ARX® Site Planning Guide
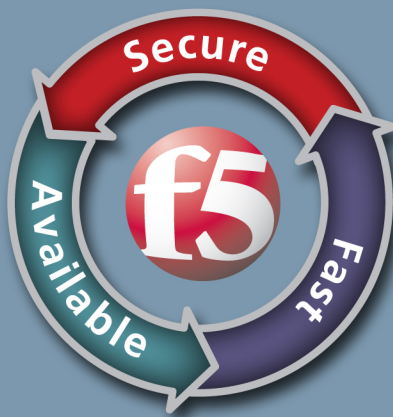
Publication Date

This manual was published on September 24, 2012.

Legal Notices

Copyright

Trademarks

Patents

Export Regulation Notice

RF Interference Warning

FCC Compliance

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

November 2005 - Rev B - hardware correction

March 2006 - Rev C - Software Release 2.3

August 2006 - Rev D, updates for Software Release 2.4

September 2006 - Rev E, added more configuration limits

September 2006 - Rev F, updates for Software Release 2.4.2

March 2007 - Rev G, added best practices for more filer types; updates for Release 2.5

May 2007 - Rev H, updates for Release 2.5.1

December 2007 - Rev J, clarification of Metadata-share best practices

February 2008 - Rev K, clarification of Console-cable pinouts for Release 2.7.1

March 2008 - Rev L, conversion to F5 format for Release 3.1.0

June 2008 - Rev M, more namespace/volume limits for Release 3.2.0

June 2008 - Rev N, updates for the ARX®4000 chassis, and for Release 4.0.0

October 2008 - Rev P, re-brand the OS and revise power specifications for ARX4000

June 2009 - Rev Q, update filer/server instructions

November 2009 - Rev R, add ARX®2000 and other updates for Release 5.01.000

January 2010 - Rev S, more updates for Release 5.01.000

March 2010 - Rev T, updates for Software Release 5.01.005
September 2010 - Rev U, add updates for Release 5.02.000
January 2011 - Rev V, add updates for Release 5.03.000
June 2011 - Rev W, add updates for Release 6.00.000
September 2011 - Rev X, add updates for Release 6.01.000
October 2011 – Rev Y, refer to licensed limits
July 2012 - Rev Z, add updates for Release 6.02.000

1

Site Planning

This manual describes network and environmental considerations for installing an Adaptive Resource Switch (ARX®). Use this document to prepare for adding an ARX to your network.

Concepts and Terminology

The ARX acts as a *resource proxy* between the current clients and servers on your network. The switch terminates client requests, determines the correct server to process the request, and then originates a new request to the server. Messages in the reverse direction, from servers to clients, also terminate and restart at the ARX. The clients are said to be at the *front end* of the ARX, and the servers are said to be at the *back end*. As you plan to add a switch to your network, it is helpful to remember the sharp division between the switch's front-end and back-end processing.

The following figure illustrates this sharp division.

Figure 1.1 ARX architecture showing front end versus backend
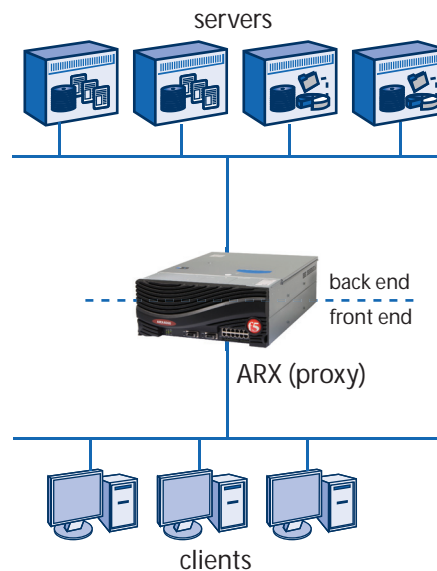
Platforms and Modules

You can purchase any of the following ARX platforms:

- Single-port ARX-VE (a virtual appliance)
- Single-port ARX-500
- 8-port ARX-1500
- 12-port ARX-2000
- 4-port ARX-2500, with 2 additional 10-Gigabit ports
- 12-port ARX-4000, with 2 additional 10-Gigabit ports

This document is relevant to all of the above platforms.

Namespaces

You can configure one or more *namespaces* for your front-end clients. Each namespace is a collection of virtual file systems, called *volumes*, under a single authentication domain. A volume is a collection of shares (or exports) hosted on the back-end file servers.

Selecting a Network Topology

You can deploy the ARX in one of the following network topologies:

- *One-armed proxy*. The ARX uses a single logical connection to reach all clients and servers.
- *Multiple subnet*. The ARX terminates one or more client subnets and a separate server subnet.

The following subsections describe each of these topologies.

One-Armed Proxy Topology

In a one-armed proxy topology, the ARX connects to a single IP subnet on a single VLAN. All of the switch's connectivity to clients and servers goes through this single subnet/VLAN. The connection is typically through a link-aggregation channel.

One-Armed Proxy: Before Installing an ARX

Figure 1.2 shows clients and servers on the same VLAN and subnet before the introduction of the ARX. The router connects the LAN to additional client and server subnets, perhaps on other campuses.

Figure 1.2 *Clients and servers on the same VLAN before using the* ARX.
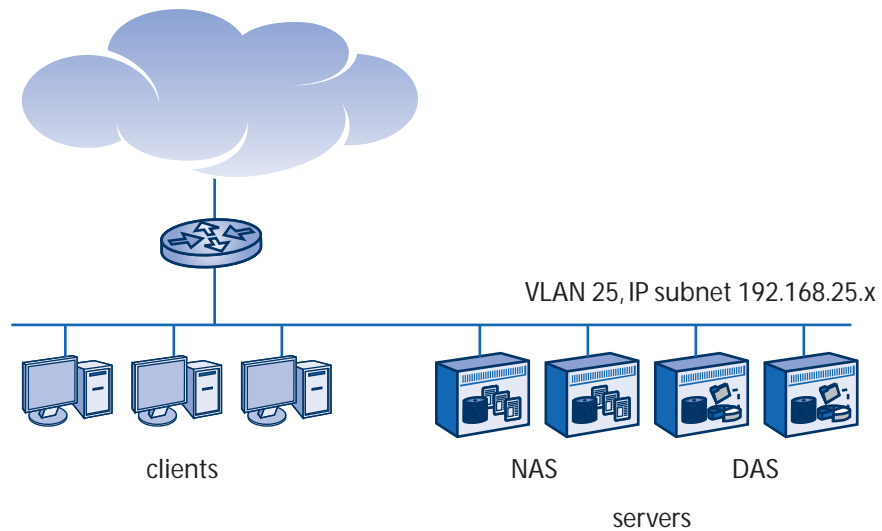
One-Armed Proxy: After Installing an ARX

The ARX has a single physical connection to the client/server subnet. On the switch, you configure the same subnet and VLAN for both front-end clients and back-end servers. You can also add static routes to additional client and/or server subnets.

Figure 1.3 shows clients and servers after cutting in an ARX.

Figure 1.3 *Clients and servers on a VLAN after cutting in an* ARX.

VLAN 25, IP subnet 192.168.25.x

clients   NAS   DAS
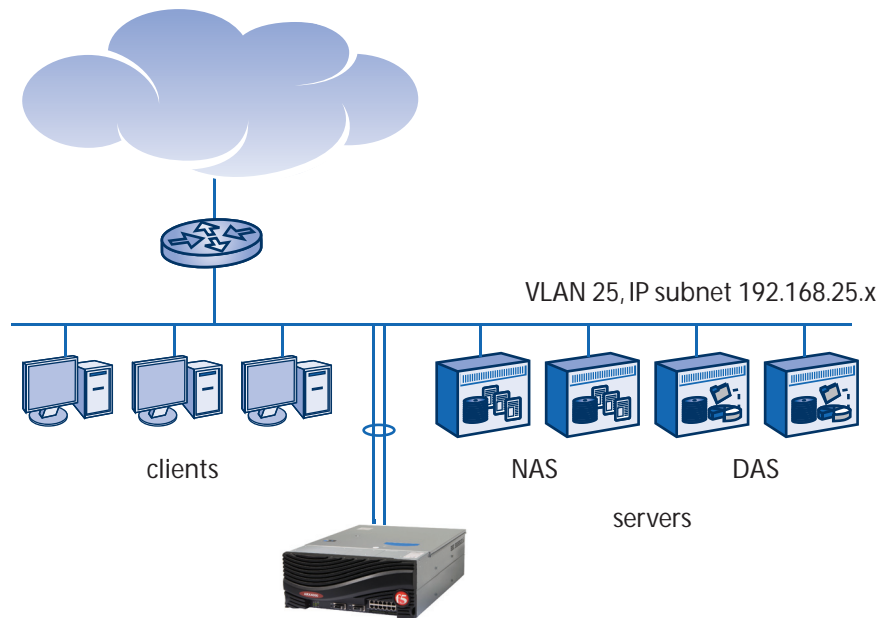
servers

Multiple Subnet Topology

A multiple subnet deployment divides clients and servers into multiple IP subnets. You can define multiple subnets, static routes, and default routes on the ARX to reach any number of subnets at the front end or back end.

Multiple Subnet: Before Installing an ARX

Figure 1.4 shows clients and servers on separate VLANs and subnets, with a router connecting the two subnets.

Figure 1.4 *Clients and servers on separate VLANs and subnets*

servers

192.168.25.x

172.100.90.x
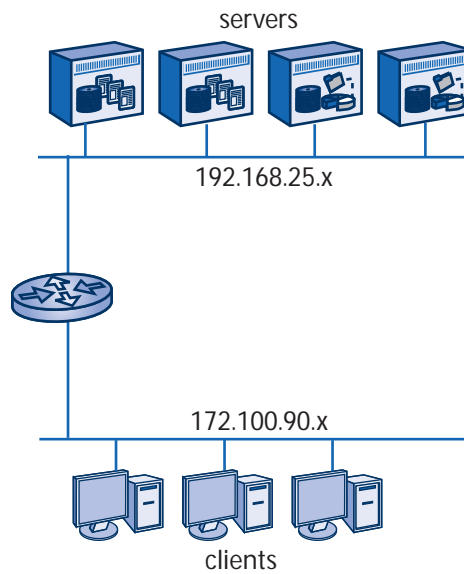
clients

Multiple Subnet: After Installing an ARX

As shown in *Figure 1.5*, the ARX has a separate connection to the client subnet and the server subnet in a multiple subnet topology. The switch serves as a proxy for CIFS and/or NFS transactions between the clients and servers.

As a proxy, the ARX has separate transactions with clients on its front end and servers on its back end. In this topology, the ARX does *not* cause looping.

Figure 1.5 *ARX as proxy for CIFS and/or NFS transactions*
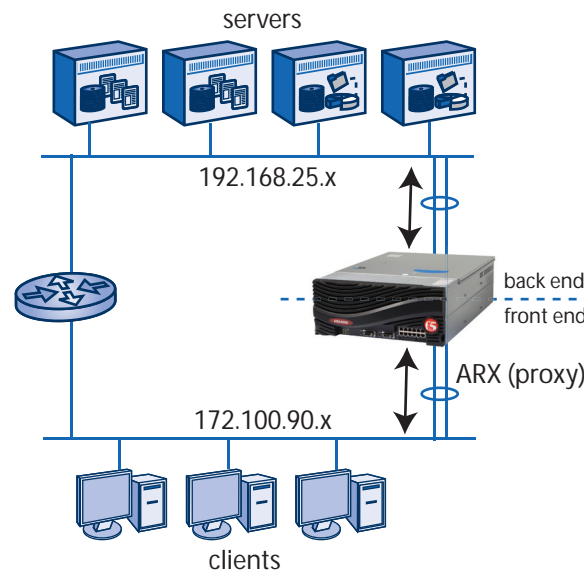
Allocating IP Addresses

As a resource proxy with distributed processors, the ARX requires several IP addresses to communicate with front-end clients and its back-end servers. Every network processor on the switch requires its own address to terminate and originate transactions with back-end servers. These IP addresses are called *proxy IP* addresses.

Communication to client subnets is less distributed: a minimum of one IP address is required per namespace to handle client requests. The client-side IP addresses are called *virtual IPs*, or *VIPs*.

The ARX also requires one or more management IP addresses to load new software images. The switch has one out-of-band management interface, and you can configure one in-band interface for each VLAN. The in-band interfaces can also be used to connect one ARX to another in a Resilient Overlay Network (RON).

