumes in Volume Group 1
--------------------------------------------

Symlink requests dereferenced from backend:                      0
Symlink requests dereferenced from cache:                        0
                                                            ---------
Total symlink requests dereferenced:                             0

Entries in symlink cache:                                        0
Symlink cache size:                                              0 KB
Cache hit rate:                                                  0%

Failed symlink requests dereferenced to dangling path:           0
Failed symlink requests dereferenced to absolute path:           0
Failed symlink requests due to filer timeout:                    0
Failed symlink requests due to other filer error:                0

Statistics for all volumes in Volume Group 2
--------------------------------------------

Symlink requests dereferenced from backend:                      0
```

```
Symlink requests dereferenced from cache:                        0
                                                               ---------
Total symlink requests dereferenced:                             0

Entries in symlink cache:                                        0
Symlink cache size:                                              0 KB
Cache hit rate:                                                  0%

Failed symlink requests dereferenced to dangling path:           0
Failed symlink requests dereferenced to absolute path:           0
Failed symlink requests due to filer timeout:                    0
Failed symlink requests due to other filer error:                0

Statistics for all volumes in Volume Group 3
--------------------------------------------

Symlink requests dereferenced from backend:                      0
Symlink requests dereferenced from cache:                        0
                                                               ---------
Total symlink requests dereferenced:                             0

Entries in symlink cache:                                        0
Symlink cache size:                                              0 KB
Cache hit rate:                                                  0%

Failed symlink requests dereferenced to dangling path:           0
Failed symlink requests dereferenced to absolute path:           0
Failed symlink requests due to filer timeout:                    0
Failed symlink requests due to other filer error:                0

Statistics for all volumes in Volume Group 4
--------------------------------------------

Symlink requests dereferenced from backend:                      3
Symlink requests dereferenced from cache:                       12
                                                               ---------
Total symlink requests dereferenced:                            15

Entries in symlink cache:                                        4
Symlink cache size:                                              0 KB
Cache hit rate:                                                 80%

Failed symlink requests dereferenced to dangling path:           0
Failed symlink requests dereferenced to absolute path:           0
Failed symlink requests due to filer timeout:                    0
Failed symlink requests due to other filer error:                0
```

Figure 41.27 *Sample Output: show statistics cifs symlinks volume-group 4*

```
bstnA# show statistics cifs symlinks volume-group 4

CIFS Symlink Resolution Statistics
==================================

Statistics for all volumes in Volume Group 4
--------------------------------------------

Symlink requests dereferenced from backend:                      3
Symlink requests dereferenced from cache:                       12
                                                               ---------
Total symlink requests dereferenced:                            15
```

```
Entries in symlink cache:                                    4
Symlink cache size:                                          0 KB
Cache hit rate:                                             80%

Failed symlink requests dereferenced to dangling path:       0
Failed symlink requests dereferenced to absolute path:       0
Failed symlink requests due to filer timeout:                0
Failed symlink requests due to other filer error:            0
```

show statistics cifs work-queues

Purpose

A CIFS volume uses internal work queues to manage its CIFS-related tasks. There is a main work queue for the bulk of client requests, an authentication queue for client-authentication tasks, two queues for communication with the data plane (networking software), and so on. The show statistics cifs work-queues command shows the time that CIFS work items have spent waiting in these work queues, as well as the time used to perform the actual work. You can use this information for troubleshooting CIFS performance issues.

Mode

(any)

Security Role(s)

network-technician, network-engineer, storage-engineer, crypto-officer, or operator

Syntax

```
show statistics cifs work-queues volume-group vg-id
show statistics cifs work-queues instance instance-id
```

volume-group ***vg-id*** (optional, 1-255) chooses a volume group. A *volume group* is a failure domain for a group of volumes in the same namespace. You can use show volume-group for a list of all volume groups on the system, and to see which volumes are assigned to each volume group.

instance ***instance-id*** (optional) chooses a namespace by its instance ID. Instance IDs often appear in syslog messages, which you can view with show logs syslog. You can can also see instance IDs with show namespace all, which shows full details on all namespaces in the system.

Guidelines

The clear statistics cifs work-queues command clears most of these statistics, and clear statistics filer clears the rest of them. A chassis reboot clears all of them at once.

There are some related CLI commands that can aid in troubleshooting CIFS connections. Use show cifs-service user-sessions to see the which CIFS clients are currently connected to the ARX. Use show cifs-service client-activity to see how many back-end session, tree connections, and files they currently have open. The show cifs-service open-files command shows details about currently-open files. For details about currently-held Kerberos tickets, use the show cifs-service kerberos-tickets command. The show statistics domain-controller command examines the NTLM-related communication with domain controllers in the network.

Guidelines: Output

The output is a series of tables. The work queues only function in managed volumes that support CIFS. The tables are empty for a volume group or namespace that only supports NFS-only volumes and/or direct volumes.

Guidelines: CIFS Work Queue Histograms

CIFS work queue pool statistics shows the usage statistics for each CIFS work queue in the chosen volume group or namespace instance.

Reset... shows the last time these statistics were cleared. The software resets these statistics whenever someone issues the clear statistics cifs work-queues command or reboots the chassis.

Under the Reset line, a separate sub table appears for each work queue:

- Work list... shows the name and purpose of the current work queue.
- *n* items active or waiting... displays the number of work items currently in the queue, along with the highest number ever in the queue at one time. The time stamp is for the most-recent time when the queue had its highest number of work items.
- *x* items completed, *y* discarded is the number of completed and discarded items processed in this work queue. A completed item may have succeeded or failed. Items are discarded from the queue only when system resources are low and traffic is high, and the system only discards items that have waited a long time. A discard occurs before any thread begins executing the work item.
- Average... shows the average time that completed items spent waiting in the queue, and the average time that the volume software spent processing the work items after they were taken off of the queue. The times are measured in microseconds; there are 1,000,000 microseconds in 1 second.
- Time waiting in queue is a histogram showing various time amounts (0.001 second, 0.01 second, 0.1 second, and so on) and the number of work queue items that waited in the queue for that amount of time. The smallest time period is 0.001 seconds, which is 1 millisecond or 1,000 microseconds. This histogram only appears if at least one work item was in the queue since the last Reset time.
- Time processing on thread is a similar histogram, showing the time that threads spent processing the work-queue items after having taken them off the queue.

Work-queue status messages, if any, appear after all of the work-queue tables. Each of these is a declaration that a work item was excessively "slow," was slow long enough to be declared "stuck," or eventually got "done" after being slow or stuck. Each event appears in its own row with the event type (SLOW, STUCK, or DONE), a time stamp for when the work item reached this state, the name of the work queue, the time spent working on the item, the name of the CIFS command sent to the data plane, and any other relevant information about the work item.

Guidelines: Processing-Time Table

Average processing time per operation type is a table of Server Message Blocks (SMBs, also known as CIFS Commands) processed through the above queues. These are commands that originate from ARX clients. You can use this table to determine if some client requests take longer than others.

Reset... shows the last time these statistics were cleared. The software resets these statistics whenever someone issues the clear statistics cifs work-queues command or reboots the chassis.

Under the Reset line, a table of SMBs appears. These are all the SMBs processed by the work queues since the Reset time. Each SMB appears in its own row with the following columns:

- CIFS SMB is the name of the SMB,
- Count is the number of these SMBs received from clients,
- Avg time (uSec) is the average number of microseconds (millionths of a second) required to fully process the request. This is the time in the work queue plus the time for the thread to process the work item.

All SMBs shows the total count for all of the above SMBs, and the average time to process them.

Guidelines: Filer-RTT Table

Average file server round-trip time per operation type is another table of SMBs forwarded to back-end filers and servers. This table helps to determine if a filer or server is slow.

Reset... shows the last time these statistics were cleared. The software resets these statistics whenever someone issues the clear statistics filer command or reboots the chassis. This is a different clear command than the one used for the other statistics.

Under the Reset line, a table of SMBs appears. These are all the SMBs sent to the filers since the Reset time. Each SMB appears in its own row with the following columns:

- CIFS SMB is the name of the SMB,
- Count is the number of these SMBs sent to filers,
- RTT (uSec) is the average round-trip time (RTT) to and from the filer.

All SMBs shows the total count for all of the above SMBs, and the overall average RTT for them.

The next two rows show the total Response Timeouts and Network Errors, if any.

Guidelines: Authentication Table

Time spent in authentication shows the time spent with CIFS-client authentication. (For counters of various authentication tasks, you can use the show statistics cifs authentication command.)

Reset... shows the last time these statistics were cleared. The software resets these statistics whenever someone issues the clear statistics cifs work-queues command or reboots the chassis.

Under the Reset line, a separate sub table appears for each authentication type: Kerberos, Netlogon (NTLM/NTLMv2 for CIFS services that use constrained delegation (see domain-join)), or ARX Secure Agent (NTLM/NTLMv2 for CIFS services that use unconstrained delegation or no delegation).

Kerberos service ticket requests focuses on Kerberos authentications.

- Total calls is the number of Kerberos tickets requested from the ARX-CIFS service.
- Elapsed time in call is a histogram showing various time amounts (0.001 second, 0.01 second, 0.1 second, and so on) and the number of service-ticket requests that took that amount of time. The smallest time period is 0.001 seconds, which is 1 millisecond or 1,000 microseconds. This histogram does not appear unless at least one client authenticated with Kerberos.
- Average time is the average time for all service-ticket requests.

Netlogon requests shows the time taken for NTLM and NTLMv2 requests through Netlogon. These occur for a front-end CIFS service that uses constrained delegation. You choose the delegation type when you run domain-join for the CIFS service.

- Total calls is the number of Netlogon requests received by the ARX-CIFS service.
- Elapsed time in call is a histogram showing various time amounts (0.001 second, 0.01 second, 0.1 second, and so on) and the number of Netlogon requests that took that amount of time. The smallest time period is 0.001 seconds, which is 1 millisecond or 1,000 microseconds. This histogram does not appear unless at least one client authenticated with NTLM or NTLMv2 through a CIFS service with constrained delegation.
- Average time is the average time for all Netlogon requests.

ARX Secure Agent requests shows the time taken for NTLM and NTLMv2 requests through the ARX Secure Agent (see ntlm-auth-server (gbl-ns)). These occur for a front-end CIFS service that uses unconstrained delegation or no delegation. You choose a delegation type when you run domain-join for the CIFS service.

- Total calls is the number of NTLM and NTLMv2 requests received by the unconstrained ARX-CIFS service.
- Elapsed time in call is a histogram showing various time amounts (0.001 second, 0.01 second, 0.1 second, and so on) and the number of NTLM + NTLMv2 requests that took that amount of time. The smallest time period is 0.001 seconds, which is 1 millisecond or 1,000 microseconds. This histogram does not appear unless at least one client authenticated with NTLM or NTLMv2 through a CIFS service with unconstrained delegation.
- Average time is the average time for all NTLM and NTLMv2 requests.

Guidelines: Authentication Table (Cont.)

Table of longest recorded round-trip times shows up to 20 clients who took the longest time for authentication. Each authentication session appears in one row of the table, with the following fields:

- Timestamp is the time that the client started the authentication process.
- Type is NL (Netlogon request), KRB (Kerberos ticket request), or ASA (ARX Secure Agent request).
- Time (uSec) is the time consumed by the authentication request, in microseconds.
- Principal Name identifies the CIFS client that invoked the authentication request. You can use show cifs-service user-sessions to see the clients that are currently connected to the ARX.

Sample

`bstnA(cfg)# show statistics cifs work-queues volume-group 2`

shows the work-queue statistics for volume group 2. See Figure 41.28 on page 41-61 for sample output.

Related Commands

clear statistics cifs work-queues
clear statistics filer
domain-join
show cifs-service user-sessions
show cifs-service client-activity
show cifs-service open-files
show cifs-service kerberos-tickets
show statistics domain-controller
show statistics cifs authentication
show active-directory status

Figure 41.28 *Sample Output: show statistics cifs work-queues*

