

BIG-IP[®] Device Service Clustering: Administration

Version 11.5



Table of Contents

Legal Notices.....	7
Acknowledgments.....	9
Chapter 1: Introducing BIG-IP Device Service Clustering.....	21
What is BIG-IP device service clustering?.....	21
DSC components.....	21
Best practices.....	22
Chapter 2: Creating an Active-Standby Configuration using the Configuration Utility.....	23
Utility.....	23
Overview: Creating an active-standby DSC configuration.....	23
About DSC configuration on a VIPRION system.....	23
DSC prerequisite worksheet.....	25
Task summary.....	26
Specifying an IP address for config sync.....	27
Specifying an IP address for connection mirroring.....	27
Specifying the HA capacity of a device.....	28
Establishing device trust.....	28
Creating a Sync-Failover device group.....	29
Syncing the BIG-IP configuration to the device group.....	30
Specifying IP addresses for failover communication.....	31
Syncing the BIG-IP configuration to the device group.....	32
Implementation result.....	33
Chapter 3: Creating an Active-Active Configuration using the Configuration Utility.....	35
Overview: Creating an active-active DSC configuration.....	35
About DSC configuration on a VIPRION system.....	36
DSC prerequisite worksheet.....	37
Configurations using Sync-Failover device groups.....	38
Task summary.....	39
Specifying an IP address for config sync.....	39
Specifying an IP address for connection mirroring.....	40
Specifying the HA capacity of a device.....	41
Establishing device trust.....	41
Creating a Sync-Failover device group.....	42
Syncing the BIG-IP configuration to the device group.....	43
Specifying IP addresses for failover communication.....	43
Creating a second traffic group for the device group.....	44
Assigning traffic-group-2 to a floating virtual IP address.....	45
Assigning traffic-group-2 to a floating self IP address.....	45

Syncing the BIG-IP configuration to the device group.....	46
Forcing a traffic group to a standby state.....	46
Implementation result.....	47
Chapter 4: Working with DSC Devices.....	49
About IP addresses for config sync, failover, and mirroring.....	49
About device properties.....	50
Viewing device properties.....	50
Specifying values for device properties.....	50
Device properties.....	51
About device status.....	51
Viewing possible status types for a device.....	52
Viewing the status of a device.....	52
Device status.....	52
Chapter 5: Managing Device Trust.....	55
What is device trust?.....	55
Types of trust authority.....	55
Device identity.....	56
Device discovery in a local trust domain.....	56
Establishing device trust.....	57
Adding a device to the local trust domain.....	58
Managing trust authority for a device.....	58
Chapter 6: Working with Device Groups.....	61
About Sync-Failover device groups.....	61
Sample Sync-Failover configuration.....	62
Sync-Failover device group considerations.....	62
Creating a Sync-Failover device group.....	63
About Sync-Only device groups.....	64
Sample Sync-Only configuration.....	64
Creating a Sync-Only device group.....	65
Viewing a list of device groups.....	66
Viewing the members of a device group.....	66
Adding a device to a device group.....	66
A note about folders and overlapping device groups.....	67
Chapter 7: Managing Configuration Synchronization.....	69
About configuration synchronization.....	69
About automatic and manual sync.....	69
Enabling and disabling automatic sync.....	70
Manually synchronizing the BIG-IP configuration.....	71
About full and incremental sync.....	71

Enabling and disabling full sync.....	72
Specifying an IP address for config sync.....	72
Viewing config sync status for the local device.....	73
Viewing config sync status for all device groups and members.....	73
Troubleshooting the config sync process.....	74
Sync status for device groups.....	74
Sync status for device group members.....	76
Advanced config sync properties for a device.....	78
Chapter 8: Managing Failover.....	79
Introduction to failover.....	79
What triggers failover?.....	79
About IP addresses for failover.....	80
Specifying IP addresses for failover communication.....	80
Failover addresses for appliance models.....	81
Failover addresses for VIPRION models.....	82
About traffic groups.....	83
Failover objects and traffic group association.....	83
Pre-configured traffic groups.....	84
Before you configure a traffic group.....	84
Creating a traffic group.....	84
Adding members to a traffic group.....	85
Viewing a list of traffic groups for a device.....	86
Viewing the members of a traffic group.....	86
Traffic group properties.....	86
Active and standby states.....	87
Viewing the failover state of a device.....	88
Viewing the state of a traffic group.....	88
Forcing a traffic group to a standby state.....	88
About active-standby vs. active-active configurations.....	89
Description of current and next-active devices.....	89
About the next-active device.....	89
What is load-aware failover?.....	90
What is an HA ordered list?.....	96
What is an HA group?.....	97
About MAC masquerade addresses.....	101
Chapter 9: Managing Connection Mirroring.....	103
About connection and persistence mirroring.....	103
Connection mirroring considerations.....	103
Configuration task summary.....	104
Specifying an IP address for connection mirroring.....	104
Enabling connection mirroring for TCP and UDP connections.....	105
Enabling connection mirroring for SNAT connections.....	105

Enabling mirroring of persistence records.....106

Chapter 10: Working with Folders.....107

 About folders on the BIG-IP system.....107

 About folder attributes for redundancy.....107

 About the root folder.....108

 Viewing redundancy attributes for the root folder.....108

 Configuring the traffic group attribute for the root folder.....108

Appendix A: Summary of tmsh Troubleshooting Tools.....111

 Summary of tmsh troubleshooting tools.....111

Legal Notices

Publication Date

This document was published on January 15, 2016.