◆ Note

At least one in-band management interface is required for any redundant switch pair.

The Server Subnet and Proxy IP Addresses

The proxy IP addresses must reside on a single server subnet. If you have multiple server subnets in your current network, you can configure static routes to the remote subnets (through a router on a local server subnet). The

chosen server subnet must have enough address space for one proxy IP address per network processor. The number of network processors varies for each platform type. The following table shows the number of proxy IP addresses for each platform type.

Table 1.1 Required number of proxy IP addresses and platform types

Number of Proxy IPs Required	Platform Type
12	ARX-4000
3	ARX-2500
4	ARX-2000
2	ARX-1500
1	ARX-500 or ARX-VE

In any chassis type, the CLI issues a warning if insufficient proxy IP addresses have been defined.

Adding Proxy IP Addresses to Domain Name Servers

During an NFS mount, some back-end servers perform domain name server (DNS) reverse-lookups against remote NFS clients. This is designed as an authentication mechanism. To prevent such file servers from denying access to the proxy IP addresses, add the proxy IP addresses to your DNS configuration.

Adding Virtual IP Addresses

Clients access ARX storage through a unique virtual IP address (VIP) or VIPs. All VIPs can reside on a single client subnet or they can be divided among several client subnets. Any client must be able to reach its respective VIP. The ARX can send responses back to clients through a client subnet, a static route, or a default route.

Configuring Management IP Addresses

You can configure in-band management interfaces, one per configured VLAN. At least one such interface is required for many installations – and adding at least one in-band management interface is strongly recommended in any case.

◆ A switch in a Resilient Overlay Network (RON) of switches requires at least one in-band management address. A RON is a network of IP tunnels that connects multiple ARXes together. Each tunnel requires an in-band IP address as an endpoint. You can re-use the same in-band management IP for multiple RON tunnels, but at least one is required.

◆ Redundant switches also require at least one in-band management interface to communicate with shared, external resources (the quorum disk).

The ARX also has an interface for out-of-band management (typically labeled *MGMT*) on its front panel. This interface is designed for installations with discrete management networks. It must have an IP address outside of the server subnet or any of the client subnets.

◆ Note

ARX-VE does not offer out-of-band management.

Ports Required by the ARX

The ARX is assumed to operate in a three-tiered network in a large data center. The first tier is core routers that provide connectivity to a campus or WAN. The second tier has redundant distribution switches that distribute all data center traffic between the access switches; the access switches constitute the third tier of the network. All LAN clients and back-end file servers connect to the access switches. Through the access switches and distribution switches, LAN clients connect to all back-end servers/storage, to one another, and to the WAN. *Figure 1.6* shows a sample network.

Figure 1.6 *Sample Network*
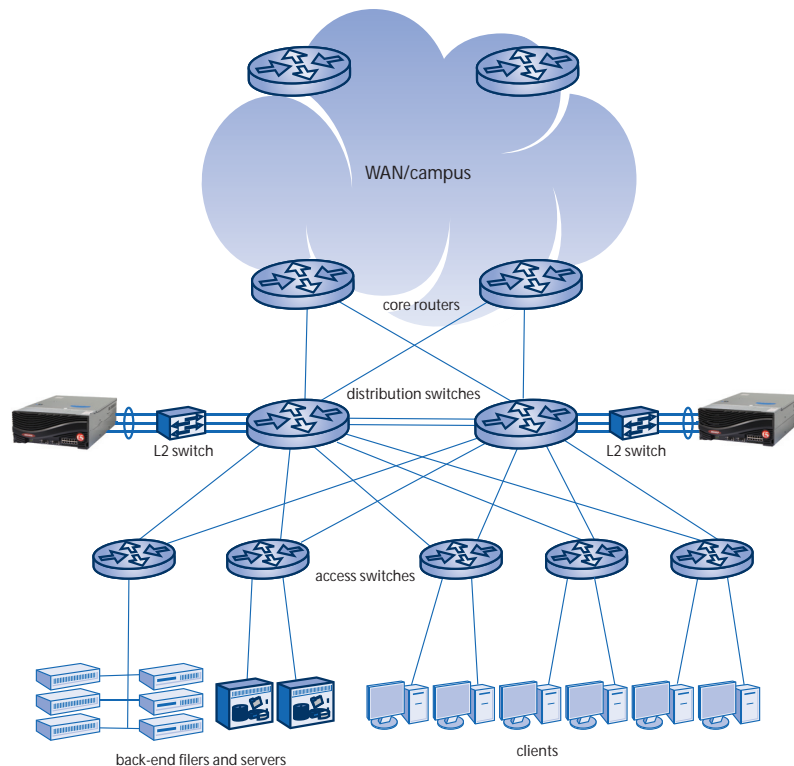
The sample network has redundant ARX devices connected at each of the distribution switches. Physically, this is a one-armed connection; conceptually, the ARX has clients in front and file servers in back.

The network file servers are the storage behind the (front-end) services of the ARX. The servers can be heterogeneous: NAS devices and file servers need only support CIFS or NFS.

Table 1.2 lists the ports required by the ARX to communicate with its authentication services, back-end storage, and front-end client services. If you plan to operate in a secure environment, these ports must be opened on fire walls in order for the ARX to function properly.

All client/server data is on the inband network. The inband network enables clients and servers to talk to the ARX. The higher layer protocols use the inband network. If traps and email home are configured, they also use the inband network.

The out-of-band network is used to connect to the CLI or the GUI for purposes of querying or configuring the ARX.

There is a third network (heartbeat) between the two ARX devices in a high-availability configuration that is used to pass redundancy information between the switches.

Table 1.2 *Ports Required by the ARX*

Port	Service/Protocol	Inbound VIP	XIP	MGMT	Outbound VIP	XIP	MGMT	Comment
ARX Management								
161	SNMP agent for polling UDP			✓				Disabled by default.
162	SNMP traps TCP/UDP			✓			✓	Disabled by default. The port is configurable.
22	SSH TCP			✓			✓	Enabled by default.
23	TELNET TCP			✓				Disabled by default.
443	HTTPS TCP			✓			✓	GUI enabled by default.
80	HTTP TCP			✓			✓	GUI disabled by default.
514	SYSLOG UDP							If External Syslog is used.
25	SMTP TCP						✓	For Email Home.
139	MMC TCP							
49803	RON UDP			✓			✓	Inband only.
20/21	FTP TCP						✓	Passive mode client only, no server.
	SCP						✓	Client, no server.
123	NTP TCP/UDP						✓	Disabled by default.
1812	RADIUS						✓	Disabled by default.
NFS Proxy								
111	rpcbind TCP/UDP	✓				✓		
2049	server V2 UDP/TCP	✓				✓		
2049	server V3 UDP	✓				✓		
2049	locked TCP/UDP	✓				✓		
635	mountd TCP/UDP	✓				✓		

Table 1.2 Ports Required by the ARX (Continued)

Port	Service/Protocol	Inbound VIP	I XIP	I MGMT	Outbound VIP	I XIP	I MGMT	Comment
637	nlockmgr TCPUDP	✓				✓		
638	status TCP/UDP	✓				✓		
CIFS Proxy/SMB								
445	CIFS (SMB) Server TCP	✓				✓		Preferred port.
139	CIFS (SMB) Server TCP CIFS (SMB) over NETBIOS	✓				✓		
CIFS Authentication/Other								
53	DNS TCP/UDP				✓	✓		Queries.
389	LDAP TCP/UDP				✓	✓		
25805	NTLM agent default TCP				✓	✓		Default, port is configurable.
464	lc passwd				✓	✓		UDP and TCP
137/138	WINS UDP	✓			✓	✓		Name service
88	Kerberos TCP/UDP				✓	✓		
137, 138	NetBIOS TCP/UDP				✓	✓		
445	Microsoft Directory Services TCP				✓	✓		
NFS Authentication (NIS)								
	NIS UDP/TCP				✓	✓		Client, no server. Outbound only.
	DNS UDP/TCP				✓	✓		Client, no server. Outbound only.
Snapshot Management to File Servers								
514	RSH (remote shell) TCP						✓	
22	SSH TCP						✓	Enabled by default.

Port	Service/Protocol	Inbound VIP	XIP	MGMT	Outbound VIP	XIP	MGMT	Comment
598, 5986	WINRM (Windows Remote Management) TCP						✓	
High-Availability (HA)								
49800	"rendezvous"							
API (new in 5.2.0)								
83	HTTP-API			✓				
843	HTTPS-API			✓				

Using NTP

To support time-based policies, the clock on the ARX must be consistent with the clock in its back-end servers. For example, an accurate clock is needed to determine when to trigger age-based file migration or replication. Accurate time is also important because the ARX keeps track of time-based file attributes, such as last-modified time and create time.

In a redundant pair of ARX devices, time consistency between the redundant peers is vital. A redundant peer may trigger a failover if an improper clock setting makes a heartbeat appear to be out-of-date.

For these reasons and others, we strongly recommend using two or more NTP servers for the servers, clients, and the ARX devices.

Automatically Discovering Your File Servers

Prior to installing an ARX, best practices recommend using F5 Data Manager, a web-based application that can give you a detailed understanding of your unique file storage configuration, contents, structure, and usage. With this understanding, you can identify and apply data management policies to create an efficient and cost-effective storage environment.

Data Manager can give you a snapshot of your current environment as it exists today, starting with file server discovery.

Data Manager performs file server discovery by examining a specified file server and collecting detailed configuration parameters and then displaying and reporting on that information through its GUI. Discovery is accomplished through a wizard that guides you during the process.

In addition, Data Manager uses discovered configuration information to assist in bringing an ARX inline and in virtualizing your NAS file server environments.

Data Manager installs and runs on any Microsoft® Windows® platform that meets a set of minimum system and network requirements. For specifics on these requirements, consult the product *Release Notes*. To download your free trial version of Data Manager, click:

http://www.f5.com/products/data-manager/

Manually Preparing the Back-End File Servers

To assist the process of a subsequent ARX import, Data Manager discovery can identify issues with file servers and their shares/exports. The following subsections describe these issues. You can use these guidelines to manually discover any issues for a site without Data Manager.

Avoiding Name Collisions

To prepare multiple file server shares for inclusion in the same namespace volume, you should avoid *name collisions*. A name collision occurs when two shares contain a file with the same path and name. The collision is resolved by renaming the second file (and all subsequent files with that path and name) before they are imported.

For example, consider two servers with /var directories that you want to combine into a single namespace volume. If server A contains /var/log/readme.txt and you import it first into a namespace, the import succeeds. If server B contains the same filename at the same path (/var/log/readme.txt), the same volume cannot import it because two different files cannot exist in the same volume path. Either the import fails, or (optionally) you permit the import process to rename the file. The renamed file includes the name of the namespace share and an ID for the import job; for example, /var/log/readme_myShare-4.txt.

To prevent the rename on import, you can change its name or location on one of the filers before the import. F5 provides tools and services to assist with finding name collisions and other issues before the ARX is installed.

NFS Servers

When a namespace imports an NFS export/share, the ARX takes inventory by reading the share's directory tree as *root*. The shares cannot squash root access by the ARX devices' proxy IPs, or this tree walk (and therefore the import) may fail. Set your NFS shares to *no-root-squash* for all of your proxy IPs.

NFS access control is based on the IP addresses of NFS clients; some IP addresses are allowed and some are not. All of the proxy IP addresses require exclusive root access to the server's NFS export.

For a list of all proxy IP addresses on the ARX, issue a show ip proxy-addresses command. Note the addresses that are in use. See the following example.

```
bstnA> show ip proxy-addresses

  Proxy Address      VLAN      Mac Address      Owner      In Use By      Proc
------------------   ----   -----------------   -------   -----------     ----
  192.168.25.31/24    25    00:0a:49:17:7c:80   bstnA        bstnA        2.1
  192.168.25.32/24    25    00:0a:49:17:7c:81   bstnA        bstnA        2.2
  192.168.25.33/24    25    00:0a:49:17:7c:82   bstnA        bstnA        2.3
  192.168.25.34/24    25    00:0a:49:17:7c:83   bstnA        bstnA        2.4
  192.168.25.141/24   25    00:0a:49:17:7c:84   bstnA        bstnA        2.5
  192.168.25.142/24   25    00:0a:49:17:7c:85   bstnA        bstnA        2.6
  192.168.25.143/24   25    00:0a:49:17:7c:86   bstnA        bstnA        2.7
  192.168.25.144/24   25    00:0a:49:17:7c:87   bstnA        bstnA        2.8
  192.168.25.145/24   25    00:0a:49:17:7c:88   bstnA        bstnA        2.9
  192.168.25.146/24   25    00:0a:49:17:7c:89   bstnA        bstnA        2.10
  192.168.25.147/24   25    00:0a:49:17:7c:8a   bstnA        bstnA        2.11
  192.168.25.148/24   25    00:0a:49:17:7c:8b   bstnA        bstnA        2.12
```

Do *not* allow access to these shares from actual clients; changes from other clients would cause confusion for the namespace software. The only exception to this rule would be a management client, which may require access for backups or troubleshooting.