```
bstnA# show statistics cifs work-queues volume-group 2

Statistics for instance id 3 collected at 03/01/2012 01:29:47 -0500


CIFS work queue pool statistics
-------------------------------
Reset at 03/01/2012 00:38:17 -0500

  Work list "main": used for most client requests
    0 items active or waiting; highest was 3 at 03/01/2012 01:29:42 -0500
    1471 items completed, 0 discarded
    Average 20 uSec queued, 1049 uSec processing
    Time waiting in queue:
      Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
      Count:      1470         1         0         0         0         0         0
    Time processing on thread:
      Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
      Count:      1348        89        31         3         0         0         0

  Work list "auth": used for client authentication
    0 items active or waiting; highest was 2 at 03/01/2012 01:13:06 -0500
    115 items completed, 0 discarded
    Average 27 uSec queued, 1321 uSec processing
    Time waiting in queue:
      Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
```

```
    Count:       115         0         0         0         0         0         0
  Time processing on thread:
    Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
    Count:        97        15         3         0         0         0         0

Work list "setup": authentication support for data plane
  0 items active or waiting; highest was 2 at 03/01/2012 01:27:55 -0500
  35 items completed, 0 discarded
  Average 37 uSec queued, 7081 uSec processing
  Time waiting in queue:
    Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
    Count:        35         0         0         0         0         0         0
  Time processing on thread:
    Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
    Count:         1        27         7         0         0         0         0

Work list "special": internal notifications, including client disconnect
  0 items active or waiting; highest was 1 at 03/01/2012 01:28:31 -0500
  36 items completed, 0 discarded
  Average 40 uSec queued, 172 uSec processing
  Time waiting in queue:
    Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
    Count:        36         0         0         0         0         0         0
  Time processing on thread:
    Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
    Count:        36         0         0         0         0         0         0

Work list "no-op": processes keepalives from data plane
  0 items active or waiting; highest was 12 at 03/01/2012 00:38:22 -0500
  3505 items completed, 0 discarded
  Average 17 uSec queued, 23 uSec processing
  Time waiting in queue:
    Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
    Count:      3503         2         0         0         0         0         0
  Time processing on thread:
    Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
    Count:      3505         0         0         0         0         0         0

No work queue status messages; no significant events have occurred


Average processing time per operation type
------------------------------------------
Reset at 03/01/2012 00:38:11 -0500

  CIFS SMB                              Count   Avg time (uSec)

  NTCreateAndX                             11             24108
  Transaction2/FindFirst2                 108              2543
  DeleteDir                                 3              2211
  SessionSetupAndX                         72              1929
  Transaction2/QueryPathInfo               66              1193
  Delete                                    6              1136
  Close                                    10               374
  TreeConnectAndX                          65               318
  Negotiate                                43               101
  Transaction2/QueryFSInfo                839                42
  TreeDisconnect                            6                42
  FindClose2                               66                14
  QueryInfoDisk                             5                14

  All SMBs                               1300               644
```

```
Average file server round-trip time per operation type
------------------------------------------------------
Reset at 03/01/2012 00:38:17 -0500

  CIFS SMB                                Count        RTT (uSec)

  SessionSetupAndX                           43              1207
  Transaction2/QueryFSInfo                 2078               680
  NTCreateAndX                              704               679
  Negotiate                                  31               667
  Delete                                      6               619
  Rename                                      1               594
  Transaction2/QueryPathInfo                 66               559
  NTTransact/Create                          31               522
  TreeConnectAndX                            45               511
  Transaction2/FindFirst2                   725               388
  DeleteDir                                   4               367
  TreeDisconnect                             32               319
  Transaction2/SetFileInfo                   35               312
  Transaction2/QueryFileInfo                 69               279
  NTTransact/SetSecurity                      4               272
  NTTransact/QuerySecurity                  421               266
  Close                                     545               264
  FindClose2                                455               259

  All SMBs                                 5295               518

  No response timeouts
  No network errors


Time spent in authentication
----------------------------
Reset at 03/01/2012 00:39:16 -0500

  Kerberos service ticket requests:
    Total calls: 65
    Elapsed time in call:
      Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
      Count:        19        39         7         0         0         0         0
    Average time: 0.005 secs

  Netlogon requests:
    Total calls: 12
    Elapsed time in call:
      Secs:    < 0.001    < 0.01     < 0.1       < 1      < 10     < 100    Longer
      Count:         0         9         3         0         0         0         0
    Average time: 0.009 secs

  ARX Secure Agent requests:
    Total calls: 0
    Average time: none

  Table of longest recorded round-trip times:
    Timestamp            Type   Time (uSec)  Principal Name
    03/01/2012 01:12:57  NL           20596  MEDARCH.ORG\lab
    03/01/2012 01:13:10  KRB           7767  mhoward_md@MEDARCH.ORG
    03/01/2012 01:13:12  KRB           7736  lfine_md@MEDARCH.ORG
    03/01/2012 01:12:57  KRB           7832  lab@MEDARCH.ORG
    03/01/2012 01:12:59  KRB           7433  lab@MEDARCH.ORG
```

```
03/01/2012 01:27:50  KRB           8922  administrator@MEDARCH.ORG
03/01/2012 01:28:01  KRB          14077  Administrator@MEDARCH.ORG
03/01/2012 01:27:51  KRB           9743  Administrator@MEDARCH.ORG
03/01/2012 01:27:55  KRB          19717  lfine_md@MEDARCH.ORG
03/01/2012 01:28:20  KRB           7980  Administrator@MEDARCH.ORG
03/01/2012 01:28:15  KRB          10719  Administrator@MEDARCH.ORG
03/01/2012 01:13:05  KRB          15207  Administrator@MEDARCH.ORG
03/01/2012 01:13:13  KRB           7944  choward_md@MEDARCH.ORG
03/01/2012 01:27:57  KRB          13125  choward_md@MEDARCH.ORG
03/01/2012 01:27:55  KRB          15093  Administrator@MEDARCH.ORG
03/01/2012 01:13:06  NL           51422  administrator@MEDARCH.ORG
03/01/2012 01:13:14  KRB           7338  lfine_md@MEDARCH.ORG
03/01/2012 01:27:53  NL           19884  mhoward_md@MEDARCH.ORG
03/01/2012 01:12:59  KRB           7438  lab@MEDARCH.ORG
03/01/2012 01:27:55  KRB          16783  juser@MEDARCH.ORG
```

show statistics domain-controller

Purpose Use the show statistics domain-controller command to view NTLM-usage statistics for a particular domain controller (DC).

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax `show statistics domain-controller` ***ip-address***

ip-address identifies a DC .

Guidelines This command shows the statistics for communication with a particular DC. The statistics are divided into two categories: statistics for NTLM Netlogon authentications and statistics for LDAP pings.

The NETLOGON statistics only apply to CIFS services that use *constrained delegation*, which is chosen when domain-join is performed for the service. If you use unconstrained delegation for your service, but still support NTLM or NTLMv2 for the service's clients, use the show ntlm-auth-server command to find NTLM-authentication statistics. The LDAP-ping statistics in this command apply to any CIFS service, whether or not it uses constrained delegation.

This command shows all the counters for this NTLM-related communication with a particular DC. You can use the clear statistics domain-controller command to clear these statistics. The statistics also clear after every chassis reboot.

Guidelines: Output The output is divided into two tables: NTLM NETLOGON Authentication Statistics and LDAP Ping Statistics.

Guidelines: Netlogon Authentication Statistics The NTLM NETLOGON Authentication Statistics table only applies to CIFS services that use constrained delegation, as explained above. The CIFS services that can use this DC are any that are in its Windows domain.

An ARX CIFS service uses Netlogon service for NTLM and NTLMv2 authentication. Through the Netlogon protocol, each CIFS service sets up a secure channel with the DC. For each client that authenticates to the CIFS service with NTLM or NTLMv2, the service sends a client's username and password through that encrypted channel. If the password is correct, the DC returns the domain security identifiers (SIDs) and access rights for the user back through the secure channel.

Domain Controller IP is the IP address you entered in the command.

Last Reset *time* shows the last time these counters were cleared, either with a chassis reboot or with the clear statistics domain-controller command. All of the counters below started at this time.

Successful NTLM and

Successful NTLMv2 count the number of successful NTLM authentications at the selected DC.

Guidelines: Netlogon Authentication Statistics (Cont.)

Failed NTLM is a total number of NTLM failures at the DC. The counters below show the specific failures.

- No Such User counts the clients who used an unidentified username for the Windows domain.
- Bad Password is the number of failures due to an incorrect password for a valid username.
- Account Disabled shows the number of failures due to an administratively-disabled user account.
- No logon server counts the number of failures due to a service interruption. Each time, the ARX service was unable to reach the DC's Netlogon server. You can use show active-directory status to check the status of the DC connection, and you can set a preference for another DC in the same Windows Domain with forest-root ... preferred, child-domain ... preferred, or tree-domain ... preferred.
- Ntlm Blocked is the number of times the DC blocked NTLM access. This can be controlled at the DC with Local Security Policy; the specific Local Policy settings are "Network Security:Restrict NTLM:*policy-name*."
- Other shows the number of failed authentications outside any of the above categories.

Failed NTLMv2 is a total number of NTLMv2 failures at the DC. The counters below this summary field break down the specific failures, as described above.

Schannel Inits are the number of times that the Secure Channel between the CIFS service an the DC has been re-initialized.

Avg Request SRT (mSec),

Min Request SRT (mSec), and

Max Request SRT (mSec) show the average, minimum, and maximum round-trip times for NTLM authentications. This measures the elapsed time the Netlogon agent uses to process an authentication request. This is the processing time within the agent itself plus the communication time between the agent and the DC. The times are measured in milliseconds (there are 1,000 milliseconds in one second).

Avg DC SRT (mSec),

Min DC SRT (mSec), and

Max DC SRT (mSec) show the average, minimum, and maximum round-trip times for NTLM authentications to this DC. This measures the round trips between the internal Netlogon Agent and the DC. These times are also measured in milliseconds.

Guidelines: LDAP Ping Statistics

The LDAP Ping Statistics table shows the results of periodic pings to the DC. The pings measure the ability to reach the DC as well as the latency of the DC connection:

Each ping-statistics table contains the following information:

Last Reset (Local Time) *time* shows the last time these counters were cleared with clear statistics domain-controller or a chassis reboot. All of the counters below started at this time.

Total Count is the number of LDAP pings to the DC.

Success Count and

Error Count show the numbers of successful and failed pings.

Average RTT (uSec) is the average round-trip time for LDAP pings. This is measured in microseconds. There are 1,000,000 microseconds in one second.

RTT Histogram is a sub-table to display the numbers of pings in certain time ranges. Each row defines a time range in microseconds: 0 to 1000 microseconds, 1000 to 2000 microseconds, and so on. At the end of each row is the number of LDAP pings whose round-trip times where in this range.

Sample

`bstnA(cfg)# show statistics domain-controller 192.168.25.102`

shows NTLM statistics for the DC at 192.168.25.102. See Figure 41.29 for sample output.

Related Commands

clear statistics domain-controller
show ntlm-auth-server

Figure 41.29 *Sample Output: show statistics domain-controller ...*