Publication Number

MAN-0375-05

Copyright

Copyright © 2013-2016, F5 Networks, Inc. All rights reserved.

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

Trademarks

AAM, Access Policy Manager, Advanced Client Authentication, Advanced Firewall Manager, Advanced Routing, AFM, APM, Application Acceleration Manager, Application Security Manager, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Certified [DESIGN], F5 Networks, F5 SalesXchange [DESIGN], F5 Synthesis, f5 Synthesis, F5 Synthesis [DESIGN], F5 TechXchange [DESIGN], Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, iApps, IBR, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, Local Traffic Manager, LTM, LineRate, LineRate Systems [DESIGN], LROS, LTM, Message Security Manager, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Protocol Security Manager, PSM, Real Traffic Policy Builder, SalesXchange, ScaleN, Signalling Delivery Controller, SDC, SSL Acceleration, software designed applications services, SDAC (except in Japan), StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TechXchange, TMOS, TotALL, Traffic Management Operating System, Traffix Systems, Traffix Systems (DESIGN), Transparent Data Reduction, UNITY, VAULT, vCMP, VE F5 [DESIGN], Versafe, Versafe [DESIGN], VIPRION, Virtual Clustered Multiprocessing, WebSafe, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

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

Patents

This product may be protected by one or more patents indicated at:

<http://www.f5.com/about/guidelines-policies/patents>

Export Regulation Notice

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

RF Interference Warning

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

FCC Compliance

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

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

Canadian Regulatory Compliance

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

Standards Compliance

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

Acknowledgments

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

This product includes software developed by the NetBSD Foundation, Inc. and its contributors.

This product includes software developed by the Politecnico di Torino, and its contributors.

This product includes software developed by the Swedish Institute of Computer Science and its contributors.

This product includes software developed by the University of California, Berkeley and its contributors.

This product includes software developed by the Computer Systems Engineering Group at the Lawrence Berkeley Laboratory.

This product includes software developed by Christopher G. Demetriou for the NetBSD Project.

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

This product includes software developed by Jochen Pohl for the NetBSD Project.

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

This product includes software developed by SigmaSoft, Th. Lockert.

This product includes software developed for the NetBSD Project by Jason R. Thorpe.

This product includes software developed by Jason R. Thorpe for And Communications, <http://www.and.com>.

This product includes software developed for the NetBSD Project by Frank Van der Linden.

This product includes software developed for the NetBSD Project by John M. Vinopal.

This product includes software developed by Christos Zoulas.

This product includes software developed by the University of Vermont and State Agricultural College and Garrett A. Wollman.

This product includes software developed by Balazs Scheidler (bazsi@balabit.hu), which is protected under the GNU Public License.

This product includes software developed by Niels Mueller (nisse@lysator.liu.se), which is protected under the GNU Public License.

Acknowledgments

In the following statement, *This software* refers to the Mitsumi CD-ROM driver: This software was developed by Holger Veit and Brian Moore for use with 386BSD and similar operating systems. *Similar operating systems* includes mainly non-profit oriented systems for research and education, including but not restricted to NetBSD, FreeBSD, Mach (by CMU).

This product includes software developed by the Apache Group for use in the Apache HTTP server project (<http://www.apache.org/>).

This product includes software licensed from Richard H. Porter under the GNU Library General Public License (© 1998, Red Hat Software), www.gnu.org/copyleft/lgpl.html.

This product includes the standard version of Perl software licensed under the Perl Artistic License (© 1997, 1998 Tom Christiansen and Nathan Torkington). All rights reserved. You may find the most current standard version of Perl at <http://www.perl.com>.

This product includes software developed by Jared Minch.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product contains software based on oprofile, which is protected under the GNU Public License.

This product includes software with glib library utility functions, which is protected under the GNU Public License.

This product includes software with grub2 bootloader functions, which is protected under the GNU Public License.

This product includes software with the Intel Gigabit Linux driver, which is protected under the GNU Public License. Copyright ©1999 - 2012 Intel Corporation.

This product includes software with the Intel 10 Gigabit PCI Express Linux driver, which is protected under the GNU Public License. Copyright ©1999 - 2012 Intel Corporation.

This product includes RRDtool software developed by Tobi Oetiker (<http://www.rrdtool.com/index.html>) and licensed under the GNU General Public License.

This product contains software licensed from Dr. Brian Gladman under the GNU General Public License (GPL).

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

This product includes Hypersonic SQL.

This product contains software developed by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, and others.

This product includes software developed by the Internet Software Consortium.

This product includes software developed by Nominum, Inc. (<http://www.nominum.com>).

This product contains software developed by Broadcom Corporation, which is protected under the GNU Public License.