NetApp Shares

A NetApp share requires read/write access and root access for every proxy IP address. You should exclude non-proxy IP addresses, other than management IP addresses.

To import a share, navigate to **NFS > Manage Exports** (in the left-hand navigation panel) and double-click the share that you want the ARX to import. *Figure 1.7* shows the NetApp Manage Exports screen.

Figure 1.7 *NetApp Manage Exports screen*
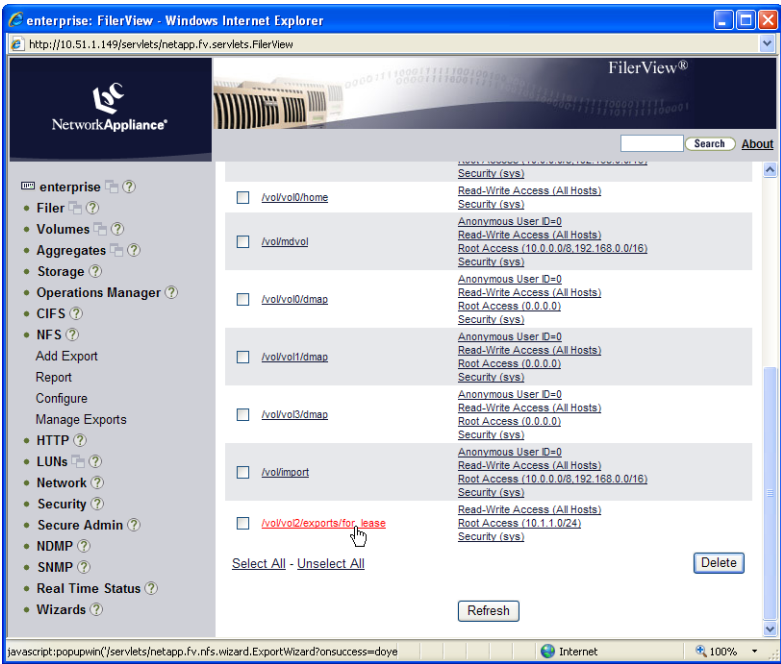
The NFS Export wizard is invoked. See the following figure. In the wizard, select the options to configure for the export.

Figure 1.8 *NFS Export wizard*
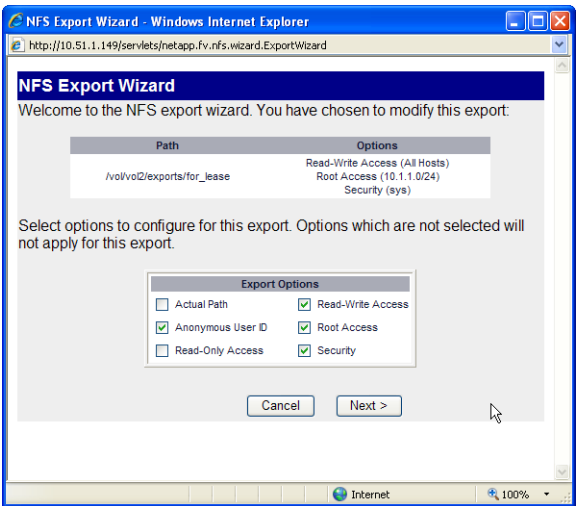
Use this wizard to add your proxy IP addresses to the following lists:

- read/write access

- root access

These lists should be limited to proxy IP addresses and management addresses only.

On the wizard Commit screen (the final screen), click **Commit**. See the following figure.

Figure 1.9 *NFS Export wizard Commit screen*
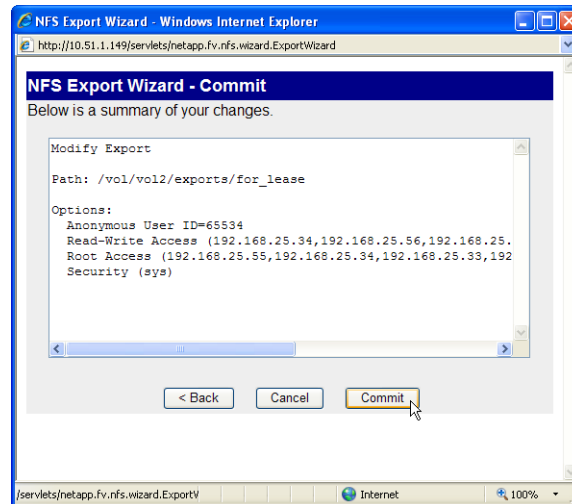
EMC Celerra Server

On the EMC Celerra server, select **NFS Exports**. See the left-hand navigation column in the following figure.

Click **New** or click an existing export name, as appropriate.

Figure 1.10 *EMC Celerra screen for an existing export*
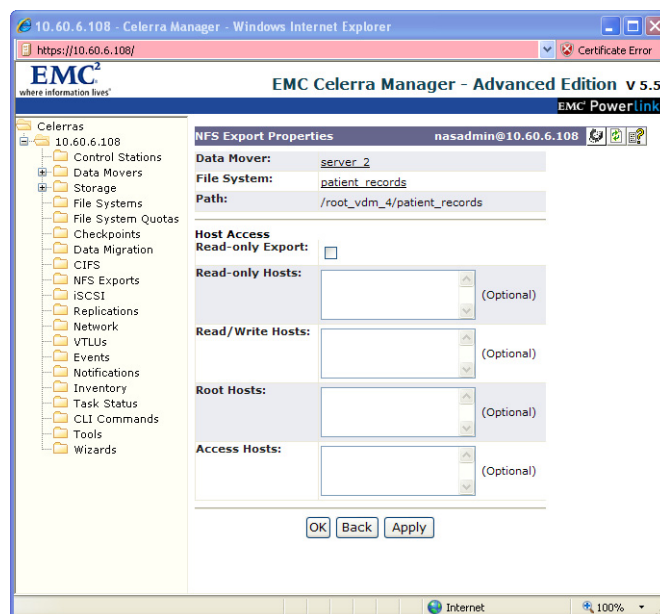
In the Root Hosts field, enter all your proxy IP addresses. Enter only proxy IP addresses (and, possibly, management IP addresses) in these fields. For an example showing these fields, see the following figure.

Figure 1.11 *EMC Celerra Read/Write Hosts or the Root Hosts fields.*
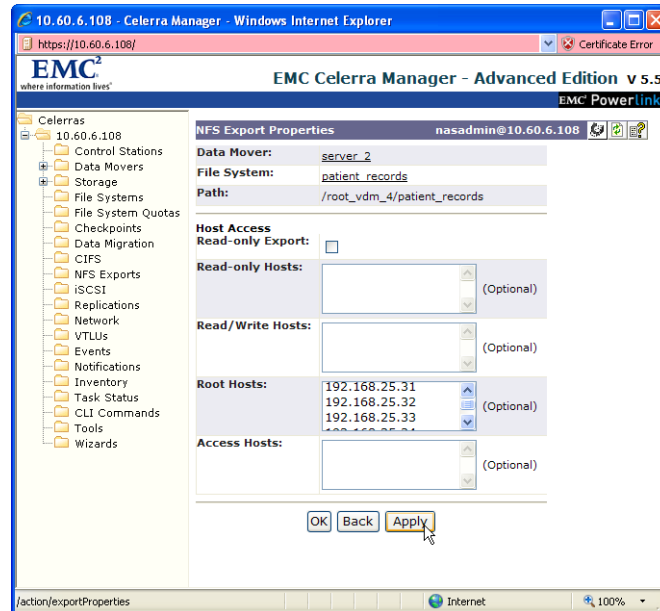
If there are any other users/groups with write access to the backend, you should remove them (unless they involve backup or well-supervised administration).

EMC Data Domain Exports

All Data Domain exports are in the /backup directory. You can use the Data Domain CLI to export one of those directories through NFS. Use the **nfs add** command to create an NFS export, specifying the ARX proxy-IP addresses as the only clients that can access the export.

For example, this Data Domain command exports the /backup/t3users directory to the proxy-IP addresses shown in previous examples:

```
sysadmin@med-dd510# nfs add /backup/t3users 192.168.25.31
192.168.25.32 192.168.25.33 192.168.25.34 192.168.25.141
192.168.25.142 192.168.25.143 192.168.25.144 192.168.25.145
192.168.25.146 192.168.25.147 192.168.25.148
```

By default, the export allows read/write access and does not have a root squash setting.

Linux

The following instructions have been tested with Debian-Linux NFS servers. The same method should apply to other Linux distributions, such as RedHat.

As root, edit the /etc/exports file to set the following for each NFS export:

- read-write (rw) access for all proxy-IP addresses
- no root squashing (no-root-squash) for all proxy-IP addresses

For example, the following file exports /lhome/budget to an ARX with 12 proxy IP addresses:

```
# /etc/exports: the access control list for filesystems which may be
exported
#         to NFS clients.  See exports(5).
/lhome/budget        192.168.25.31(rw,no_root_squash)
/lhome/budget        192.168.25.32(rw,no_root_squash)
/lhome/budget        192.168.25.33(rw,no_root_squash)
/lhome/budget        192.168.25.34(rw,no_root_squash)
/lhome/budget        192.168.25.141(rw,no_root_squash)
/lhome/budget        192.168.25.142(rw,no_root_squash)
/lhome/budget        192.168.25.143(rw,no_root_squash)
/lhome/budget        192.168.25.144(rw,no_root_squash)
/lhome/budget        192.168.25.145(rw,no_root_squash)
/lhome/budget        192.168.25.146(rw,no_root_squash)
/lhome/budget        192.168.25.147(rw,no_root_squash)
/lhome/budget        192.168.25.148(rw,no_root_squash)
```

This limits access to the proxy IPs only. Only the ARX should be able to access this export, except for backups or other read-only activities.

To use the edited exports file, use the following commands at the Linux file server (still as root):

```
exportfs -a
```

```
/etc/init.d/nfs-kernel-server restart
```

To lean more about the configuration options in the /etc/exports file, use the following command:

```
man 5 exports
```

BSD

The configuration of NFS on a BSD system has the following goals:

- Allow access from all proxy-IP addresses.
- Limit access from other IP addresses to read-only.
- The namespace software gets confused if the directory structure changes unexpectedly on one of its back-end shares.
- Allow the proxy-IP addresses to access the export/share as root.
- Allow mounts in subdirectories.

  The namespace software makes NFS mounts below the root of the share.

As root, edit the /etc/exports file to accomplish all of these goals. To allow mounts below the root of the share, you must use the -alldirs flag. For security reasons, BSD only allows this flag for shares that map to block devices. On the BSD machine, use the df command for a list of block devices.

Consider the following example:

```
BSD2-[/root]# df
Filesystem  1K-blocks     Used    Avail Capacity  Mounted on
/dev/da0s1a    495726    73274   382794    16%    /
devfs               1        1        0   100%    /dev
/dev/da0s1e   4058062    10390  3723028     0%    /users
/dev/da0s1f  29356354 10575988 16431858    39%    /usr
/dev/da0s1d    297326    11040   262500     4%    /var
BSD2-[/root]#
```

On the above machine, you can export /, /dev, /users, /usr, and/or /var. Consider the following example of a /etc/exports file that exports /usr to the ARX:

```
# Note that in order to allow the Acopia Switch to do
# "submounting", you must set the export with a -alldirs
# flag. This means you must also set the export dir to be a
# top level filesystem found on a /dev/<slice>.
/usr -alldirs -maproot=0 -network 192.168.25.0 -mask 255.255.255.0
```

This limits access to IPs only in the 192.168.25.0/24 subnet, the server subnet for the ARX. Only the ARX should be able to access this export, except for backups or other read-only activities.

To use the edited exports file, force the mountd process to re-read the file. As root, find the PID for mountd and then send a kill -HUP to the PID.