```
bstnA# show statistics domain-controller 192.168.25.102

Domain Controller IP : 192.168.25.102

NTLM NETLOGON Authentication Statistics
---------------------------------------
Last Reset    06/22/2011 01:13:07 -0400

    Successful NTLM               4
    Successful NTLMv2             0
    Failed NTLM                   0
      No Such User                0
      Bad Password                0
      Account Disabled            0
      No logon server             0
      Ntlm Blocked                0
      Other                       0
    Failed NTLMv2                 0
      No Such User                0
      Bad Password                0
      Account Disabled            0
      No logon server             0
      Ntlm Blocked                0
      Other                       0
    Schannel Inits                1
    Avg Request SRT (mSec)      166
    Min Request SRT (mSec)        1
    Max Request SRT (mSec)      664
    Avg DC SRT (mSec)           166
    Min DC SRT (mSec)             1
```

```
    Max DC SRT (mSec)            664


LDAP Ping Statistics
--------------------

  Last Reset (Local Time): 06/22/2011 01:02:57 -0400

  Total Count                    57
  Success Count                  56
  Error Count                     1
  Average RTT (uSec)            611


  RTT Histogram
  -------------

  Lower (uSec)       Upper (uSec)       Counts

         0               1000             55
      1000               2000              1
      2000               4000              0
      4000               8000              0
      8000              16000              0
     16000              32000              0
     32000              64000              0
     64000             128000              0
    128000             256000              0
    256000             512000              0
    512000            1024000              0
   1024000            2048000              0
   2048000           Infinite              0
```

show statistics domain-controller load-balancing

Purpose Whenever a Windows domain is served by multiple domain controllers (DCs), the ARX rotates its queries between the "active" DCs in that group. Use the show statistics domain-controller load-balancing command to view the statistics for using these DCs. You can use this output to assess the usage of one DC over another.

Mode (any)

Security Role(s) crypto-officer, storage-engineer, network-engineer, network-technician, or operator

Syntax **`show statistics domain-controller load-balancing`**

Guidelines This command shows the statistics for communication with your DCs. You can use this output to compare the number of Kerberos requests to each DC.

The output contains one table for every Windows Domain in the Active Directory. (You can use the active-directory update seed-domain command to discover all of the domains and DCs in the current Active Directory.) Each table contains one row per DC in that domain, with the following columns:

DC IP Address identifies the DC.

Kerberos Requests is the number of Kerberos requests sent to the DC since the last reboot, or since someone cleared the statistics with clear statistics domain-controller load-balancing.

Preferred is either Yes or No. The ARX always directs its Kerberos requests to a preferred DC if any are reachable. DCs at the same AD site as the ARX are preferred by default. You can also use the child-domain, forest-root, or tree-domain command to manually set a DC preference.

These statistics all restart at 0 (zero) after the ARX reboots. You can manually reset these statistics with the clear statistics domain-controller load-balancing command.

Sample `bstnA(cfg)# show statistics domain-controller load-balancing`
shows Kerberos-usage statistics for all DCs behind the "bstnA" ARX. See Figure 41.30 for sample output.

Related Commands clear statistics domain-controller load-balancing

Figure 41.30 *Sample Output: show statistics domain-controller load-balancing*

```
bstnA# show statistics domain-controller load-balancing

Domain: ADK.WELLS.ME.ORG

DC IP Address    Kerberos Requests   Preferred
--------------   -----------------   ----
172.16.110.8     0                   No
172.16.110.5     0                   No

Domain: YORK.WELLS.ME.ORG

DC IP Address    Kerberos Requests   Preferred
--------------   -----------------   ----
```

```
172.16.120.22    0                   No
172.16.120.5     0                   No

Domain: BOSTONCIFS.FDTESTNET.NET

DC IP Address    Kerberos Requests   Preferred
--------------   ------------------  ----
10.19.230.88     0                   Yes
10.19.230.94     0                   Yes
10.51.7.2        0                   Yes

Domain: BOSTONMED.ORG

DC IP Address    Kerberos Requests   Preferred
--------------   ------------------  ----
172.16.74.89     0                   Yes
172.16.74.88     0                   Yes

Domain: WESTCOAST.MEDARCH.ORG

DC IP Address    Kerberos Requests   Preferred
--------------   ------------------  ----
192.168.202.9    0                   Yes
192.168.202.10   0                   Yes
192.168.202.11   0                   Yes
10.51.2.4        0                   Yes

Domain: MEDARCH.ORG

DC IP Address    Kerberos Requests   Preferred
--------------   ------------------  ----
192.168.25.109   592                 Yes
192.168.25.110   624                 Yes
192.168.25.111   625                 Yes
192.168.25.102   632                 Yes
192.168.25.104   633                 Yes

Domain: FDTESTNET.NET

DC IP Address    Kerberos Requests   Preferred
--------------   ------------------  ----
172.16.168.22    0                   Yes
172.16.168.21    0                   Yes

Domain: MA.NE.MEDARCH.ORG

DC IP Address    Kerberos Requests   Preferred
--------------   ------------------  ----
192.168.25.105   0                   Yes
192.168.25.103   0                   Yes

Domain: MCNIELS.VT.COM

DC IP Address    Kerberos Requests   Preferred
--------------   ------------------  ----
172.16.240.70    0                   No
172.16.240.88    0                   No

Domain: ATLANTIC.ME.ORG

DC IP Address    Kerberos Requests   Preferred
--------------   ------------------  ----
```

```
172.16.210.14    0                  No
172.16.210.7     0                  No

Domain: NE.MEDARCH.ORG

DC IP Address    Kerberos Requests   Preferred
--------------   -----------------   ----
172.16.124.19    0                  Yes
172.16.124.73    0                  Yes

Domain: WELLS.ME.ORG

DC IP Address    Kerberos Requests   Preferred
--------------   -----------------   ----
172.16.108.139   0                  No
172.16.108.136   2                  No

Domain: BSH.ATLANTIC.ME.ORG

DC IP Address    Kerberos Requests   Preferred
--------------   -----------------   ----
172.16.110.11    0                  No
172.16.210.9     0                  No

Domain: VT.COM

DC IP Address    Kerberos Requests   Preferred
--------------   -----------------   ----
172.16.213.8     0                  No
172.16.213.79    0                  No
```

show subshare-cache

Purpose A managed volume that supports filer-subshares keeps all of its subshare information in cache memory. This decreases the number of RPC calls between the ARX and its back-end filers during sync subshare operations. You can use this command to show the current contents of this cache.

Mode priv-exec

Security Role(s) network-technician, network-engineer, storage-engineer or crypto-officer

Syntax `show subshare-cache [filer ext-filer-name] [report prefix]`

filer *ext-filer-name* (optional, 1-64 characters) focuses on a particular external filer. This is the name of the filer in the ARX configuration. For a list of configured external filers, use show external-filer.

report *prefix* (optional, 1-64 characters) sends the subshare-cache output to a report instead of the screen. The CLI displays the full report name after you enter the command. The report is named as follows:

prefix_yyyymmddHHMM.rpt, where *prefix* is chosen here and the rest of the filename is the current date and time.

Default(s) None

Guidelines A *filer subshare* is any CIFS share that is inside an imported CIFS share. A client who connects to a front-end subshare, if the subshares are configured as described in the filer-subshares documentation, is passed through the managed volume directly to a corresponding subshare on a back-end filer. The filer can then enforce its subshare ACL, as opposed to the top-level ACL of the imported share. Each front-end subshare (visible to your CIFS clients) maps to one or more back-end subshares. The state of every filer share (and share ACL) resides in a memory cache. You can use this command to show the cache, or the cache for one filer.

You have the option to clear this cache for a filer where the ACLs and/or shares have changed since import. Use the clear subshare-cache command to clear the subshare cache.

Guidelines: Output

The output is a series of nested tables, one per back-end filer. Each filer has a table with the following heading:

Subshare cache for filer *ip-address*

Each filer table contains a separate table for each of its imported shares:

Subshares nested below import share bulkstorage on filer *ip-address* is the heading for the first table. Each share contains its own sub table, with several fields to describe the share:

- Share is the name of the share in "*ip-address**share-name*" format. CIFS clients can use this path to access the storage.
- Path is the file-system path on the filer itself, including the drive letter. If you log into the filer, you can use this path to directly access the storage.
- Import is the name of the import share, the one that was imported into a managed volume. This is typically the container share for the current subshare. This says "[This is the imported share.]" for each top-level share.
- RelPath only appears for a subshare. This is the relative path to the subshare, starting from the root of the imported share.
- ACL is the Access Control List (ACL) for this share or subshare. This must match the ACL on all matching subshares. (A matching subshare is a subshare at the same relative path of another import share.)

After all imported shares and subshares, each filer is re-iterated to show all of its un imported shares (called *orphan* shares):

Orphan (unnested) share caches for filer *ip-address* is the heading for the second filer table. These are shares outside the imported share(s). Each of these shares contains a similar sub table with these fields:

- Share and.
- Path are described above.
- Remark only appears for special default shares. This is a note about a share, such as "Default share" for the root share on a disk drive or "Remote Admin" for the special ADMIN$ share.
- ACL is also described above.

Sample

```
bstnA# show subshare-cache
```

shows the full subshare-cache. See Figure 41.31 on page 41-74 for sample output.

```
bstnA# show subshare-cache filer fs2 report fs2_sbshrs

Creating subshare-cache report: fs2_sbshrs_201004290537.rpt
```

shows the subshare-cache for the external filer named "fs2," and sends it to a report.

Related Commands

filer-subshares
sync subshares from-namespace
sync subshares from-service
clear subshare-cache

Figure 41.31 *Sample Output: show subshare cache*

```
bstnA# show subshare cache

Subshare cache for filer 192.168.25.27
------------------------------------------------------------------------------

  Subshares nested below import share bulkstorage on filer 192.168.25.27
  ------------------------------------------------------------

      Share : \\192.168.25.27\bulkstorage
       Path : e:\exports\bulkstorage
     Import : [[This is the imported share.]]
        ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
              37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
              0{Everyone})

      Share : \\192.168.25.27\Y2004
       Path : e:\exports\bulkstorage\2004
     Import : bulkstorage
    RelPath : 2004
        ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
              37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
              0{Everyone})

      Share : \\192.168.25.27\Y2005
       Path : e:\exports\bulkstorage\2005
     Import : bulkstorage
    RelPath : 2005
        ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
              37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
              0{Everyone})

      Share : \\192.168.25.27\Y2010
       Path : e:\exports\bulkstorage\2010
     Import : bulkstorage
    RelPath : 2010
        ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
              37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
              0{Everyone})

      Share : \\192.168.25.27\MP3S
       Path : e:\exports\bulkstorage\2011\mp3downloads
     Import : bulkstorage
    RelPath : 2011\mp3downloads
        ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
              37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
              0{Everyone})

      Share : \\192.168.25.27\CELEBS$
       Path : e:\exports\bulkstorage\VIP_wing
     Import : bulkstorage
    RelPath : VIP_wing
        ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
              37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
              0{Everyone})

Subshare cache for filer 192.168.25.20
------------------------------------------------------------------------------

  Subshares nested below import share histories on filer 192.168.25.20
  ------------------------------------------------------------
```

```
     Share : \\192.168.25.20\histories
      Path : d:\exports\histories
    Import : [[This is the imported share.]]
       ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
             37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
             0{Everyone})

     Share : \\192.168.25.20\Y2004
      Path : d:\exports\histories\2004
    Import : histories
   RelPath : 2004
       ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
             37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
             0{Everyone})

...