For example:

```
BSD2-[/root]# ps ax | grep mountd
  347  ??  Is     0:03.21 /usr/sbin/mountd -r
24817  p5  S+     0:00.00 grep mountd
BSD2-[/root]# kill -HUP 347
BSD2-[/root]#
```

CIFS Servers and Client Authentication

Whereas NFS uses a client's IP address for access control, CIFS uses the client's user account. Each CIFS share, directory, and file has an associated Access Control List (ACL), a list of user accounts and the access privileges of each. The ARX passes a CIFS client's credentials through to each back-end server, and the server applies its ACLs to the client's identity. The server has the authority to deny any of the client's actions based on its ACL configuration, just as though the client was accessing the server directly. No special preparation is required for CIFS client authentication.

Autonomous ARX operations, such as migrating files between back-end CIFS shares, require a Windows user identity so that the ARX can similarly access servers. This identity, called a *proxy user*, is a valid user account in

the file servers' Windows domain. The proxy user requires strong privileges on all CIFS-supporting servers, This user account must belong to the Backup Operators group or a group with equivalent privileges, and it must have full control (defined as both read and change control) over all files and directories in the share.

To support CIFS subshares and their share-level ACLs, the proxy user credentials must belong to the more-privileged Administrator's group. These higher privileges are required for RPC queries, such as queries about the physical paths of the filer's shares and subshares.

You may also need to address potential issues with licensing, local groups, and mapping a proxy user to the UNIX root user. These issues are discussed in the following sections.

Windows Server Licensing

Each namespace can use up to 16 threads for replicating and migrating files between back-end CIFS shares. Since a thread can move files between shares on the same Windows server, using one connection to the source share and a second connection to the destination share, this means that each namespace can use up to 32 concurrent server connections.

If your installation supports *Per-Seat* licensing, this is not an issue. For *Per-Server* licensing, you must configure each back-end server with 32 licenses per namespace.

Support for Local Groups

If you plan to use a CIFS server with Local Groups in an ARX volume, you must duplicate the local groups on each of the volume's servers. This only applies to a volume where you plan to use policy to migrate files. If a client has access to a file on server A because he is a member of the "doctors" local group, he cannot access the file if it moves to server C, where the "doctors" group is undefined. If the volume has policy running, the policy engine is likely to migrate many files between server A and server C.

You can also configure a volume to copy its files to a remote ARX volume, called a *shadow volume*. In this case, you must copy all local groups behind the source volume to all servers behind the shadow volume. This facilitates client access to the local copy as well as the remote one. If the shadow volume site is in a different Windows domain altogether, you must duplicate all of the Windows user groups in the source volume file servers in all of the shadow volume servers.

EMC Data Domain

The EMC Data Domain system has a particular CLI command designed to support the ARX proxy user: **cifs option set F5**. This command accepts the domain and username of any valid Windows account, defined externally on your Windows Domain Controllers (DCs):

```
cifs option set F5 domain\username
```

where

> ***domain*** is the Windows domain for the proxy-user name, as defined on your DCs.

> ***username*** is the name of an existing Windows user, also as defined on your DCs.

Then, run the **filesys restart** command. This gives the Windows user Backup Operator privileges on the Data Domain system. For example, the following command sequence provides Backup Operator privileges to "jqpublic" on the Data Domain system named "med-dd510:"

```
sysadmin@med-dd510# cifs option set F5 medarch.org\jqpublic
sysadmin@med-dd510# filesys restart
```

Options for CIFS Shares

By default, the EMC Data Domain system allows access to a small number of CIFS clients with read-only access. However, an ARX volume facilitates read/write access to the Data-Domain share by *any* of its CIFS clients; even if CIFS clients access this share infrequently, there is typically a wide variety of them. When you use the **cifs share create** command to create a CIFS share, use the following options to prepare it for an ARX volume:

```
cifs share create share-name path /backup/path clients *
browsing enabled writeable enabled
```

where

> ***share-name*** is the name of the Data Domain share, as seen by the ARX volume.

> ***path*** is directory path, which is always a subdirectory of /backup.

> **clients *** allows all CIFS clients to access the share.

> **browsing enabled** allows the ARX software to perform some necessary management functions.

> **writeable enabled** allows CIFS clients and the ARX software to write to the share.

For example, the following Data Domain command shares the "/backup/buusers" directory so that it can be used in an ARX-CIFS volume:

```
sysadmin@swic-dd510# cifs share create BUUSERS path
/backup/buusers clients * browsing enabled writeable enabled
```

Linux Samba

The ARX is often required to support file servers other than Windows, Network Appliance, and EMC. Some environments can include file servers that are Samba-based.

Samba is an open source suite of programs for Linux and UNIX that has been in use since 1992. Samba provides file and print services for clients using the SMB/CIFS protocol. Samba uses the winbind daemon to integrate Linux/UNIX servers and desktops into Active Directory environments.

If you plan to deploy the ARX in a Samba environment, disable Samba name mangling.

Windows 3.x and earlier used an 8.3 naming scheme, meaning that files had eight places for a name, following by a dot, followed by three places for an extension. To deal with files that have more than 8 characters followed by a dot and then three characters for a file extension, Samba uses *name mangling*.

To disable name mangling, insert the following line in your *mb.conf* file in the [global] section:

mangled names = no

The ARX supports Samba for multi-protocol namespaces without persistent ACLs. Without POSIX ACL and extended attribute support, the ARX namespace cannot support persistent ACLs.

If you intend to deploy Samba, work with your F5 SE and consult the latest version of the F5 Data Solutions Compatibility Matrix for details on the level of support provided.

The Data Solutions Compatibility Matrix (also called the Interoperability Matrix) is included in the standard ARX documentation (v5.01.000 and higher). There is also an F5 solution on askf5 linking to the document:

https://support.f5.com/kb/en-us/solu.../sol10909.html

Configuring Windows Domain Controllers

If you plan to offer CIFS services, your Windows Domain Controllers (DCs) must have the following ports accessible from the ARX:

• 389/TCP, used by the ARX for a Kerberos health check.

• 88/TCP, required for Kerberos queries and responses.

To support NTLM or NTLMv2, each DC requires a separate Secure Agent installation. Refer to the *ARX® Secure Agent Installation Guide* for details on this type of installation.

To support NTLM authentication and the Secure Agent, the DC firewall must also allow the ARX to access the following ports:

• 139/TCP or 445/TCP, used by an ARX health check.

- 25805/TCP, used by default for NTLM or NTLMv2 communication. You can change this port from the ARX and from the Secure Agent's management interface.

If you have a Windows firewall on your DCs, you may need to manually allow the ARX to access the above ports. We recommend leaving the firewall disabled on Windows 2003 and Windows 2000 DCs; the firewall is disabled by default on those releases.

Planning for a Namespace

If you have a single authentication domain for your network, a single namespace for all of your storage typically makes sense. You can offer the namespace through a single, consolidated VIP or through a series of VIPs, depending on the configuration of your client subnets.

If you have multiple authentication domains, you must set up one namespace for each.

Namespace Volumes

A namespace can contain up to three types of volumes: managed volumes, direct volumes, and/or shadow volumes. A *managed volume* keeps metadata about all of its files and directories. Metadata is information about the location of the files on back-end filers, among other things. The managed volume uses the metadata to manage these file locations through namespace policy. A *direct volume* does not have any metadata and therefore does not support any policy. A *shadow volume* is a copy of a managed volume, possibly located on another ARX.

Maximum Namespaces and Volumes per Platform

The number of supportable namespaces, volumes, or shares is dependent on the amount of memory and number of processors available for namespace software. The ARX-1500 and ARX-2500 each have multi-core processors, where each processor has a dedicated memory region. The ARX-4000 uses an 8-way symmetric multiprocessor, which runs as 8 independent processors each with a dedicated memory region. The ARX-2000 uses a 4-way multiprocessor, which runs as 4 independent processors. The ARX-500 has a single processor that runs management software in addition to namespace processes. The ARX-VE is a virtual machine, and has a configurable amount of processor time and memory.

The maximum namespaces, volumes, shares, and files per platform are controlled by these system resources and by the license that you purchase fromF5. The specific limits for each model are published here:

http://support.f5.com/kb/en-us/solutions/public/13000/100/sol13129.html

Refer to this chart to determine the limits for each ARX-appliance model. Specifically, the chart shows the following maximums:

- Namespaces (shown as the maximum Volume Groups),
- Volumes,
- Managed-Volume Shares,
- Managed-Volume Files,
- Direct (or Presentation) Volume Shares, and
- Direct (or Presentation) Volume Attach Points.

Direct Volumes and NFSv3

NFSv3 contains 64-bit fields for each of the following file identifiers:

- the NFS *file handle*, and
- the back-end filer's *file ID*.

An NFSv3 direct volume supports only 32 bits from its filers' file handles. Many NFS filers limit their file handles to this length for backward compatibility with NFSv2.

An NFSv3 direct volume supports 30 bits or fewer in its filers' File IDs. Many popular NFS clients support a maximum of 32 bits total, and the direct volume requires at least 2 of those bits (sometimes more) to differentiate the file IDs from different back-end filers.

A managed volume does not have these limitations with respect to Filehandles and File IDs.

The ARX Inter-Operability Matrix lists all of a direct volume's supported NFS filers and clients. This matrix is available from the ARX GUI, along with this document; click the Documentation link in the navigation panel to access all of the ARX documentation.

Choosing Storage for Metadata

The ARX namespace *metadata* is data that contains the physical location of files that the managed volumes are managing. It is used to find files and directories on the physical file systems. Each managed volume in a namespace maintains its own metadata for the file systems it manages. Metadata is maintained on a per managed volume basis with a volume being a collection of physical file systems. A namespace can consist of multiple managed volumes and each volume can have a unique metadata location.

You can completely rebuild the ARX metadata at any time by re-scanning the physical file systems. There is nothing proprietary in the metadata that would prevent it from being recreated if it is lost or corrupted. For this reason, F5 Networks does not recommend backing up the metadata. F5 Networks provides utilities to fix most metadata inconsistencies, and a rebuild utility to create a new metadata database.

Using the NFS or CIFS protocols, metadata is stored persistently on file servers external to the ARX so that both ARXes in a redundant pair will have access to the metadata database. Each volume within the namespace must have a metadata location specified although more than one managed volume can share the same metadata location.

In fact, best practice dictates that managed volumes sharing similar data locations should also have a common metadata location. Typically, all managed volumes within the same VPU/Domain will share the same metadata location as they are a single resource domain. If one volume within the VPU/Domain has a metadata outage it can affect other volumes within the domain. Thus, it is best to group like managed volumes that have data on the same file server together within a single domain, and typically co-locate the metadata on the same file server as the data so that the failure domain can be limited.

Metadata can be accessed using the CIFS or NFS protocols, and the metadata access protocol does not need to match the namespace protocol. If NFS is available it is considered the better protocol for metadata access because the protocol is more efficient and better suited for database applications due to its hard mount semantics.

F5 Networks recommends a dedicated file system on a file server that has adequate performance and redundancy characteristics, as well as adequate space. For details on estimating the size required, see *Sizing the Metadata Share*, on page 1-28.

The namespace software frequently accesses its metadata, so latency should be very low at the metadata share. A file server with a large amount of cache would be ideal. The location is configurable by the administrator.

The ARX will create a special hidden directory, .acopia, at the root of each imported share. Also, a .acopia directory is created on the metadata share where the switch will persistently store metadata. To optimize transaction performance, the ARX also maintains a cached copy of the metadata in its internal memory in addition to the copy kept in the external database.

◆ Note

Failure to specify a metadata location will prevent the namespace from becoming active. If an error message (such as % ERROR: VOLUME_METADATA_MISCONFIGURATION (38011052)), specify a metadata location and re-enable the managed volume.

File systems that the ARX will manage must be inventoried to populate the metadata database with file and directory location information. The inventory process is a background task run by the ARX that walks the tree of all the file systems in the managed volume once it is enabled. As the tree-walk learns of files and directories, it populates the metadata database with data about them.

Clients need not wait until the tree-walk is complete and the metadata database is fully populated. Client access to the namespace is instantaneous as soon as the managed volume and shares are enabled.

◆ **Important**

Choosing a metadata share that is too small or unreliable can seriously jeopardize managed-volume performance. A poor configuration for the metadata can possibly lead to metadata corruption. Choose your metadata shares carefully, as instructed in this section.

You must choose at least one metadata-storage location for each managed volume. Viable metadata storage options are as follows:

•   Metadata-only share for the volume

•   Metadata-only share for the namespace

Alternatively, a very large storage volume can host metadata-only shares for multiple volumes or even multiple namespaces. In general, each managed volume requires 3-5 gigabytes of storage space for its metadata, so you should be able to use a 140-gigabyte volume for a total of 28 managed volumes (140G/5G per managed volume).

If multiple switches use the same metadata-only shares, each switch must have its own directory tree (for example, /vol/vol1/switch1 and /vol/vol1/switch2).

The sections below present some considerations in choosing metadata-storage candidates for a namespace. Use the information in these sections as guidelines before configuring any namespaces with managed volumes.

Sizing the Metadata Share

A metadata share requires multiple gigabytes of storage. If a managed volume's metadata share falls below 512 megabytes (0.5 gigabytes) of free space, the volume software sends an SNMP trap and issues a warning to the syslog. If it falls below 256 megabytes, stronger warnings appear. If it falls to 128 megabytes or less, the volume is inaccessible to clients until more free space is added to the share. This insures against metadata corruption.

Each file requires roughly 300 to 800 bytes of metadata, and each directory requires roughly 2 Kbytes. These numbers vary for each protocol: for example, the NFSv2 protocol uses a fixed-length 32-byte file handle, whereas NFSv3 uses a variable length file handle from 8-64 bytes. If the namespace supports multiple protocols, each file and directory requires additional metadata for each protocol. The numbers also vary for each file, since the metadata contains the file name. Use the byte requirements in the following table for rough estimations. The following table shows the byte size for each protocol.

Table 1.3 *Size of protocol in bytes.*

Protocol	Bytes
NFSv2	395
NFSv3	395
NFSv2 + NFSv3	485
CIFS	330
NFSv3 + CIFS or NFSv2 + CIFS	395
NFSv2 + NFSv3 + CIFS	815

For example, consider a volume with 5 million files and 1 million directories. The estimated disk capacity needed for the volume would be 3.5 gigabytes:

- Files: (5 million) x (300 bytes per file) = 1,500,000,000 bytes or 1.5 gigabytes

- Directories: (1 million) x (2 Kbytes per directory) = 2 billion bytes or 2 gigabytes

For a volume in a single-protocol namespace this estimate would probably be high. In a three-protocol namespace, the estimate would probably be low.

Check the Limit on File Size

The metadata for each managed volume is stored in a single database file. As shown, a volume with millions of files can require multiple gigabytes of metadata. However, a file system has a limit on the size of each file. Thus, you must verify that the file system for the metadata share can support an individual file that is large enough to hold all metadata in its largest volume.

Performance and Availability Requirements

F5 recommends that all metadata-share candidates have good performance and high availability. Namespace performance suffers if its metadata is housed on a physical device that is overwhelmed, and a namespace cannot function at all if its metadata is unavailable. Ideally, all metadata shares are housed in a high-performance storage cluster.

If the share's server acknowledges a write, that acknowledgement must amount to a guarantee that the data is written to disk. Use the following steps to ensure that this guarantee extends from the NFS/CIFS mount point all the way to the disk drive:

1. Configure the export/share for synchronous writes.

   All CIFS shares support synchronous writes, as do most implementations of NFS. For NFS servers that do not support synchronous writes by default (such as many Linux implementations), set the *sync* option for the NFS export. To improve volume performance, also set the *no_wdelay* option on the export.

2. Mount the server's file system for synchronous writes.

3. Configure the server's hard drives for write-through (no disk-caching).

Do not use any disk quotas on a metadata share, as they often report incorrect information or do not conform to appropriate synchronous write semantics.

NFS Metadata Shares Can Support Any Volume

An NFS metadata share can support any volume (NFS-only, CIFS-only, or multi-protocol), whereas a CIFS metadata share can only support a volume that also supports CIFS. This is because an NFS-only volume does not have the CIFS credentials required to access a CIFS metadata-only share.

Avoid /vol/vol0 on NetApp

By default, a NetApp file server uses /vol/vol0 for its operating system, making it a poor candidate for metadata storage. Typically, /vol/vol0 is designed for very high reliability at the expense of fast access. Use a NetApp volume that is configured for both reliability and speed.

Best Practices (Summary)

Choose a highly-available, synchronous file system that is dedicated exclusively to metadata. The metadata share should be on a file system that is not shared by any other device, should have multiple gigabytes of free space, and should not be subjected to disk quotas. To increase the space, we recommend configuring a discrete metadata-only share for each volume. Choose a share with a balance of reliability and speed. If you use a NetApp share, do not use /vol/vol0. An NFS metadata share supports all types of managed volumes, whereas a CIFS metadata share is limited to volumes that also support CIFS.

F5 Networks' web-based application, Data Manager, can help you inventory the shares/exports you plan to import and calculate the ARX metadata size required. If you have access to Data Manager, you can use it to

automatically discover key attributes about your file storage environment, such as volumes, shares, exports, security settings, file system settings, and estimated size required for metadata storage. For more information and to download a free trial version, see:

http://www.f5.com/products/data-manager/

Choosing a Filer Share for the Volume Root

When a client creates a new file in the root (/) of a managed volume, the namespace software cannot deterministically choose an intended back-end share for the new file. By default, it places the new file in the volume's first imported share.

This filer share is said to hold the *volume root*. Choose the root share for each of your volumes in advance. Configure the root share first as you configure your namespace volumes.

As an alternative, you can configure namespace policy to distribute new files and directories amongst several shares in the volume. This policy is called *new-file placement* policy in a *share farm*, and it is described in the *ARX® CLI Storage-Management Guide*.

Best Practices for Share Management

Once an external share is imported into a managed volume, you should manage the share entirely through the ARX. Specifically, do not edit any files in the share (especially the .acopia directory at the share's root), and do not remove the share or make it unavailable.

No Overlapping Shares

If one share contains another (a fairly common scenario in a CIFS environment), only import *one* of the shares. Overlapping shares, imported into one or more namespaces, invariably cause namespace corruption. For example, suppose you have a C:\home share that contains C:\users\jrandom and C:\users\juser. You can import C:\home into a namespace volume, or you can import one or both of the subdirectories. Do not import *both* C:\home *and* any of its subdirectories.

For a CIFS subshare with a different Access-Control List (ACL) than its parent share, you can use a special *CIFS subshare* feature on the ARX. You use this feature to identify any subshares after the parent share is imported, and share them out to your CIFS clients.

Planning for a Multi-Protocol Namespace

A multi-protocol namespace supports both NFS and CIFS. You can use it with a heterogeneous set of back-end filers with multi-protocol shares. Currently-supported filer vendors include NetApp and EMC.

The ARX passes NFS and/or CIFS operations through to the back-end shares, which permit or deny the operations based on client identity and file/directory permissions. Each filer permits or denies access to files based on its own rules, and the ARX passes back the response back to the client.

◆ **Important**

This has far-reaching implications for client access in a heterogeneous volume. If you mix multi-protocol shares from multiple vendors in the same volume, client access can change when files migrate from one filer to another. This is an unavoidable consequence of different multi-protocol implementations from the filer vendors.

Consult your filer documentation (from all vendors) concerning client access and the recommended security configurations. Pay particular attention to *non-native access* to the filer. Non-native access means accessing a UNIX file through CIFS, or an NTFS (Windows) file through NFS. Of particular interest are the following questions:

◆ How will the client's identity (UID/GID in UNIX, owner SID and primary-group SID in NTFS) be interpreted in the file's environment?

◆ Is the client's access restricted in unexpected ways (for example, can an NFS client delete NTFS files in a directory where they apparently have UNIX write and execute permissions)?

◆ Is the client allowed to change permissions on the non-native file? If so, how are the new permissions interpreted in the file's environment?

Our best practice is to use NTFS Qtrees for NetApp filers whenever you mix them with EMC file servers. The following subsections discuss specific best practices for various vendors.

NetApp Best Practices

Client Identity

A NetApp filer maps UNIX usernames to NT usernames. If a client has the same username for both operating systems, the mapping is straightforward. For clients with different usernames in each environment, you can configure a map on the NetApp, /etc/usermap.cfg. After the NetApp maps the username, it consults the operating system to find the user's Group ID (in UNIX) or primary-group SID (in NT).

Non-Native File Permissions (and Qtree Configuration)

Each NetApp qtree can support UNIX-based permissions or NTFS-based ACLs. The client's identity is mapped as discussed above, then their identity is applied to the file's native permissions. For example, consider an NFS client who accesses a file on an NTFS qtree: the client's UNIX username is mapped to an NT username, then the NT identity is applied to the file's ACLs.

◆ **Note**

NetApp mixed-mode qtrees are not currently qualified.

A multi-protocol namespace passes all permissions checks back to the qtrees, which permit or deny access to each file or directory. All NetApp shares behind a multi-protocol volume must be configured with the same permissions type (UNIX, or NTFS).

F5 recommends that you use NTFS qtrees, which offer the richest set of file-access permissions.

Timestamp Skew in a Unix Qtree

A NetApp share with a Unix Qtree can create a unique timestamp skew: some files may show different timestamps to Unix clients than to CIFS clients. The timestamps differ by one hour. This is an indication that the timestamp was set before daylight savings time started or ended; it is a known issue on NetApp filers. We recommend that you follow NetApp's best practices and set the filer to GMT.

Mapping the Proxy User

As mentioned earlier, the configured proxy user must have full read/write privileges from both NFS and CIFS. The NetApp's NT/UNIX user map must equate the proxy-user credentials on the NT side with *root* on the UNIX side. The user map is in /etc/usermap.cfg, which you can access from an NFS client by mounting /vol/vol0. For example, this command sequence mounts the NetApp filer at 192.168.25.21 and lists the usermap.cfg file:

```
rh1:/mnt# mount 192.168.25.21:/vol/vol0/etc netapp/
rh1:/mnt# ls -l netapp/usermap.cfg
-rwxrwx---   1 root     root         1385 Apr 25  2005 netapp/usermap.cfg
rh1:/mnt#
```

One line maps the Windows proxy user to *root*. Follow this syntax:

DOMAIN\proxy-username == root

where

DOMAIN is the Windows domain for the proxy user (use the short version; for example, MYDOMAIN instead of MYDOMAIN.MYCOMPANY),

proxy-username is the Windows username, and

the spaces before and after **==** are required.

If the proxy-username has spaces or a pound-sign (#) character in it, you must enclose it in quotation marks. Enclose only the username, not the entire *DOMAIN\proxy-username* string. For example:

```
MYDOMAIN\"random user" == root # correct
```

is correct, but

```
"MYDOMAIN\random user" == root # incorrect
```

is incorrect.

As the final example, this line maps a Windows user, MEDARCH\jqpublic, to root:

```
MEDARCH\jqpublic == root
```

EMC Best Practices: Creating the Proxy-User Account

EMC Celerra servers require a new, unused account for a proxy-user, immediately mapped to *root* on the UNIX side. If a client has already authenticated with a particular username and password, it would be prohibitively difficult to re-map the username to *root* on an EMC. EMC Release 5.5.24.2 introduces a command to resolve this problem; these instructions apply to prior releases.

Start by creating a new Windows account for the proxy user. Remember to add it to the Backup Operator's or Administrator's group on the EMC, or give it equivalent privileges. Then map the above user to *root* on the UNIX side, as described below.

From the EMC CLI, use the following command to get a copy of the UNIX "passwd" file:

server_file *data-mover* -get passwd passwd.new

> where *data-mover* identifies the data mover behind the ARX.

This places copy of the passwd file in a local file, "passwd.new." Edit this file and add the following line to the bottom:

username::0:0:*anything*:*home-dir*:*path-to-shell*

> where *username* is the name for the proxy-user account that you created from Windows. The two zeros in the third and fourth fields are the required UID and GID for *root*. The values for the remaining fields are outside the scope of this document; you can use **man 5 passwd** from the EMC CLI to access the EMC documentation.