Subshare cache for filer 192.168.25.29
------------------------------------------------------------------------------

  Subshares nested below import share prescriptions on filer 192.168.25.29
  ------------------------------------------------------------

     Share : \\192.168.25.29\prescriptions
      Path : d:\exports\prescriptions
    Import : [[This is the imported share.]]
       ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
             37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
             0{Everyone})

     Share : \\192.168.25.29\Y2004
      Path : d:\exports\prescriptions\2004
    Import : prescriptions
   RelPath : 2004
       ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
             37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
             0{Everyone})

...

Subshare cache for filer 192.168.25.27
------------------------------------------------------------------------------

  Orphan (unnested) share caches for filer 192.168.25.27
  ------------------------------------------------------------

     Share : \\192.168.25.27\C$
      Path : C:\
    Remark : Default share
       ACL : D:P(A;;0x1f01ff;;;S-1-1-0{Everyone})

     Share : \\192.168.25.27\ADMIN$
      Path : C:\WINDOWS
    Remark : Remote Admin
       ACL : D:P(A;;0x1f01ff;;;S-1-1-0{Everyone})

     Share : \\192.168.25.27\D$
      Path : D:\
    Remark : Default share
       ACL : D:P(A;;0x1f01ff;;;S-1-1-0{Everyone})

     Share : \\192.168.25.27\BABYLON_ft_md
```

```
       Path : d:\exports\BABYLON_ft_md
        ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
              37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
              0{Everyone})

      Share : \\192.168.25.27\BAGEND_ft_md
       Path : d:\exports\BAGEND_ft_md
        ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
              37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
              0{Everyone})

...

Subshare cache for filer 192.168.25.20
------------------------------------------------------------------------------

  Orphan (unnested) share caches for filer 192.168.25.20
  ------------------------------------------------------------

      Share : \\192.168.25.20\C$
       Path : C:\
     Remark : Default share
        ACL : D:P(A;;0x1f01ff;;;S-1-1-0{Everyone})

      Share : \\192.168.25.20\ADMIN$
       Path : C:\WINDOWS
     Remark : Remote Admin
        ACL : D:P(A;;0x1f01ff;;;S-1-1-0{Everyone})

      Share : \\192.168.25.20\D$
       Path : D:\
     Remark : Default share
        ACL : D:P(A;;0x1f01ff;;;S-1-1-0{Everyone})

      Share : \\192.168.25.20\chem_results
       Path : d:\exports\chem_results
        ACL : O:S-1-5-32-544{Administrators}G:S-1-5-21-448539723-18442\
              37615-725345543-513{Domain Users}D:(A;;0x1f01ff;;;S-1-1-\
              0{Everyone})

...

      Share : \\192.168.25.29\F$
       Path : F:\
     Remark : Default share
        ACL : D:P(A;;0x1f01ff;;;S-1-1-0{Everyone})

      Share : \\192.168.25.29\G$
       Path : G:\
     Remark : Default share
        ACL : D:P(A;;0x1f01ff;;;S-1-1-0{Everyone})
```

42

ARX Manager (GUI) Maintenance

This chapter contains reference information for maintaining the ARX GUI, also known as the ARX Manager interface.

cipher

Purpose This command is designed for use with a running-config script only. As part of running-config playback, this command re-establishes ciphers used when the ARX Manager (the GUI) performs SSL negotiation.

Mode cfg-ssl

Security Role(s) crypto-officer or network-engineer

Syntax `cipher {all | defaults | cipher-name}`

all chooses all of the ciphers

defaults chooses the default (strongest) ciphers only.

cipher-name identifies one cipher to use in SSL negotiation. Type **?** for a full list of available ciphers. The running-config script invokes this command multiple times, once for each cipher.

Default(s) half of the ciphers (the strongest) are used by default

Guidelines This command is not intended for ARX administrators.

The ssl command leads to the cfg-ssl mode with this command and one other. The output of show running-config includes both of these commands so that the GUI is properly certified after re-playing the running-config script.

Related Commands
ssl
show running-config

gui

Purpose The gui restart command shuts down and restarts the ARX Manager, the Graphical User Interface (GUI) for the ARX.

You can use the no gui command to shut down the ARX GUI without immediately restarting it. The gui command starts the GUI after it was previously shut down.

Mode priv-exec

Security Role(s) crypto-officer or network-engineer

Syntax

```
gui [restart]
no gui
```

restart (optional) stops and restarts the GUI with a single command.

Default(s)

```
gui
```

Guidelines The GUI runs automatically when management access (along with its sub command, authentication) is set up for HTTP or HTTPS. You can use no gui to shut down the GUI, or gui restart to stop and restart it.

Samples

```
bstnA# no gui
```

shuts down the GUI for the current ARX.