Then use the following command to install the passwd.new file and put it into service:

server_file *data-mover* -put passwd.new passwd

For example, the following command sequence accesses the CLI of an EMC Celerra and maps the "jqpublic" user to *root*. (The name of the data mover is "server_2" in this example.)

```
rh1:/# ssh nasadmin@192.168.25.51
nasadmin@192.168.25.51's password: password
Last login: Thu Mar 22 09:37:42 2007 from juser.wwmed.com
EMC Celerra Control Station Linux Mon Nov 20 12:42:47 EST 2006
```

```
            *** slot_0 primary control station ***

[nasadmin@emc01-mgt nasadmin]$ server_file server_2 -get passwd passwd.new
server_2 : done
[nasadmin@emc01-mgt nasadmin]$ vi passwd.new
```

Add the following line to the bottom of the "passwd.new" file:

`jqpublic::0:0:jqpublic:/:/bin/bash`

Then install the new passwd file and exit the EMC CLI:

```
[nasadmin@emc01-mgt nasadmin]$ server_file server_2 -put passwd.new passwd
server_2 : done
[nasadmin@emc01-mgt nasadmin]$ exit
Connection to 192.168.25.51 closed.
rh1:/# ...
```

Preparing for Use in an Active Directory Domain

This section applies to Windows installations that support Active Directory (AD) domains. These installations use Kerberos authentication and require some special configuration at the Domain Controller (DC) machine.

Figure 1.12 shows how the ARX acts as a resource proxy between clients and servers. In its role as a proxy, it must carry the identity of a client user through to the back-end servers. This allows for already-established Access Control Lists (ACLs) to continue their role in controlling access to files. This also makes the ARX transparent to users in an AD domain. The ARX authenticates a client once, using Kerberos, then uses the client's credentials to access any server that contains a requested file.

Figure 1.12 ARX® as a resource proxy between clients and servers
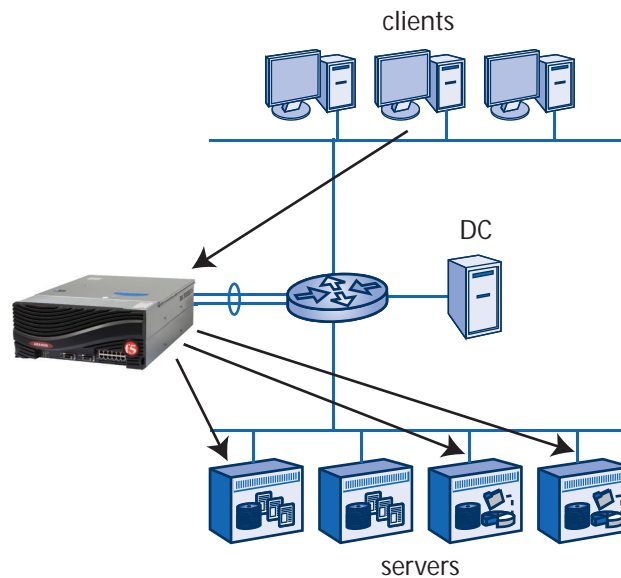
Kerberos authentication simplifies security management: all ACLs on all servers can remain the same, and all clients retain the same rights and restrictions that they had before inserting the ARX. This proxy mechanism, called *delegation* in Windows terminology, is possible with Kerberos but not with NTLM.

Required Administrative Privileges

Special administrative privileges are required to join an F5 front-end CIFS server (*F5 server*) to an AD domain. The domain-join operation has two major steps: add the F5 server to the AD domain and raise the "Trusted for Delegation" flag for the server. Each of these steps requires a distinct administrative privilege:

- "Add workstations to domain" (where the "workstation" is the F5 server), and

- "Enable computer and user accounts to be trusted for delegation."

An administrator in the Domain Admins group has both of these privileges. You need the username and password of one of these administrators to join an F5 server to an AD domain.

◆ **Note**

Trusting an F5 server for delegation poses no security threat to your network. Kerberos authentication was designed with delegation in mind to provide a clean way of carrying identity through n-tiered application systems. For more information, refer to IETF RFC 1510 or the Microsoft

*white paper on Kerberos authentication
(http://www.microsoft.com/windows2000/techinfo/howitworks/security/kerb
eros.asp).*

Front-End Service Limitations

The ARX aggregates all of your back-end storage and offers it through
front-end (CIFS and/or NFS) services to your clients. The ARX supports a
maximum of 64 services, each with a distinct FQDN and virtual-IP (VIP)
address. This maximum is for the sum of all CIFS and NFS services. The
smaller ARX-500 platform offers up to 16 distinct services.

The maximum number of CIFS shares that you can offer to your ARX
clients is 16,000.

Physical Site Preparation

Table 1.4 describes the basic features of each ARX. For more detailed
hardware requirements, refer to *System Specifications and Requirements*, on
page 1-39 or refer to the *ARX Hardware Reference Guide*.

Table 1.4 shows an overview of the features for each ARX.

Table 1.4 *ARX Models and Description*

Model	Description	Hardware Feature
500	The ARX-500 is a cost-effective, small-form-factor Adaptive Resource Switch designed for use in small data centers and branch/remote offices. It combines application processing and control, switch fabric throughput, and external interfaces into a single FRU compact design.	• Supports Fast Ethernet and Gigabit Ethernet throughput including: – OOB management port – 100/1000 BASE-T client/server port for connecting network infrastructure and NAS – 100/1000 BASE-T port for a dedicated link to a redundant peer (called the redundancy link) – RJ-45 connector for console connection • 1U compact design • Weight: 31 pounds • LEDs on the front panel for overall system status • LEDs on the rear panel for ACM status • One IDE hard disk drive • Throughput: 100 MB/s • Power supply: Not redundant
1500	The ARX-1500 is a cost-effective, small-form-factor Adaptive Resource Switch designed for use in small data centers and branch/remote offices. Similar to the ARX-500, the ARX-1500 combines application processing and control, fabric throughput, and external interfaces into a single FRU compact design. It offers the same software features as the other ARX® platforms, differing only in performance and scale.	• Ports: 8 x 1GbE copper (1 mgmt/data and 7 data) • Throughput: 500MB/s • Files: 768M • Power supply: redundant • Disk drives: dual, hot-swappable • NVRAM with SuperCap (No Battery required) • Height:1U • Weight: 22.5 • Licensing and entitlement tracking using F5 Licensing • Weight: 35 pounds
2000	The ARX-2000 is a cost-effective, small-form-factor Adaptive Resource Switch designed for use in small data centers and branch/remote offices. The ARX-2000 combines application processing and control, switch fabric throughput, and external interfaces into a single FRU compact design. It offers the same software features as the other ARX® platforms, differing only in performance and scale.	• Enables Fast Ethernet and Gigabit Ethernet throughput • Provides 12 100/1000 BASE-T external ports for connectivity to network infrastructure • (2U) Compact design • Weight: 35 pounds • Box-to-box failover capability to provide redundancy. See Configuring Redundant Pairs, on page 2-4. • External interfaces, including: – Serial console port – 10/100/1000 Mbps out-of-band (OOB) Ethernet management port – 100/1000 BASE-T Ethernet ports (12, copper) • Power Supply: Auto-sensing, redundant (110-220V) • 146GB internal SAS hard drives (2) hot-swappable, configured as RAID1 • AC power cords. The power cords (2) are rated for 15 A/120 VAC with IEC-320 type connector.

Table 1.4 *ARX Models and Description (Continued)*

Model	Description	Hardware Feature
2500	The ARX-2500 is a cost-effective, small-form-factor Adaptive Resource Switch designed for use in small data centers and branch/remote offices. Similar to the ARX-2000, the ARX-2500 combines application processing and control, switch fabric throughput, and external interfaces into a single FRU compact design. It offers the same software features as the other ARX® platforms, differing only in performance and scale.	• Ports: 4 x 1GbE copper (1 mgmt/data and (3) 1 GbE and 2 x 10GbE • 10GbE ports are optical with SFP+ transceivers • Throughput: 400 to 500 MB/s • Files: 1536M • Power supply: redundant • Disk drives: dual, hot-swappable • NVRAM with SuperCap (No Battery required) • Height: 1U • Weight: 22.5 • Licensing and entitlement tracking using F5 Licensing • Weight: 35 pounds
4000	The ARX-4000 is a cost-effective Adaptive Resource Switch that combines application processing and control, switch fabric throughput, and external interfaces into a compact design consisting of two FRUs. The switch includes a removable front bezel which, when attached, covers all components except the Ethernet ports and a set of data plane LEDs.	• Compact design (2U control plane, 2U data plane) • Weight: 96 pounds • Throughput: 1050 MB/s • Box-to-box failover capability to provide redundancy. • Single bezel covering control and data plane components. • AC power cords (4): 15 amp, 110 volts and 220 volts. • PCI-E interconnect cable (control plane to data plane). • External interfaces, including: – Serial console port – 10/100 Mbps out-of-band (OOB) Ethernet management port – 100/1000 gigabit Ethernet ports (12 ports, copper) – 10-gigabit Ethernet ports (2 ports, fiber-optic) • Front control panel LEDs that indicate port activity, disk drive activity, system status, and power. • Front data panel LEDs that indicate PCI link status, NVRAM battery status, power supply status, and various operational states. • Power supply: (4) with the following characteristics: 1:1 redundant, load-sharing, universal, auto-sensing, hot-swappable (110 volts and 220 volts). • 146 GB internal SAS hard disks (2), hot-swappable, configured as RAID1.

System Specifications and Requirements

This section details the system specifications and requirement of each ARX. Choose from the following:

ARX-500 System Specifications and Requirements

The following table describes the ARX-500 system specifications.

Table 1.5 *ARX-500 System Specifications*

Component	Specification
Dimensions	Height: 1.703 in. Width: 16.930 in. Depth: 26.457 in.
Weight	31 lb (14.061 kg)
Power Load	8.55 amps @ 110Vac, 4.3 amps @ 220Vac
Environmental Requirements	Altitude: -200 ft. (-60 m) min. to 8000 ft. (2500 m) max.
	Humidity Operating: 90% relative humidity (non condensing) at 30 deg. C Storage: 5% to 95%
	Temperature Operating: 50 deg. to 95 deg. F (10 deg. to 35 deg. C) Storage: -40 deg. to 158 deg. F (-40 deg. to 70 deg. C)

ARX-1500 System Specifications and Requirements

The following table describes the ARX-1500 system specifications.

Table 1.6 *ARX-1500 System Specifications*

Component	Specification
Chassis Dimensions (includes front bezel)	Height: 44.5 mm (1.75 inches) Width: 443 mm (17.44 inches) Depth: 481 mm (18.93 inches)
Weight	Weight: 22.5 lb (10.2 kg)
Power Load	2.5A @ 110V, 1.3A @ 220V
AC/DC Power Supply	Input: 100Vac @ 4 Amps/240Vac @ 2Amps, 47-63Hz 100-240VAC, 47-63Hz, 4-2A

Table 1.6 *ARX-1500 System Specifications (Continued)*

Component	Specification
Environmental Requirements	Altitude: 60m (197ft) min. to 1800m (6000 ft) max.
	Humidity Operating: 10% min. to 95% max. (non condensing) Storage: 5% to 95%
	Temperature Operating: 32 deg. to 104 deg. F (0 deg. to 40 deg. C) Storage: -40 deg. to 149 deg. F (-20 deg. to 65 deg. C)

ARX-2000 System Specifications and Requirements

The following table describes the ARX-2000 system specifications:

Table 1.7 *ARX-2000 System Specifications*

Component	Specification
Chassis Dimensions (includes front bezel)	Height: 3.5 in. Width: 19.00 in. (including mounting ear assemblies) Depth: 24 in.
Chassis Weight	40 lb (18.14 kg)

Power Consumption	Typical power consumption:
	    2.5A @ 110VAC (280W; 0.98 PF)
	Maximum power consumption:
	    8.4A @ 100VAC (maximum rated current)
	    450W (maximum rated power)

AC/DC Power Supply	700 Watts 1+1 redundancy > 80% efficiency Input: 100VAC – 240VAC, 47/63HZ Output: -12V @ 1A, +3.3V @ 32A, +5V 30A, +12V 62A
Environmental Requirements	Altitude -200 ft. (-60 m) min. to 8000 ft. (2500 m) max. Humidity Operating: 5 % min. to 95% max. (non condensing) Storage: 5% to 95% Temperature Operating: 41 deg. to 95 deg. F (5 deg. to 35 deg. C) Storage: -4 deg. to 149 deg. F (-20 deg. to 65 deg. C)

ARX-2500 System Specifications and Requirements

The following table describes the ARX-2500 system specifications.

Table 1.8 *ARX-2500 System Specifications*

Component	Specification
Chassis Dimensions (includes front bezel)	Height: 44.5 mm (1.75 inches) Width: 443 mm (17.44 inches) Depth: 481 mm (18.93 inches)
Weight	Weight: 22.5 lb (10.2 kg)
Power Load	2.5A @ 110V, 1.3A @ 220V
AC/DC Power Supply	Input: 100Vac @ 4 Amps/240Vac @ 2Amps, 47-63Hz 100-240VAC, 47-63Hz, 4-2A
Environmental Requirements	Altitude: 60m (197ft) min. to 1800m (6000 ft) max
	Humidity Operating: 5% min. to 95% max. (non condensing) Storage: 5% to 95%
	Temperature Operating: 32 deg. to 104 deg. F (0 deg. to 40 deg. C) Storage: -40 deg. to 149 deg. F (-20 deg. to 65 deg. C)

ARX-4000 System Specifications and Requirements

The following table describes the ARX-4000 system specifications.

Table 1.9 *ARX-4000 System Specifications*

Component	Specification
Chassis Dimensions (includes front bezel)	Height: 7.00 in. Width: 19.00 in. (including the fixed mounting ears) Depth: 24.0 in.
Weight	96 lb (43.54 kg)
Power Load	Control plane: 5.5 amps @ 110VAC and 2.75 amps @ 220VAC Data plane: 3.5 amps @ 110VAC and 1.75 amps @ 220VAC

Table 1.9 *ARX-4000 System Specifications (Continued)*

Component	Specification
Environmental Requirements	Altitude: –200 ft. (–60 m) min. to 8000 ft. (2500 m) max.

	Humidity:
	    Operating: 5 % min. to 95% max. (non condensing)
	    Storage: 5% to 95%

	Temperature:
	    Operating: 50° to 95° F (10° to 35° C)
	    Storage: -40° to 149° F (-20° to 65° C)

System Power Requirements

This section lists the system power requirements for all the ARX models. Choose from the following:

- *ARX-500 System Power Requirements*, on page 1-43

- *ARX-1500 System Power Requirements*, on page 1-43

- *ARX-2000 System Power Requirements*, on page 1-43

- *ARX-2500 System Power Requirements*, on page 1-44

- *ARX-4000 System Power Requirements*, on page 1-44

ARX-500 System Power Requirements

The ARX-500 power supply distributes up to 600 Watts of DC power to the chassis components. The power supply runs at 72% efficiency, so it consumes up to 833 Watts of AC power (833/0.72) to meet the 600-Watt demand. This is equivalent to 2,843 BTUs/hour.

ARX-1500 System Power Requirements

The ARX-1500 is powered by two power supplies (1+1 redundancy) system) consisting of two power modules and one power system frame. Two power modules are recommended for full redundancy and load-sharing. The power supplies require a 10A / 220VAC input cord, which is provided with the chassis. The AC outlet to the switch must be properly grounded.

ARX-2000 System Power Requirements

The ARX-2000 uses two power supplies (1+1 redundancy). The power supplies require a 10A / 220VAC input cord, which is provided with the chassis. The AC outlet to the switch must be properly grounded.

The ARX-2000 power supply can distribute up to 700 Watts of DC power to the chassis components. The power supply runs at a minimum of 80% efficiency, so it is capable of consuming up to 563 Watts of AC power (700 * 0.80). This is equivalent to 1,536BTUs/hour.

ARX-2500 System Power Requirements

The ARX-2500 is powered by two power supplies (1+1 redundancy) consisting of two power modules and one power system frame. Two power modules are recommended for full redundancy and load-sharing. The power supplies require a 10A / 220VAC input cord, which is provided with the chassis. The AC outlet to the switch must be properly grounded.

ARX-4000 System Power Requirements

The ARX-4000 chassis dissipates 975 watts of power. The control plane dissipates 600 watts and the data plane dissipates 375 watts. This is equivalent to 3328 BTUs/hour.

Cable Requirements

This section contains the cable requirements for all the ARX models. Choose from the following:

- *ARX-500 Cable Requirements*, on page 1-44
- *ARX-1500 Cable Requirement*, on page 1-45
- *ARX-2000 Cable Requirements*, on page 1-46
- *ARX-4000 Cable Requirements*, on page 1-47

ARX-500 Cable Requirements

The following table lists the required cables and power cords for the ARX-500. All cables *except* the AC power cord and console cable are customer-supplied.

Table 1.10 ARX-500 Required Power and Data Cables

Qty.	Cable/Cord	Used on...	Specification
1	AC power cord	AC/DC Power Supply	You can choose from the two cables shipped with the switch: • 20 A/250 Vac or • 15 A/120 Vac Both have IEC-320 type connectors. For sites in Europe and Great Britain, we ship a power cord that is compatible with local standards.

Table 1.10 *ARX-500 Required Power and Data Cables (Continued)*

Qty.	Cable/Cord	Used on...	Specification
1	Console cable with RJ-45-to-DB9 adapter	Serial console interface (labeled "10101" on the rear panel)	100BASE-T Category 5 unshielded twisted pair (UTP); 24 AWG
1	Ethernet cables for connection to 10/100 Mbps Ethernet management port (RJ-45 connector);	OOB management interface (labeled "2" on the rear panel)	
2	Ethernet cables for connection to 100/1000 Mbps Ethernet (RJ-45 connectors)	Two copper Gigabit Ethernet ports: • one for client/server traffic (Port 1/1), and • one for the dedicated redundancy link (Port 1/2)	[a]100/1000BASE-T Category 5/6, unshielded twisted pair (UTP) cable; 24 AWG.

a.Gigabit Ethernet ports support automatic MDI/MDIX cross-over. This feature automatically corrects the polarity of the attached CAT5 cable, regardless whether it is a cross-over or straight-through type. However, for this feature to work, the port speed must be set to auto (auto-negotiate) through the CLI. When the port speed/duplex is forced (auto-negotiate is disabled), automatic MDI/MDIX cross-over is disabled, and you must cable the port using standard cross-over or straight-through cabling.

ARX-1500 Cable Requirement

The Following table lists the required cables and power cords for the switch. All cables *except* the AC power cord and console cable are customer-supplied.

Table 1.11 *ARX-1500 Required Power and Data Cables*

Qty.	Cable/Cord	Used on...	Specification
2	AC power cord	AC/DC power supply	10A/220 VAC
1	Console cable with RJ-45-to-DB9 adapter	Console port (labeled CONSOLE)	100BASE-T Category 5 unshielded twisted pair (UTP); 24 AWG
1	Ethernet cable (RJ-45 connector)	Management port (labeled MGMT)	
7	Ethernet cable (RJ-45 connectors)	Gigabit Ethernet ports, copper	[a]100/1000BASE-T Category 5/6, unshielded twisted pair (UTP) cable; 24 AWG.

a.Gigabit Ethernet ports support automatic MDI/MDIX cross-over. This feature automatically corrects the polarity of the attached CAT5 cable, regardless whether it is a cross-over or straight-through type. However, for this feature to work, the port speed must be set to auto (auto-negotiate) through the CLI. When the port speed/duplex is forced (auto-negotiate is disabled), automatic MDI/MDIX cross-over is disabled, and you must cable the port using standard cross-over or straight-through cabling.

ARX-2000 Cable Requirements

The following table lists the required cables and power cords for the switch. All cables *except* the AC power cord and console cable are customer-supplied.

Table 1.12 ARX-2000 Required Power and Data Cables

Qty.	Cable/Cord	Used on...	Specification
2	AC power cord	AC/DC power supply	15 A/120 VAC with IEC-320 type connector
1	Console cable with RJ-45-to-DB9 adapter	Console port (labeled CONSOLE)	100BASE-T Category 5 unshielded twisted pair (UTP); 24 AWG
1	Ethernet cable (RJ-45 connector)	Management port (labeled MGMT)	
4	Ethernet cable (RJ-45 connectors)	Gigabit Ethernet ports, copper	[a]100/1000BASE-T Category 5/6, unshielded twisted pair (UTP) cable; 24 AWG.

a.Gigabit Ethernet ports support automatic MDI/MDIX cross-over. This feature automatically corrects the polarity of the attached CAT5 cable, regardless whether it is a cross-over or straight-through type. However, for this feature to work, the port speed must be set to auto (auto-negotiate) through the CLI. When the port speed/duplex is forced (auto-negotiate is disabled), automatic MDI/MDIX cross-over is disabled, and you must cable the port using standard cross-over or straight-through cabling.

ARX-2500 Cable Requirements

The following table lists the required cables and power cords for the switch. All cables *except* the AC power cord and console cable are customer-supplied.

Table 1.13 ARX-2500 Required Power and Data Cables

Qty.	Cable/Cord	Used on...	Specification
2	AC power cord	AC/DC power supply	10A/220 VAC
1	Console cable with RJ-45-to-DB9 adapter	Console port (labeled CONSOLE)	100BASE-T Category 5 unshielded twisted pair (UTP); 24 AWG
1	Ethernet cable (RJ-45 connector)	Management port (labeled MGMT)	

Table 1.13 *ARX-2500 Required Power and Data Cables (Continued)*

Qty.	Cable/Cord	Used on...	Specification
2	Fiber-optic cables for connection to 10-Gbps Ethernet X2 MSA-compliant form factor	10-gigabit Ethernet ports	10GBASE-SR (gigabit Ethernet) fiber cable: Short-reach multi-mode fiber (MMF) with duplex LC-style connectors. Distances up to 82m on 50/125um MMF, or 26m on 62.5/125um MMF.
3	Ethernet cable (RJ-45 connectors)	Gigabit Ethernet ports, copper	[a]100/1000BASE-T Category 5/6, unshielded twisted pair (UTP) cable; 24 AWG.

a.Gigabit Ethernet ports support automatic MDI/MDIX cross-over. This feature automatically corrects the polarity of the attached CAT5 cable, regardless whether it is a cross-over or straight-through type. However, for this feature to work, the port speed must be set to auto (auto-negotiate) through the CLI. When the port speed/duplex is forced (auto-negotiate is disabled), automatic MDI/MDIX cross-over is disabled, and you must cable the port using standard cross-over or straight-through cabling.

ARX-4000 Cable Requirements

The following table lists the required power cords and cables. All cords are customer-supplied *except* the AC power cords and the console cable.

Table 1.14 *ARX-4000 Required Cables and Power Cords*

Qty.	Cord/Cable	Used on...	Specification
4	AC power cords	AC/DC power supplies	You can choose from the following types of cables shipped with the switch: • 20 A/250 VAC or • 15 A/120 VAC Both types have IEC-320 type connectors.
1	Console cable (flat, crossover) with RJ-45-to-DB9 adapter	Console port	100BASE-T Category 5 unshielded twisted pair (UTP); 24 AWG
1	PCI-E interconnect cable	Control plane to data plane	
1	Ethernet cable for connection to 10/100/1000 Mbps Ethernet management port (RJ-45 connector);	Management interface	

Table 1.14 *ARX-4000 Required Cables and Power Cords (Continued)*

Qty.	Cord/Cable	Used on...	Specification
12	Ethernet cables for connection to 100/1000 Mbps Ethernet (RJ-45 connectors)	Gigabit Ethernet ports	[a]100/1000BASE-T Category 5/6, unshielded twisted pair (UTP) cable; 24 AWG.
2	Fiber-optic cables for connection to 10-Gbps Ethernet X2 MSA-compliant form factor	10-gigabit Ethernet ports	10GBASE-SR (gigabit Ethernet) fiber cable: Short-reach multi-mode fiber (MMF) with duplex SC-style connectors. Distances up to 300m on 50/125um MMF, or 33m on 62.5/125um MMF.

a.Gigabit Ethernet ports support automatic MDI/MDIX cross-over. This feature automatically corrects the polarity of the attached CAT5 cable, regardless of whether it is a cross-over or straight-through type. However, for this feature to work, the port speed must be set to auto (auto-negotiate) through the CLI. When the port speed/duplex is forced (auto-negotiate is disabled), automatic MDI/MDIX cross-over is disabled, and you must cable the port using standard cross-over or straight-through cabling.

Cable Connectors and Pinout Assignments

This section shows the cable connectors and the pinout assignments for all models of the ARX.

The serial console port requires a rollover cable (RJ-45 to RJ-45) that is included with the ARX-installation kit. This cable is sufficient for connecting to a Terminal Server. For a direct connection to the serial port on a management station (such as a laptop), an RJ-45 to DB9 adapter is also included in the kit.