```
bstnA# gui restart
```

stops and restarts the GUI.

Related Commands management access
authentication

ssl

Purpose This command is designed for use with a running-config script only. As part of running-config playback, this command re-establishes the SSL certificate for the ARX Manager (that is, the GUI).

Mode cfg

Security Role(s) crypto-officer or network-engineer

Syntax `ssl`

Default(s) None

Guidelines This command is not intended for ARX administrators.

The ssl command leads to a submode with commands for entering an encrypted SSL certificate and choosing the ciphers for SSL negotiation. The output of show running-config includes these commands so that the GUI is properly certified after re-playing the running-config script.

Related Commands show running-config

ssl-key-store

Purpose This command is designed for use with a running-config script only. As part of running-config playback, this command re-establishes the SSL certificate for the ARX Manager (or GUI).

Mode cfg-ssl

Security Role(s) crypto-officer or network-engineer

Syntax `ssl-key-store` ***key-data***

key-data (1-2048 characters) is one part of the encrypted SSL certificate. If the encrypted data exceeds 2048 characters, the running-config script invokes this command again with the next segment of the string.

Default(s) None

Guidelines This command is not intended for ARX administrators.

The ssl command leads to the cfg-ssl mode with this command and one other. The output of show running-config includes both of these commands so that the GUI is properly certified after re-playing the running-config script.

Related Commands ssl
show running-config

Index

A

ABE 22-8
  enabling or disabling for all back-end shares behind a volume 22-11
  excluding one share from ABE processing 22-14
activate configs 35-3
active-directory alias 23-3
active-directory forest-trust 18-5
active-directory proxy-user 23-5
active-directory update forest 18-7
active-directory update seed-domain 16-3
active-directory-forest 18-3
Adaptive Resource Switch (ARX) 1-3
additional-command 38-13
Address Resolution Protocol (ARP) 10-3, 10-4, 10-5, 10-28
age-fileset 26-3
anonymous-gid (gbl-nfs-acl) 19-3
anonymous-uid (gbl-nfs-acl) 19-4
API
  applying a snapshot schedule 34-14
  clearing API-usage statistics 34-3
  clearing notification statistics for 34-4
  creating a notification rule 34-6
  showing notification rules and snapshots 34-15
  showing usage statistics 34-21
archive 31-3
arp 10-3
arp gratuitous 10-4
ARX 1-3
at 7-3
attach 22-3
authentication 17-3
auth-port (gbl-radius) 17-6
auto reserve files 22-4
auto sync files 22-5
auto-diagnostics 38-15
auto-diagnostics test 38-17
auto-migrate 28-3

B

Backups
  for running-config 33-19, 33-43
  for volume
    preparing for 40-9
    restoring 32-6
balance 28-4
bandwidth-limit 29-3
blocked-vlan (cfg-vlan) 9-3
Boot process 5-23, 5-43
boot system 36-3
browsing 24-3

C

cancel import 22-7
cancel migrate-metadata 37-3
cancel migrate-volume 37-4
cancel migration 40-3
cancel nsck report 37-5
cancel remove 40-4
cancel restore data 32-3
cancel snapshot archive 31-4
cancel sync 37-6
capture merge 39-5
capture session 39-6
Cards
  showing configuration 5-37, 5-39, 5-43, 5-47
  showing processors 5-37, 5-39, 5-43, 5-47
cfg mode 1-6
channel 9-4
character-encoding nfs 21-3
Chassis
  clear metalog usage 5-4
child-domain 18-11
CIFS
  allowing SMB2 on the system 21-77
  allowing SMB2 with clients 24-82
  authentication
    kerberos auto-realm-traversal 18-17
    kerberos health-check threshold 18-18
  file-system type 22-71, 22-74, 22-94
  front-end CIFS service 24-5
    probe delegate-to 24-40
  service
    active-directory update forest 18-7
    active-directory update seed-domain 16-3
  setting up a DNS alias 24-19
  upgrading a CIFS service to constrained delegation 24-83
cifs 24-5
cifs access-based-enum 22-8
cifs access-based-enum (priv-exec) 22-11
cifs access-based-enum exclude 22-14
cifs anonymous-access 21-5
cifs authentication 21-6
cifs case-sensitive 22-16
cifs connection-limit 20-3
cifs deny-symlinks 22-18
cifs filer-signatures 21-8
cifs file-system-name 22-19
cifs notify-change-mode 22-21
cifs oplocks-disable 22-22
cifs path-cache 22-23
cifs promote-subshares 41-3
cifs rekey 24-8
cifs-8dot3-resolution 29-4
cifs-port 20-4
cipher only for use in running-config scripts 42-3
clear active-license 11-3

clear arp 10-5
clear at 7-5
clear counters channel 9-6
clear counters gigabit 9-7
clear counters lacp 9-8
clear counters mgmt 10-6
clear counters redundancy 15-3
clear counters redundancy network 9-10
clear counters ten-gigabit 9-9
clear dynamic-dns 24-10
clear file-history archive 31-5
clear global-config 33-3
clear health 38-18
clear metalog usage 5-4
clear nlm locks 24-11
clear nsck 37-7
clear nvr 5-5
clear restore data 32-4
clear session 17-7
clear smtp queue 13-3
clear statistics api 34-3
clear statistics authentication 17-8
clear statistics cifs authentication 41-6
clear statistics cifs path-cache 41-7
clear statistics cifs symlinks 41-8
clear statistics cifs work-queues 41-9
clear statistics domain-controller 41-10
clear statistics domain-controller load-balancing 41-11
clear statistics filer 40-5
clear statistics filer connections 39-8
clear statistics global server 39-9
clear statistics metadata 40-6
clear statistics metalog 40-7
clear statistics migration 40-8
clear statistics notification 34-4
clear statistics snapshot 30-3
clear subshare-cache 41-12
clear sync 37-8
CLI conventions 1-6
  no 1-7
clock set 5-7
clock timezone 5-8
Clock, system
  setting manually 5-7
  showing clock 5-35
  using an NTP server 10-25
close cifs file 41-13
Cluster configuration
  creating a config-replication rule 35-15
  full process 35-3
  naming the cluster 35-13
  preparing for 35-15
  removing a cluster name from the config 35-31
cluster-name 35-13
collect 38-19
collect diag-info 38-19
collect logs 38-23
Command scheduler 7-3
  clearing a scheduled CLI job 7-5
  showing all scheduled jobs 7-43
compressed-files 22-24
config-replication 35-15
constrain-directories 28-6
Constrained delegation
  finding all filers to which a CIFS service can delegate 24-40
constrain-files 28-8
contents 31-7
Conventions in documentation 1-10
Copiers
  supporting the 'Save As' feature in CIFS namespaces 21-5
copy {nfs|cifs} 7-8
copy ftp 7-6
copy global-config 33-4
copy ron 7-11
copy running-config 33-8
copy scp 7-13
copy smtp 7-16
copy startup-config 33-12
copy tftp 7-18
Core files
  NSM 15-9
CPLDs 5-23
CPU
  NSM-core restarts 15-11
  redundancy on the NSM 15-10
critical 22-25
critical route 15-5
Ctrl-z to exit a mode 1-7

D

database-location 29-5
Default(s), in the CLI Reference 1-11
delete 7-20
delete startup-config 33-16
delta-threshold 29-6
deny (gbl-nfs-acl) 19-5
description (cfg-channel) 9-11
description (cfg-email-event) 13-4
description (cfg-if-gig) 9-12
description (cfg-if-ten-gig) 9-13
description (cfg-if-vlan) 10-7
description (cfg-mgmt) 10-8
description (cfg-vlan) 9-14
description (gbl-archive) 31-9
description (gbl-cfg-repl) 35-18
description (gbl-cifs) 24-12
description (gbl-filer) 20-5
description (gbl-gs) 23-6

description (gbl-nfs) 24-13
description (gbl-nfs-acl) 19-6
description (gbl-ns) 21-10
description (gbl-ns-vol) 21-10
description (gbl-ns-vol-shr) 21-10
description (gbl-proxy-user) 16-8
description (gbl-schedule) 25-3
direct 22-26
Disaster recovery
  activate configs 35-3
  cluster-name 35-13
  config-replication 35-15
  enable (gbl-cfg-repl) 35-19
  load configs 35-20
  removing a cluster from the config 35-31
  show cluster 35-35
  show config-replication 35-36
  target-cluster 35-39
  target-file 35-40
  user (gbl-cfg-repl) 35-41
Document conventions 1-10
Domain Controller
  showing statistics for load-balancing 41-69
domain-join (gbl-cifs) 24-14
drop cifs-service user-session 41-14
drop filer-connections 39-10
dual-reboot 5-10
duration (gbl-schedule) 25-4
Dynamic DNS
  clear dynamic-dns 24-10
  dynamic-dns update 24-21
  name-server 18-19
dynamic-dns 24-19

E

Email notifications
  See SMTP.
email-event 13-5
email-severity 13-7
enable (cfg-email-event) 13-8
enable (cfg-ns) 21-11
enable (cfg-redundancy) 15-7
enable (gbl-cfg-repl) 35-19
enable (gbl-cifs, gbl-nfs) 24-22
enable (gbl-gs) 23-7
enable (gbl-gs-vs) 23-7
enable (gbl-ns-vol-ntfy) 34-5
enable (gbl-ns-vol-plc) 27-3, 28-10
enable (gbl-ns-vol-sfarm) 27-3, 28-10
enable (gbl-ns-vol-shdwcp) 29-7
enable (gbl-ns-vol-shr) 22-27
enable (gbl-ns-vol-snap) 30-4
enable slot firmware upgrade 36-6
enable, use no enable to go from priv-exec mode back to exec mode 1-7
enable/no enable CLI convention 1-7
end 1-7
Entering cfg or gbl mode 1-9
every 25-5
every (gbl-ns-vol-fs-age) 26-4
exclude 30-5
Exec mode 1-6
exit 1-7
expect monitor 38-25
expect nslookup 39-12
expect show firewall 39-13
expect show ifconfig 39-16
expect show netstat 39-19
expect traceroute 39-20
expect ttcp 39-21
export (gbl-cifs) 24-23
export (gbl-nfs) 24-27
export offline-access 24-29
export-mapping 40-9
External filers
  cifs connection-limit 20-3
  cifs-port 20-4
  external-filer 20-6
  ext-filer-ip-addrs activate 20-8
  filer-type 20-9
  ignore-name 20-12
  ip address 20-14
  ip address ... change-to 20-16
  manage snapshots 20-18
  nfs tcp connections 20-19
  probe authentication 20-20
  probe exports 20-22
  show exports 20-27
  show external-filer 20-34
external-filer 20-6
ext-filer-ip-addrs activate 20-8

F

FDB 9-62
File system type
  of a CIFS volume 22-71, 22-74, 22-94
File tracking
  canceling an archive operation 31-4
  clearing out a file-history archive 31-5
  description (gbl-archive) 31-9
  location 31-16
  querying file locations 31-25
file-history archive 31-10
filename-fileset 26-6
filer 22-31
filer-subshares 22-34
filer-type 20-9
Files
  backing up and restoring 32-6
Files, for switch maintenance

clear global-config 33-3
copy (ARX volume) 7-8
copy ftp 7-6
copy global-config 33-4
copy ron 7-11
copy running-config 33-8
copy scp 7-13
copy smtp 7-16
copy startup-config 33-12
copy tftp 7-18
delete 7-20
delete startup-config 33-16
grep 7-22
move 7-29
move {nfs|cifs} 7-33
move ftp 7-31
move scp 7-35
move tftp 7-37
rename 7-41
restore startup-config 33-17
save boot-config 33-18
show directories 7-53
show global-config 33-19
show running-config 33-43
Fileset placement
applying a schedule 27-36
automatically closing files to enable migration 27-11
choosing the target storage 27-51
disabling inline notifications 27-6
disabling volume scans 27-53
migrating all hard links 27-12
Filesets
age-fileset 26-3
filename-fileset 26-6
filesize-fileset 26-7
from (gbl-ns-vol-plc) 27-4
from fileset 26-8
inline notify 27-6
intersection-fileset 26-9
last 26-10
limit-migrate 27-10
name 26-11
offline attribute, CIFS 26-13
path 26-14
place-rule 27-14
policy-age-fileset 26-17
policy-cifs-attributes-fileset 26-18
policy-filename-fileset 26-20
policy-filesize-fileset 26-21
policy-intersection-fileset 26-22
policy-union-fileset 26-23
recurse 26-24
select-files (gbl-ns-vol-fs-age) 26-26
select-files (gbl-ns-vol-fs-filesize) 26-25
show all global filesets 26-28
source share 27-50
source share-farm 27-50
start 26-32
union-fileset 26-33
volume-scan 27-53
filesize-fileset 26-7
find 31-12, 40-12
firmware upgrade 36-6
flowcontrol 9-16
forest-root 18-13
forward-delay 9-15
FPGAs 5-24
freespace adjust 22-41
freespace apparent-size 22-43
freespace calculation dir-master-only 22-45
freespace calculation manual 22-47
freespace cifs-quota 22-49
freespace ignore 22-52, 22-62, 40-72
from (cfg-smtp) 13-9
from (gbl-ns-vol-plc) 27-4
from fileset (gbl-ns-vol-fs-age) 26-8
from fileset (gbl-ns-vol-fs-isect) 26-8
from fileset (gbl-ns-vol-fs-union) 26-8
from fileset (gbl-ns-vol-shdwcp) 29-8
Front-end services 1-3
active-directory forest-trust 18-5
active-directory-forest 18-3
anonymous-gid (gbl-nfs-acl) 19-3
anonymous-uid (gbl-nfs-acl) 19-4
child-domain 18-11
cifs 24-5
cifs rekey 24-8
clear dynamic-dns 24-10
clear nlm locks 24-11
close cifs file 41-13
deny (gbl-nfs-acl) 19-5
description (gbl-cifs) 24-12
description (gbl-nfs-acl) 19-6
domain-join (gbl-cifs) 24-14
drop cifs-service user-session 41-14
drop filer-connections 39-10
dynamic-dns 24-19
dynamic-dns update 24-21
enable (gbl-cifs, gbl-nfs) 24-22
export (gbl-cifs) 24-23
export (gbl-nfs) 24-27
export offline-access 24-29
forest-root 18-13
name-server 18-19
nfs 24-31
nfs-access-list 19-8
nlm enable 24-34
offline-behavior 24-36
performance
metalog 15-37, 40-41
show active-directory 16-11
show active-directory status 16-15

show cifs-service 24-45
show cifs-service client-activity 41-15
show cifs-service exports 41-21
show cifs-service filer-capabilities 24-54
show cifs-service kerberos-tickets 41-24
show cifs-service open-files 41-27
show cifs-service path-cache 41-30
show cifs-service subshares 24-58
show cifs-service transactions 41-33
show cifs-service user-sessions 41-34
show dynamic-dns 24-62
show global service 24-64
show nfs tcp 24-72
show nfs-service 24-65
show nfs-service mounts 24-70
show nlm client 24-73
show nlm file 24-75
show nlm statistics 24-77
show statistics cifs authentication 41-38, 41-42
show statistics cifs fastpath 41-47
show statistics cifs path-cache 41-49
show statistics cifs symlinks 41-53
show statistics cifs work-queues 41-57
show statistics domain-controller 41-65
show statistics domain-controller load-balancing 41-69
show virtual service 24-79
wins 23-26
wins-alias 23-27
wins-name 23-29
wins-name-encoding 24-84
FTP
copy ftp 7-6
move ftp 7-31

G

gbl mode 1-6
Global commands, accessible from any mode 1-6
Global Server
active-directory alias 23-3
active-directory proxy-user 23-5
Global server
enable (gbl-gs) 23-7
enable (gbl-gs-vs) 23-7
global server 23-8
show global server 23-10
virtual server 23-21
windows-domain (gbl-gs) 23-25
global server 23-8
Gratuitous ARP 10-4
grep 7-22
group 4-3
group ... event 13-12
group (cfg-email-event) 13-10
group (gbl-user) 4-5
Groups, administrative
adding 4-3
adding user 4-18
setting role (security level) 4-8
showing 4-11
showing group roles 4-12
showing group users 4-13
windows-domain 4-19
gui 42-4
GUI maintenance 42-1
Guidelines, in the CLI Reference 1-11

H

HA
See Redundancy.
Hard links
migrating 27-12
heartbeat failure 14-3
heartbeat interval 14-4
hello-time 9-17
hostname 5-11

I

ignore-name 20-12
ignore-sid-errors 22-53
import priority 22-54
import rename-directories 22-56
import rename-files 22-59
import skip-managed-check 22-61
import sync-attributes 22-62
inline notify 27-6
inline report 27-7
inline-notify 29-9
interface gigabit 9-18
interface mgmt 10-9
interface ron 14-5
interface ten-gigabit 9-19
interface vlan 10-11
intersection-fileset 26-9
ip address 20-14
ip address ... change-to 20-16
ip address (cfg-if-vlan) 10-12
ip address (cfg-mgmt) 10-14
ip address (cfg-ron) 14-6
ip address (gbl-nis-dom) 19-7
ip address (gbl-ntlm-auth-srv) 18-15
ip domain-list 10-15
ip ftp-user 7-26
ip name-server 10-16
ip private subnet reassign 14-7
ip private vlan 9-20
ip proxy-address 10-17
ip ron-user 7-27
ip route 10-19
ip route ... per-vlan 10-21

ip route ... source-ip 10-23
ip scp-user 7-28

J
jumbo mtu 9-22

K
Keeper CPLD 5-23
Kerberos
for namespace clients 21-6
kerberos auto-realm-traversal 18-17
kerberos health-check threshold 18-18
key (gbl-radius) 17-9

L
lacp active 9-23
lacp passive 9-25
lacp rate 9-27
last 26-10
Layer 2
channel 9-4
clear counters channel 9-6
clear counters gigabit 9-7
clear counters lacp 9-8
clear counters redundancy network 9-10
clear counters ten-gigabit 9-9
description (cfg-channel) 9-11
description (cfg-if-gig) 9-12
description (cfg-if-ten-gig) 9-13
description (cfg-vlan) 9-14
flowcontrol 9-16
forward-delay 9-15
gigabit statistics 9-54
hello-time 9-17
interface gigabit 9-18
interface ten-gigabit 9-19
interface vlan 10-11
mac-address aging-time 9-29
management source 38-30
max-age 9-32
members (cfg-channel) 9-33
members (cfg-vlan) 9-35
priority (cfg-channel) 9-37
priority (cfg-stp) 9-38
protocol (cfg-stp) 9-39
redundancy (cfg-if-vlan) 10-26
redundancy protocol 9-40
redundancy protocol (cfg-channel) 9-42
show channel 9-44, 9-51
show load-balancing 9-59
show mac-address-table 9-60
show mac-address-table summary 9-62
show spanning-tree detailed 9-65
show spanning-tree interface 9-66
show spanning-tree summary 9-67
show vlan 9-68
show vlan summary 9-70
shutdown (cfg-channel) 9-71
shutdown (cfg-if-gig) 9-72
shutdown (cfg-if-ten-gig) 9-73
shutdown (cfg-stp) 9-74
spanning-tree 9-75
spanning-tree cost 9-77
spanning-tree edgeport 9-78
spanning-tree priority 9-79
spanning-tree shutdown 9-80
speed (cfg-if-gig) 9-81
switch-forwarding enable 9-82
tag 9-83
trap shutdown 9-85
vlan 9-86
vlan (cfg-channel) 9-88
vlan-tag 9-89
Layer 3
arp 10-3
arp gratuitous 10-4
clear arp 10-5
clear counters mgmt 10-6
description (cfg-if-vlan) 10-7
description (cfg-mgmt) 10-8
interface mgmt 10-9
ip address (cfg-if-vlan) 10-12
ip address (cfg-mgmt) 10-14
ip domain-list 10-15
ip ftp-user 7-26
ip name-server 10-16
ip private vlan 9-20
ip proxy-address 10-17
ip ron-user 7-27
ip route 10-19
ip route ... per-vlan 10-21
ip route ... source-ip 10-23
ip scp-user 7-28
show arp 10-28
show interface mgmt 10-42
show interface vlan 10-44
show ip address 10-45
show ip domain 10-46
show ip proxy-addresses 10-47
show ip route 10-49
shutdown (cfg-if-vlan) 10-58
speed 10-61
wait-for ip-routes 10-62
wait-for vip-disable 23-23
wait-for vip-enable 23-24
LBA FPGA 5-24
License
automatically activating 11-4
clearing the current license 11-3
creating a dossier file (for manual activation) 11-8

manually activating 11-6
showing the active license 11-11
license activate 11-4
license activate file 11-6
license create license-dossier 11-8
licensing
pinging the license server 11-10
limit-migrate 27-10
load configs 35-20
load-balance 9-29
location 31-16
logging destination 38-26
logging fastpath component 39-23
logging fastpath processor 39-27
logging fastpath processor ... filter 39-25
Logging in 1-9
logging level 38-27
Logging, syslog 38-1
login-banner 5-12
ls, for switch-maintenance files 5-19

M

mac-address aging-time 9-29
Machine name, setting 5-11, 5-12
mail-server 13-14
mail-to (cfg-email-event) 13-15
mail-to (gbl-auto-diag) 38-29
manage snapshots 20-18
managed-volume 22-64
management access 17-10
Management interfaces
in-band (VLAN)
adding 10-11
listing 10-44
out-of-band (OOB) MGMT
configuring 10-9
showing the configuration 10-42
management source 38-30
Master directory
establishing mastership during import with import priority 22-54
max-age 9-32
maximum age 13-16
max-volume-groups 22-65
members (cfg-channel) 9-33
members (cfg-vlan) 9-35
metadata cache-size 21-14
metadata critical 22-66
metadata share 22-67
Metalog (MTL) FGPA 5-24
migrate close-file 27-11
migrate hard-links 27-12
migrate retain-files 27-13
Migration
showing migration statistics 40-45
waiting for 27-54
MMC
adding a WMA group to a namespace 21-78
enabling MMC browsing 24-3
windows-mgmt-auth 18-43
Mode
in the CLI Reference 1-11
Modes
cfg
entering 1-9
config 1-6
exec 1-6
exiting 1-7
gbl
entering 1-9
global commands 1-6
priv-exec 1-6
prompts 1-7
modify 22-69
Modules
showing configuration 5-37, 5-39, 5-43, 5-47
showing processors 5-37, 5-39, 5-43, 5-47
slot locations 5-3
monitor 39-29
move 7-29
move {nfs|cifs} 7-33
move ftp 7-31
move scp 7-35
move tftp 7-37
moving-average algorithm for stats-monitor 38-31
Multi-protocol namespace
strict-attribute-consistency 40-72
MUX CPLD 5-23

N

name 26-11
named-streams 22-71
name-server 18-19
Namespace 1-3
adding a shadow volume 29-19
attach 22-3
character-encoding nfs 21-3
cifs access-based-enum 22-8
cifs access-based-enum (priv-exec) 22-11
cifs access-based-enum exclude 22-14
cifs anonymous-access 21-5
cifs case-sensitive 22-16
cifs deny-symlinks 22-18
cifs notify-change-mode 22-21
cifs oplocks-disable 22-22
cifs path-cache 22-23
clear restore data 32-4
clear sync 37-8
compressed-files 22-24
critical share 22-25

direct 22-26
enable (cfg-ns) 21-11
enable (gbl-ns-vol-shr) 22-27
filer 22-31
filer-subshares 22-34
ignore-sid-errors 22-53
import priority 22-54
import rename-directories 22-56
import rename-files 22-59
import skip-managed-check 22-61
import sync-attributes 22-62
managed-volume 22-64
max-volume-groups 22-65
metadata cache-size 21-14
metadata critical 22-66
metadata share 22-67
migrate retain-files 27-13
modify 22-69
named-streams 22-71
namespace 21-16
ntlm-auth-db (gbl-ns) 21-18
ntlm-auth-server (gbl-ns) 21-19
persistent-acls 22-74
policy migrate-attempts 27-22
policy migrate-delay 27-23
policy migrate-retry-delay 27-26
policy treewalk-threads 27-31
protocol 21-21
proxy-user (gbl-ns) 21-22
reimport-modify 22-76
remove-share migrate 40-17
remove-share nomigrate 40-19
remove-share offline 40-21
removing 21-23
removing full namespace service 40-15
replica-snap 22-77
reserve files 22-79
restarting a volume 22-80
sam-reference 21-25
share (gbl-ns-vol) 22-81
share (gbl-ns-vol-sfarm) 28-13
share-farm 28-14
show namespace 21-27
show namespace status 21-71
show nsck 37-36
show policy history for a volume 40-23
show policy settings for a namespace 27-37
show restore data 32-11
show share status 22-84
show sid-translation 22-86
show sync 37-40
show volume-group 22-87
sid-translation 22-92
sparse-files 22-93
sync-attributes 22-64
unicode-on-disk 22-94
volume 22-96
volume group 22-98
wait-for nsck 37-59
wait-for remove 40-74
wait-for report 7-81
wait-for restore data 32-14
wait-for shares-online 22-100
wait-for sync 37-60
wait-for volume-disable 22-101
wait-for volume-enable 22-102
namespace 21-16
Namespace-check utility
See nsck.
Namespaces
adding new CIFS subshares to a volume 37-55
backing up and restoring files 32-6
finding new CIFS subshares behind a volume 37-51