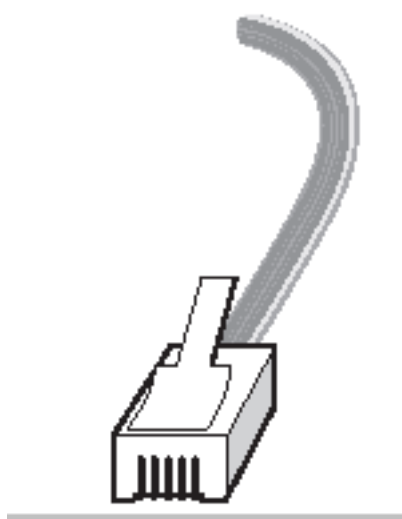*Figure 1.13 RJ-45 Male Connector*

Figure 1.14 RJ-45 to Serial DB9 Adapter

Pinout Assignments

The pinout assignments for the ARX-500, ARX-2000, and ARX-4000. The ARX-1500 and ARX-2500 have a different pinout.

Pinout Assignments for ARX-500, ARX-2000, and ARX-4000

The following table lists the RJ-45 pinout assignments for the rollover cable and the adapter. The left column shows the transmit (TxD), ground (GND), and receive (RxD) signals. and the right column shows the signals reversed at the console device. The intervening columns show the pins that carry each of those signals.

Table 1.15 *SCM/ACM console port signaling/cabling using a rollover cable for ARX-500, ARX-2000, and ARX-4000.*

SCM/ACM Console Port	RJ-45 Rollover Cable			RJ-45 to DB9 Adapter			Console Device
DTE Signal	RJ-45 Pinout	USOC Color	RJ-45 Pinout	RJ-45 Pinout	T568 Color	DB9F Pinout	DTE Signal
TxD	3	yellow	6	6	yellow	2	RxD
GND	4	green	5	5	green	5	Signal Ground
GND	5	red	4	4	red		
RxD	6	black	3	3	black	3	TxD

Pinout Assignments for the ARX-1500 and the ARX-2500

The console port uses a RJ45 connector on the system. Here is the table that describes the console port pinout.

The following table shows the pinout assignments for the ARX-1500 and the ARX-2500.

Table 1.16 *Pinout assignments for theARX-1500 and the ARX-2500*

Signal Type	RJ45 Pin Numbers
RTS	1
DTR	2
TD	3
SGND	4
RI	5
RD	6
DSR	7
CTS	8

SFP Optical Connector for the ARX-2500

The Gigabit Ethernet optical ports on the ARX-2500 use small form-factor pluggable (SFP) optical transceivers that accept LC-style multi-mode fiber connectors. These are for connection to Ethernet over fiber-optic cable.
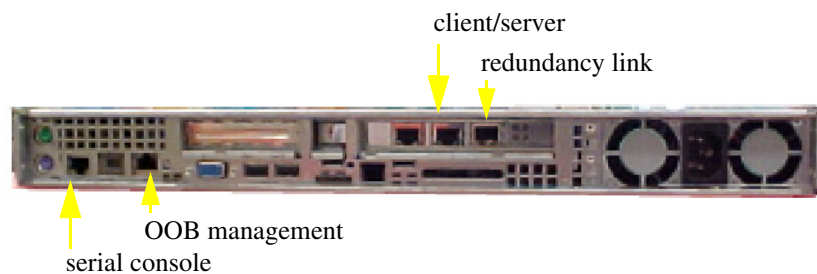
◆ **Important**

Fiber-optic ports are shipped with SFP optics installed. These ports must be protected by a rubber grommet filler or a cable connector at all times to prevent dust from collecting in the transceiver.

ARX-500 Connectors

The ARX-500 has four rear-panel ports for cabling, as shown in *Figure 1.15*:

Figure 1.15 *ARX-500 rear-panel ports.*

The following table describes the cable connectors used on the ARX-500.

Table 1.17 *ARX-500 Cable Connectors*

Interface	Connector	Purpose
serial console (labeled "10101")	RJ-45	Serial port for CLI access
OOB management (labeled "2")	RJ-45	100/1000BASE-T Ethernet (copper) port for CLI access
client/server port (labeled "Port 1/1")	RJ-45	100/1000BASE-T Ethernet (copper) port for client/server traffic and inband management
redundancy-link port (labeled "Port 1/2")	RJ-45	100/1000BASE-T Ethernet (copper) port for dedicated connection to a redundant peer

ARX-1500 Connectors

The ARX-1500 is powered by two power supplies (1+1 redundancy) system) consisting of two power modules and one power system frame. Two power modules are recommended for full redundancy and load-sharing. The power supplies require a 10A / 220VAC input cord, which is provided with the chassis. The AC outlet to the switch must be properly grounded.

Table 1.18 *ARX-1500 Cable Connectors*

Interface	Connector	Purpose
Console (labeled CONSOLE on back of switch)	RJ-45	Serial port for CLI access
Management (labeled MGMT on back of switch)	RJ-45	Ethernet port for CLI or GUI access
Gigabit Ethernet (on front of switch)	RJ-45	100/1000BASE-T Ethernet ports for access to client/server networks

ARX-2000 Connectors

The following table describes the cable connectors used on the ARX-2000.

Table 1.19 *ARX-2000 Cable Connectors*

Interface	Connector	Purpose
Console (labeled CONSOLE on back of switch)	RJ-45	Serial port for CLI access
Management (labeled MGMT on back of switch)	RJ-45	Ethernet port for CLI or GUI access
Gigabit Ethernet (on front of switch)	RJ-45	100/1000BASE-T Ethernet ports for access to client/server networks

ARX-2500 Connectors

The ARX-2500 is powered by two power supplies (1+1 redundancy) consisting of two power modules and one power system frame. Two power modules are recommended for full redundancy and load-sharing. The power supplies require a 10A / 220VAC input cord, which is provided with the chassis. The AC outlet to the switch must be properly grounded.

Table 1.20 ARX-2500 Cable Connectors

Interface	Connector	Purpose
Console (labeled CONSOLE on back of switch)	RJ-45	Serial port for CLI access
Management (labeled MGMT on back of switch)	RJ-45	Ethernet port for CLI or GUI access
Gigabit Ethernet:		
Optical ports	small form-factor Pluggable (SFP)	Two optical ports for 10-Gbps Ethernet connections over multi mode fiber.
Copper ports	RJ-45	Seven 100/1000 Base-T Ethernet ports.

ARX-4000 Connectors

The following table describes the cable connectors used on the ARX-4000.

Table 1.21 ARX-4000 Cable Connectors

Interface	Connector	Purpose
Serial console port	RJ-45	Serial console port for CLI access. Requires a serial rollover cable (RJ-45 to RJ-45) that is included in the Accessory Kit. This cable is sufficient for connecting to a terminal server. Do not use an Ethernet cable. For a direct connection to the serial port on a management station (such as a laptop), an RJ-45 to serial DB9 adapter is also included in the Accessory Kit.
OOB management port	RJ-45	Ethernet port for CLI or GUI access.
Gigabit Ethernet copper ports	RJ-45	100/1000BASE-T Ethernet ports.

Table 1.21 ARX-4000 Cable Connectors (Continued)

Interface	Connector	Purpose
10-gigabit Ethernet optical ports	X2 MSA form factor	Optical ports (2) for 10-gigabit Ethernet connections over multi-mode fiber. Shipped with small form-factor pluggable (SFP) optical transceivers installed. Transceivers accept SC-style multi-mode fiber connectors for connection to Ethernet over fiber-optic cable.
	◆ **Important**	
	Ports must be protected by blank covers, a rubber grommet filler, or a cable connector at all times to prevent dust from collecting in the transceiver.	

Bringing an ARX Inline

The process of bringing an ARX inline begins with installation. Consult the hardware installation guides for each ARX model for the tasks involved with installation.

Once you reach the installation *initial interview*, you can access the ARX through the serial console to configure a default administrator, a switch identifier, and out-of-band (OOB) management. Again, consult the hardware installation guides for each ARX model for the tasks involved.

If you have used Data Manager to discover the configuration information of a file server, you can move on to use Data Manager to generate an ARX configuration definition (script) that can be installed and run on the ARX.

Creating a Configuration Definition Using Data Manager

Using file server discovery information, the Data Manager Create ARX Configuration wizard lets you create a configuration definition that is subsequently stored in the Data Manager database. From this definition, you can generate the automated DOS and ARX CLI command scripts and step-by-step deployment instructions that constitute an ARX deployment workflow.

When implemented, the result is a virtualized NAS environment, enabling you to effectively manage unstructured file data without disrupting systems, applications, or users.

ARX configuration definitions generated by Data Manager are intended for initial switch deployments. Do not merge these configurations with previously deployed configurations without thoroughly reviewing both

existing and new configurations for naming collisions and hardware limitations. For further details on some of these issues, see *Manually Preparing the Back-End File Servers*, on page 1-15.

Accessing the ARX Using the CLI or the GUI

The process of configuring an ARX begins with the initial interview. During the initial interview, the ARX is accessed through the serial console to configure a default administrator, a switch identifier, and out-of-band (OOB) management. Consult the hardware installation guides for each ARX model for the tasks involved with each of these tasks.

After the initial interview, you must complete the network configuration phase, which includes site preparation and planning.

Finally, you must complete the storage configuration (or global configuration) phase. During the storage configuration phase, the storage-specific configuration and the associated virtualization scripts are created.

Among other tasks, this phase includes:

- Designing the namespace storage and its volumes
- Configuring policies
- Setting up virtual servers

Data Manager is not involved during the network configuration phase but can provide significant support during the storage configuration phase. However, if you choose not to use Data Manager, there are two other ways to access an ARX after installation: through the ARX Manager (GUI) or through the CLI. The following briefly describes each of these methods:

- ARX Manager Graphical User Interface (GUI).

  After hardware installation, add the ARX to your network using the *Quick Start: Network Setup Use Case*.

  Next, add CIFS and/or NFS services using the *Quick Start: CIFS Storage Use Case* or the *Quick Start: NFS Storage Use Case.* Use these documents to quickly aggregate your file server storage with the ARX Manager. Using the instructions in these quick starts, you can create a namespace, one managed volume, connect the volume to multiple file server shares, then offer the single volume (which aggregates your share storage) to clients as a single, virtual-service share.

  You can access these guides through the ARX GUI. A link to the documentation appears at the bottom of the navigation panel on the left-hand side of the GUI and across the top of the GUI.

- ARX CLI.

  After hardware installation, add the ARX to your network using the *ARX CLI Network-Management Guide*.

This manual contains instructions to set up and maintain networking and administration on a new ARX. After installing the switch, setting up its management IP, and preparing the switch for CLI and/or GUI provisioning, you can follow the order of the chapters in this manual to:

1. Set up administrative users and groups

2. Configure layer-2

3. Configure layer-3 (the network layer)

4. Join with other switches to form a RON

5. (Possibly) join with another switch to form a redundant pair

The final chapter has instructions to configure security for management access (for example, to allow other users to access the CLI and/or the GUI).

Next, you can configure your storage environment using the *ARX CLI Storage-Management Guide*. Consult the guide for the tasks involved in configuring the storage environment.

Best Practice: Regularly Saving the Configuration

We recommend copying the latest ARX configuration to a remote host every time the configuration changes. You can use the configuration parameters later to recover from a disaster and/or to facilitate switch replacement.

Every ARX keeps its local network parameters in its running-config, and every redundant pair of ARX devices share a single global-config with all namespace and service parameters. An ARX's startup-config contains both its running-config and global-config in a single file. Therefore, a redundant pair requires two or more saved configuration files:

*   two startup-configs (one per ARX), or

*   two running-configs (one per ARX) and a single global-config

In addition, to prepare for replacing a single switch in the unlikely event of a failure, you should also save the following:

*   Master key (extracted and wrapped). Save it to a remote host along with the configs.

*   Master key wrapping key password. Save it to a secure location.

To save configurations regularly from the CLI:

Use the following commands to save each configuration to a remote site:

*   copy running-config

*   copy global-config

*   copy startup-config

For detailed instructions on each of these commands, consult the *ARX CLI Reference*.

To save configurations regularly from the GUI:

1. From the left-hand navigation panel, expand **Maintenance**.

2. Click **Configs**.

3. On the Configs page, click **Save Config**.

4. On the Save Configuration page, click the **Help** button in the upper right-hand corner and follow the instructions.

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/