performance
metalog 15-37, 40-41
repairing metadata inconsistencies 37-33, 37-34, 37-43, 37-46
Netlogon
showing statistics for 41-65
NFS
clear nlm locks 24-11
nfs tcp timeout 24-33
nlm enable 24-34
offline-behavior 24-36
show nfs tcp 24-72
show nlm client 24-73
show nlm file 24-75
show nlm statistics 24-77
nfs 24-31
nfs tcp connections 20-19
nfs tcp timeout 24-33
nfs-access-list 19-8
nfs-param 22-73
NIS
ip address for the NIS server 19-7
nis domain 19-10, 19-12
nis domain (gbl-nfs-acl) 19-11
NLM
clearing locks 24-11
disabling at an NFS service 24-34
showing clients 24-73
showing locked files 24-75
showing statistics 24-77
nlm enable 24-34
No Convention, to undo a CLI command 1-7
notification rule 34-6
Notifications
applying a snapshot schedule 34-14
clearing API-usage statistics 34-3
clearing statistics for 34-4
creating a notification rule 34-6
showing, and showing snapshots for 34-15

notify in case of ARX/device communication issue 38-33
nsck
allow modify on rebuild 22-76
clear nsck 37-7
show nsck 22-3, 37-36
wait-for nsck 37-59
nsck ... destage 37-10
nsck ... migrate-metadata 37-12
nsck ... migrate-volume 37-14
nsck ... rebuild 37-16
nsck ... report dir-structure 37-18
nsck ... report inconsistencies 37-22
nsck ... report metadata-only 37-27
nsck ... report symlinks 37-30
nsck ... sync directories 37-33
nsck ... sync files 37-34
NSM
fastpath log
logging fastpath component 39-23
logging fastpath processor 39-27
logging fastpath processor ... filter 39-25
show fastpath logging 39-36
fastpath resources
show fastpath resources 39-38
fastpath statistics 40-49
nsm binary-core-files 15-9
nsm recovery 15-10
nsm warm-restart 15-11
NTLM
for namespace clients 21-6
ip address 18-15
ntlm-auth-db 18-20
password (gbl-ntlm-auth-srv) 18-24
port (gbl-ntlm-auth-srv) 18-27
Secure Agent 18-22
show ntlm-auth-db 18-28
show ntlm-auth-server 18-29
user (gbl-ntlm-auth-db) 18-40
windows-domain (gbl-ntlm-auth-srv) 18-41
ntlm-auth-db 18-20
ntlm-auth-db (gbl-ns) 21-18
ntlm-auth-server 18-22
ntlm-auth-server (gbl-ns) 21-19
NTLMv2
for namespace clients 21-6
ntp server 10-25
NVRAM
clearing 5-5
NVR FPGA 5-24

O

Offline access, for CIFS clients 24-29
offline attribute, CIFS 26-13
Offline storage removal
remove-share offline 40-21
offline-behavior 24-36
Online storage removal
remove namespace 21-23
remove service 40-15
remove-share migrate 40-17
remove-share nomigrate 40-19

P

password 4-6
password (gbl-ntlm-auth-srv) 18-24
password (gbl-user) 4-7
path 26-14
pause 7-39
peer 15-13
peer address 14-8
Performance
metalog latency affects volumes 15-37, 40-41
permit 17-12
permit (gbl-mgmt-auth) 18-25
permit (gbl-nfs-acl) 19-14
persistent-acls 22-74
Photocopiers
supporting the 'Save As' feature in CIFS namespaces 21-5
ping 39-31
ping license-server 11-10
place-rule 27-14
Policy
adding a schedule 25-7
auto-migrate 28-3
balance 28-4
config-replication rule 35-15
constrain files 28-8
constrain-directories 28-6
enable (gbl-ns-vol-plc) 27-3, 28-10
enable (gbl-ns-vol-sfarm) 27-3, 28-10
policy pause 27-29
policy pause (gbl-ns-vol) 27-30
removing from a namespace 27-32
shadow-copy-rule 29-21
show policy 27-37
show policy files-closed 27-48
show policy filesets 26-28
show policy history 40-23
show policy queue 40-25
show schedule 25-8
showing migration statistics 40-45
target 27-51
waiting for a migration to finish 27-54
policy freespace 27-18
policy freespace (gbl-ns) 27-20
policy freespace (gbl-ns-vol) 27-20
policy freespace (gbl-ns-vol-sfarm) 28-11
policy migrate-attempts 27-22
policy migrate-delay 27-23

policy migrate-method 27-24
policy migrate-retry-delay 27-26
policy order-rule 27-27
policy pause 27-29
policy pause (gbl-ns-vol) 27-30
policy treewalk-threads 27-31
policy-age-fileset 26-17
policy-cifs-attributes-fileset 26-18
policy-filename-fileset 26-20
policy-filesize-fileset 26-21
policy-intersection-fileset 26-22
policy-union-fileset 26-23
port (gbl-ntlm-auth-srv) 18-27
Port mirroring/monitoring 39-29
Port speed, setting 9-81
Power CPLD 5-23
priority (cfg-channel) 9-37
priority (cfg-stp) 9-38
Priv-exec mode 1-6
probe authentication 20-20
probe delegate-to 24-40
probe exports 20-22
probe metalog latency 5-13
Processors
  showing 5-37, 5-39, 5-43, 5-47
Prompts, show position in mode hierarchy 1-7
protocol 21-21
protocol (cfg-stp) 9-39
Proxy user
  proxy-user 16-9
  show proxy-user 16-20
  user 16-22
  windows-domain 16-24
proxy-user 16-9
proxy-user (gbl-filer) 20-25
proxy-user (gbl-ns) 21-22
prune-target 29-10
publish 29-11
Purpose, in the CLI Reference 1-11

Q

quorum-disk 15-14

R

RADIUS
  auth-port 17-6
  key (gbl-radius) 17-9
  radius-server 17-13
  retries (gbl-radius) 17-14
  show radius-server 17-16
  timeout (gbl-radius) 17-26
radius-server 17-13
raid offline 6-3
raid rebuild 6-4
raid rebuild-rate 6-6
raid silence 6-7
raid verification-mode 6-8
raid verification-rate 6-10
raid verify 6-12
RAID, commands for managing 6-1, 11-1
rconsole 14-9
Reboot 5-10, 5-14, 15-9, 15-10, 15-11, 15-22, 15-24
recurse 26-24
Redundancy
  among NSM cores 15-10
  critical route 15-5
  critical shares 22-25
  enable (cfg-redundancy) 15-7
  metadata critical 22-66
  NSM core restarts 15-11
  peer 15-13
  preparing a channel as a redundant-pair link 9-42
  preparing a port as one end of a redundant-pair link 9-40
  quorum-disk 15-14
  redundancy 15-16
  redundancy force-active 15-19
  resilver-timeout 15-21
  show redundancy 15-27, 15-41
  show redundancy all 15-28
  show redundancy ballots 15-31
  show redundancy critical-services 15-33
  show redundancy history 15-35
  show redundancy license 15-36
  show redundancy metalog 15-37
  show redundancy peer 15-40
  show redundancy reboot-history 15-44
  show redundancy resilver-timeout 15-45
  suspend failover 15-46
redundancy 15-16
redundancy force-active 15-19
redundancy protocol 9-40
redundancy protocol (cfg-channel) 9-42
reimport-modify 22-76
Related commands, in the CLI Reference 1-11
reload 5-14
remark 7-40
remove cluster-config 35-31
remove namespace 21-23
remove namespace ... policy-only 27-32
remove namespace ... volume ... exports-only 24-44
remove service 40-15
remove-share migrate 40-17
remove-share nomigrate 40-19
remove-share offline 40-21
rename 7-41
replica-snap 22-77
report (gbl-ns-vol-ntfy) 34-11
report (gbl-ns-vol-plc) 27-34
report (gbl-ns-vol-shdwcp) 29-12, 35-32
report (gbl-ns-vol-snap) 30-7

reserve files 22-79
Reset CPLD 5-23
resilver-timeout 15-21
Resource proxy 1-3
resource-profile 5-16
restart namespace ... volume 22-80
restore data 32-6
restore startup-config 33-17
retain count for notification (API-supporting) snapshots 34-13
retain count for snapshots 30-11
retries (gbl-radius) 17-14
retry attempts 29-16
retry delay 29-17
retry interval 13-17
role 4-8, 4-9
RON
    heartbeat failure 14-3
    heartbeat interval 14-4
    interface ron 14-5
    ip address 14-6
    ip private subnet reassign 14-7
    ip ron-user 7-27
    peer address 14-8
    ron evict 14-10
    ron tunnel 14-11
    show ron 14-12
    show ron conflicts 14-14
    show ron database 14-15
    show ron route 14-18
    show ron tunnel 14-19
    shutdown (cfg-if-vlan-ron-tnl) 14-23
ron tunnel 14-11

S

Samples, in the CLI Reference 1-11
sampling interval for stats-monitor 38-35
sam-reference 21-25
save boot-config 33-18
save profile 3-3
Scanners
    supporting the 'Save As' feature in CIFS namespaces 21-5
Schedule
    duration 25-4
    every 25-5
    every (gbl-ns-vol-fs-age) 26-4
    start 25-10
    stop 25-11
schedule 25-7
schedule (gbl-auto-diag) 38-36
schedule (gbl-ns-vol-ntfy) 34-14
schedule (gbl-ns-vol-plc) 27-36
schedule (gbl-ns-vol-shdwcp) 29-18, 35-34
schedule (gbl-ns-vol-snap) 30-12
Schedules
    showing 25-8
SCP
    copy ron 7-11
    copy scp 7-13
    move scp 7-35
Security level
    setting group role 4-8, 4-9
Security Roles
    in the CLI Reference 1-11
select-files (gbl-ns-vol-fs-age) 26-26
select-files (gbl-ns-vol-fs-filesize) 26-25
Selective Authentication for Kerberos 18-17
Shadow volume
    bandwidth-limit 29-3
    cifs-8dot3-resolution 29-4
    database-location 29-5
    delta-threshold 29-6
    enable (gbl-ns-vol-shdwcp) 29-7
    from fileset 29-8
    inline-notify 29-9
    prune-target 29-10
    publish 29-11
    shadow 29-19
    shadow-copy-rule 29-21
    show shadow 29-24
    sid-translation 29-29
    target (gbl-ns-vol-shdwcp) 29-31
shadow-copy-rule 29-21
share (gbl-ns-vol) 22-81
share (gbl-ns-vol-sfarm) 28-13
Share import error conditions 21-49
Share import status conditions 21-35, 21-36
Share statistics 40-49
share-farm 28-14
show active-directory 16-11
show active-directory status 16-15
show active-license 11-11
show arp 10-28
show at 7-43
show auto-diagnostics 38-37
show baudrate 5-18
show boot 36-8
show capture 7-44
show capture sessions 39-34
show channel 9-44, 9-51
show chassis 5-19
show chassis software 36-9
show cifs-service 24-45
show cifs-service client-activity 41-15
show cifs-service exports 41-21
show cifs-service filer-capabilities 24-54
show cifs-service kerberos-tickets 41-24
show cifs-service open-files 41-27
show cifs-service path-cache 41-30
show cifs-service subshares 24-58

show cifs-service transactions 41-33
show cifs-service user-sessions 41-34
show clock 5-35
show cluster 35-35
show config-replication 35-36
show configs 7-48
show cores 7-50
show diag-info 7-52
show directories 7-53
show documentation 38-39
show dynamic-dns 24-62
show email-event 13-18
show email-severity 13-21
show exports 20-27
show external-filer 20-34
show fastpath logging 39-36
show fastpath resources 39-38
show file-history archive 31-18
show file-history archive ... contents 31-21
show file-history virtual-server
See show file-history virtual-service
show file-history virtual-service 31-25
show filer connections 39-43
show firmware upgrade 36-10
show global server 23-10
show global service 24-64
show global-config 33-19
show group all 4-11
show group roles 4-12
show group users 4-13
show health 38-40
show health time-skew 38-42
show history 3-4
show hostname 5-36
show id-mappings 38-44
show interface 10-34
show interface gigabit 9-54
show interface mgmt 10-42
show interface ten-gigabit 9-57
show interface vlan 10-44
show ip address 10-45
show ip domain 10-46
show ip proxy-addresses 10-47
show ip route 10-49
show license 7-57
show license-dossier 11-14
show load-balancing 9-59
show logging destination 38-46
show logging levels 38-47
show logs 7-59
show mac-address-table 9-60
show mac-address-table summary 9-62
show master-key 8-3
show memory usage 5-37
show metalog usage 5-39
show monitor 39-45
show namespace 21-27
show namespace mapping 21-68
show namespace status 21-71
show nfs tcp 24-72
show nfs-access-list 19-17
show nfs-service 24-65
show nfs-service mounts 24-70
show nis domain 19-20
show nis netgroup 19-22
show nlm client 24-73
show nlm file 24-75
show nlm statistics 24-77
show notification 34-15
show nsck 22-3, 37-36
show nsm 15-22
show nsm warm-restart history 15-24
show ntlm-auth-db 18-28
show ntlm-auth-server 18-29
show ntp servers 10-55
show ntp status 10-56
show policy 27-37
show policy files-closed 27-48
show policy filesets 26-28
show policy history 40-23
show policy queue 40-25
show processors 5-43
show processors usage 5-47
show proxy-user 16-20
show radius-server 17-16
show redundancy 15-27, 15-41
show redundancy all 15-28
show redundancy ballots 15-31
show redundancy critical-services 15-33
show redundancy history 15-35
show redundancy license 15-36
show redundancy metalog 15-37
show redundancy network 9-63
show redundancy peer 15-40
show redundancy reboot-history 15-44
show redundancy resilver-timeout 15-45
show releases 7-60
show replicated-configs 7-62
show reports 7-64
show reports status 7-70
show restore data 32-11
show ron 14-12
show ron conflicts 14-14
show ron database 14-15
show ron route 14-18
show ron tunnel 14-19
show running-config 33-43
show schedule 25-8
show scripts 7-71
show server-mapping 23-13
show shadow 29-24
show share status 22-84

show sid-translation 22-86
show smtp queue 13-26
show smtp status 13-28
show smtp welcome 13-31
show snapshot 30-14
show snmp-server 12-3
show software 7-73
show spanning-tree detailed 9-65
show spanning-tree interface 9-66
show spanning-tree summary 9-67
show ssh-host-key 17-19
show ssh-user 4-14
show statistics api 34-21
show statistics authentication 17-21
show statistics cifs authentication 41-38, 41-42
show statistics cifs fastpath 41-47
show statistics cifs path-cache 41-49
show statistics cifs symlinks 41-53
show statistics cifs work-queues 41-57
show statistics domain-controller 41-65
show statistics domain-controller load-balancing 41-69
show statistics filer 40-27
show statistics filer connections 39-46
show statistics global server 39-51
show statistics metalog 40-37, 40-41
show statistics migration 40-45
show statistics namespace ... fastpath 40-49
show statistics namespace ... request-detail 40-52
show statistics namespace ... response-detail 40-61
show statistics namespace ... summary 40-67
show stats-logs 7-75
show stats-monitor 38-49
show subshare-cache 41-72
show sync 37-40
show system tasks 38-51
show terminal 3-5
show uptime 5-52
show users 4-16
show version 5-54, 5-55, 40-38, 40-42, 40-44
show virtual path-history 31-31
show virtual service 24-79
show vlan 9-68
show vlan summary 9-70
show volume-group 22-87
show windows-mgmt-auth 16-22, 18-34
showing results from diagnostic tests 38-51
shutdown 5-57
shutdown (cfg-channel) 9-71
shutdown (cfg-if-gig) 9-72
shutdown (cfg-if-ten-gig) 9-73
shutdown (cfg-if-vlan) 10-58
shutdown (cfg-if-vlan-ron-tnl) 14-23
shutdown (cfg-stp) 9-74
sid-translation 22-92
sid-translation (gbl-ns-vol-shdwcp) 29-29
signatures 24-81
Slots
  locations 5-3
SMB signing
  from ARX to clients 24-81
  from ARX to filers 21-8
SMB2
  allowing for a CIFS service 24-82
  allowing, system-wide 21-77
smb2-allowed 21-77
smb2-allowed (gbl-cifs) 24-82
SMTP
  clearing the SMTP queue 13-3
  configuring email events 13-5
  copy smtp 7-16
  retrying all emails in the SMTP queue 13-34
  sending a introductory email to the users of an email event role 13-39
  sending a test email 13-37
  sending a test email to an email-event role 13-35
  showing email events 13-18
  showing email severities 13-21
  showing SMTP status 13-28
  showing the SMTP queue 13-26
  showing the template for SMTP-welcome emails 13-31
  testing the connection to the email server 13-38
smtp 13-32
smtp retry 13-34
smtp test email-event 13-35
smtp test message 13-37
smtp test server 13-38
smtp welcome 13-39
snapshot clear 30-22
snapshot consistency 30-24
snapshot create 30-25
snapshot directory cifs-name 30-27, 30-31
snapshot directory display 30-29
snapshot manage 30-32
snapshot privileged-access 30-34
snapshot remove 30-36
snapshot replica-snap-rule 30-39
snapshot rule 30-46
snapshot verify 30-53
snapshot vss-mode 30-56
Snapshots
  allowing or preventing the ARX from managing a filer's snapshots 20-18
  applying a schedule 30-12
  clearing statistics for 30-3
  creating a replica-snapshot rule 30-39
  creating a snapshot rule 30-46
  identifying the filer type 20-9
  manually invoking a snapshot rule 30-25
    waiting for 30-58
  proxy-user 20-25

removing from filers 30-36
waiting for 30-59
removing from the ARX configuration 30-22
showing 30-14
storing snapshots in a replica-snap share 22-77
verifying integrity 30-53
waiting for 30-60
SNMP
allowing sets and gets from a remote host 12-11
clearing an active alarm 38-18
enabling traps 12-10
enabling traps for a channel 9-85
identifying a target host for traps 12-7
permitting access to the SNMP agent 17-12
showing active alarms 38-40
showing SNMP configuration 12-3
snmp-server community 12-5
snmp-server contact 12-6
snmp-server host 12-7
snmp-server location 12-8
snmp-server name 12-9
snmp-server traps 12-10
snmp-server trusthost 12-11
Software release
showing 5-54, 5-55, 36-9, 40-38, 40-42, 40-44
source share 27-50
source share-farm 27-50
spanning-tree 9-75
spanning-tree cost 9-77
spanning-tree edgeport 9-78
spanning-tree priority 9-79
spanning-tree shutdown 9-80
sparse-files 22-93
speed 10-61
speed (cfg-if-gig) 9-81
speed (cfg-mgmt) 10-61
spn 20-47
SSB FPGA 5-24
SSH
enabling SSHv1 17-25
manually setting the switch's host keys 17-22
regenerating the switch's host keys 17-24
showing the switch's host keys 17-19
using to run a restore data from a remote host 32-8
ssh-host-key 17-22
ssh-host-key generate 17-24
ssh-key 4-9
ssh-v1 enable 17-25
ssl only for use in running-config scripts 42-5
ssl-key-store only for use in running-config scripts 42-6
start 25-10, 26-32
Statistics
raw packet counts per share 40-49
request counts, namespace to filer 40-52
request/response counts and RTTs 40-67
RTT from namespace to filer 40-61
Layer 2
show interface gigabit 9-54
show interface ten-gigabit 9-57
stats-monitor for ARX/device communication issues 38-53
Status MUX CPLD 5-23
stop 25-11
strict-attribute-consistency 40-72
suspend failover 15-46
Switch upgrades
boot system 36-3
show boot 36-8
show chassis software 36-9
switch-forwarding enable 9-82
sync cifs delegation 24-83
sync directories 37-43
sync files 37-46
sync subshares from-namespace 37-51
sync subshares from-service 37-55
sync-attributes 22-64
Syntax, in the CLI Reference 1-11
Syslog 38-1
logging destination 38-26
logging level 38-27
remark 7-40
show logging destination 38-46
showing logging levels and categories 38-47, 39-36, 39-38

T

tag 9-83
tail 7-79
target (gbl-ns-vol-plc) 27-51
target (gbl-ns-vol-shdwcp) 29-31
target-cluster 35-39
target-file 35-40
tentative 27-52
terminal beta 3-7
terminal character-set 3-8
terminal clear 3-9
terminal confirmation 3-10
terminal expand-prompt 3-11
terminal expert 3-12
terminal history 3-13
terminal length 3-14
terminal logging 3-15
terminal stop-on-error 3-16
terminal timeout 3-17
terminal width 3-18
TFTP
copy tftp 7-18
move tftp 7-37
Tiered storage
import priority 22-54

timeout (gbl-radius) 17-26
to (cfg-smtp) 13-41
Traceroute 39-20
trap enabled for stats-monitor 38-54
trap shutdown 9-85
Traps
  monitor for ARX/device communication issues 38-53
  shutting down for a channel 9-85
  trigger for ARX/device communication issues 38-33
tree-domain (gbl-forest) 18-36
Troubleshooting
  cifs promote-subshares 41-3
  clear health 38-18
  clear statistics cifs authentication 41-6
  clear statistics cifs path-cache 41-7
  clear statistics cifs symlinks 41-8
  clear statistics cifs work-queues 41-9
  clear statistics domain-controller 41-10
  clear statistics domain-controller load-balancing 41-11
  clear statistics filer 40-5
  clear statistics filer connections 39-8
  clear statistics global server 39-9
  clear statistics metadata 40-6
  clear statistics metalog 40-7
  clear statistics migration 40-8
  clear subshare-cache 41-12
  close cifs file 41-13
  drop cifs-service user-session 41-14
  drop filer-connections 39-10
  monitor 39-29
  moving-average algorithm for stats-monitor 38-31
  notify in case of ARX/device communication issue 38-33
  ping 39-31
  ping license-server 11-10
  sampling interval for stats-monitor 38-35
  show auto-diagnostics 38-37
  show cifs-service client-activity 41-15
  show cifs-service exports 41-21
  show cifs-service kerberos-tickets 41-24
  show cifs-service open-files 41-27
  show cifs-service path-cache 41-30
  show cifs-service user-sessions 41-34
  show documentation 38-39
  show filer connections 39-43
  show id-mappings 38-44
  show monitor 39-45
  show statistics cifs authentication 41-38, 41-42
  show statistics cifs fastpath 41-47
  show statistics cifs path-cache 41-49
  show statistics cifs symlinks 41-53
  show statistics cifs work-queues 41-57
  show statistics domain-controller 41-65
  show statistics domain-controller load-balancing 41-69
  show statistics filer 40-27
  show statistics filer connections 39-46
  show statistics global server 39-51
  show statistics migration 40-45
  show stats-monitor 38-49
  show subshare-cache 41-72
  stats-monitor for ARX/device communication issues 38-53
  trap enabled for stats-monitor 38-54
truncate-report 7-80
TTCP 39-21

U

unicode-on-disk 22-94
union-fileset 26-33
Upgrading software 36-1
user 4-17
user (gbl-cfg-repl) 35-41
user (gbl-group) 4-18
user (gbl-mgmt-auth) 18-38
user (gbl-ntlm-auth-db) 18-40
user (gbl-proxy-user) 16-22
Users, administrative
  adding 4-5, 4-6, 4-7, 4-17
  adding a group role 4-8
  adding an SSH key for 4-9
  adding to a group 4-18
  showing group roles 4-12
  showing group users 4-13
  showing SSH users 4-14
  showing users 4-16

V

Valid Platforms, in the CLI Reference 1-11
Version, showing 5-54, 5-55, 40-38, 40-42, 40-44
VIP 23-21
virtual server 23-21
vlan 9-86
vlan (cfg-channel) 9-88
VLANs
  adding 9-86
  adding tagged ports 9-83
  assigning to a channel 9-88
vlan-tag 9-89
Volume
  cifs file-system-name 22-19
  nfs-param 22-73
  performance
    metalog 15-37, 40-41
  policy migrate-method 27-24
  restarting 22-80
  shadowing with another volume 29-19
volume 22-96

volume group 22-98
volume-scan 27-53

W

wait-for ip-routes 10-62
wait-for migration 27-54
wait-for nsck 37-59
wait-for remove 40-74
wait-for report 7-81
wait-for restore data 32-14
wait-for shares-online 22-100
wait-for snapshot create 30-58
wait-for snapshot remove 30-59
wait-for snapshot verify 30-60
wait-for sync 37-60
wait-for vip-disable 23-23
wait-for vip-enable 23-24
wait-for volume-disable 22-101
wait-for volume-enable 22-102
Windows
    unsupported features 24-25
windows-domain (gbl-group) 4-19
windows-domain (gbl-gs) 23-25
windows-domain (gbl-ntlm-auth-srv) 18-41
windows-domain (gbl-proxy-user) 16-24
windows-mgmt-auth 18-43
windows-mgmt-auth (gbl-ns) 21-78
wins 23-26
wins-alias 23-27
wins-name 23-29
wins-name-encoding 24-84