This product contains software developed by MaxMind LLC, and is protected under the GNU Lesser General Public License, as published by the Free Software Foundation.

This product includes software developed by Andrew Tridgell, which is protected under the GNU Public License, copyright ©1992-2000.

This product includes software developed by Jeremy Allison, which is protected under the GNU Public License, copyright ©1998.

This product includes software developed by Guenther Deschner, which is protected under the GNU Public License, copyright ©2008.

This product includes software developed by www.samba.org, which is protected under the GNU Public License, copyright ©2007.

This product includes software from Allan Jardine, distributed under the MIT License.

This product includes software from Trent Richardson, distributed under the MIT License.

This product includes vmbus drivers distributed by Microsoft Corporation.

This product includes software from Cavium.

This product includes software from Webroot, Inc.

This product includes software from Maxmind, Inc.

This product includes software from OpenVision Technologies, Inc. Copyright ©1993-1996, OpenVision Technologies, Inc. All Rights Reserved.

This product includes software developed by Matt Johnson, distributed under the MIT License. Copyright ©2012.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software from NLnetLabs. Copyright ©2001-2006. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of NLnetLabs nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes GRand Unified Bootloader (GRUB) software developed under the GNU Public License, copyright ©2007.

Acknowledgments

This product includes Intel QuickAssist kernel module, library, and headers software licensed under the GNU General Public License (GPL).

This product includes gd-libgd library software developed by the following in accordance with the following copyrights:

- Portions copyright ©1994, 1995, 1996, 1997, 1998, 2000, 2001, 2002 by Cold Spring Harbor Laboratory. Funded under Grant P41-RR02188 by the National Institutes of Health.
- Portions copyright ©1996, 1997, 1998, 1999, 2000, 2001, 2002 by Boutell.Com, Inc.
- Portions relating to GD2 format copyright ©1999, 2000, 2001, 2002 Philip Warner.
- Portions relating to PNG copyright ©1999, 2000, 2001, 2002 Greg Roelofs.
- Portions relating to gdtf.c copyright ©1999, 2000, 2001, 2002 John Ellson (ellson@lucent.com).
- Portions relating to gdft.c copyright ©2001, 2002 John Ellson (ellson@lucent.com).
- Portions copyright ©2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007 2008 Pierre-Alain Joye (pierre@libgd.org).
- Portions relating to JPEG and to color quantization copyright ©2000, 2001, 2002, Doug Becker and copyright ©1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001, 2002, Thomas G. Lane. This software is based in part on the work of the Independent JPEG Group.
- Portions relating to WBMP copyright 2000, 2001, 2002 Maurice Szmurlo and Johan Van den Brande. Permission has been granted to copy, distribute and modify gd in any context without fee, including a commercial application, provided that this notice is present in user-accessible supporting documentation.

This product includes software developed by Oracle America, Inc. Copyright ©2012.

1. Java Technology Restrictions. Licensee shall not create, modify, change the behavior of, or authorize licensees of licensee to create, modify, or change the behavior of, classes, interfaces, or subpackages that are in any way identified as "java", "javax", "sun" or similar convention as specified by Oracle in any naming convention designation. In the event that Licensee creates an additional API(s) which: (a) extends the functionality of a Java Environment; and (b) is exposed to third party software developers for the purpose of developing additional software which invokes such additional API, Licensee must promptly publish broadly an accurate specification for such API for free use by all developer.
2. Trademarks and Logos. This License does not authorize an end user licensee to use any Oracle America, Inc. name, trademark, service mark, logo or icon. The end user licensee acknowledges that Oracle owns the Java trademark and all Java-related trademarks, logos and icon including the Coffee Cup and Duke ("Java Marks") and agrees to: (a) comply with the Java Trademark Guidelines at <http://www.oracle.com/html/3party.html>; (b) not do anything harmful to or inconsistent with Oracle's rights in the Java Marks; and (c) assist Oracle in protecting those rights, including assigning to Oracle any rights acquired by Licensee in any Java Mark.
3. Source Code. Software may contain source code that, unless expressly licensed for other purposes, is provided solely for reference purposes pursuant to the terms of your license. Source code may not be redistributed unless expressly provided for in the terms of your license.
4. Third Party Code. Additional copyright notices and license terms applicable to portion of the Software are set forth in the THIRDPARTYLICENSEREADME.txt file.
5. Commercial Features. Use of the Commercial Features for any commercial or production purpose requires a separate license from Oracle. "Commercial Features" means those features identified in Table I-I (Commercial Features In Java SE Product Editions) of the Software documentation accessible at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>.

This product includes utilities developed by Linus Torvalds for inspecting devices connected to a USB bus.

This product includes perl-PHP-Serialization software, developed by Jesse Brown, copyright ©2003, and distributed under the Perl Development Artistic License (<http://dev.perl.org/licenses/artistic.html>).

This product includes software developed by members of the CentOS Project under the GNU Public License, copyright ©2004-2011 by the CentOS Project.

This product includes software licensed from Gerald Combs (gerald@wireshark.org) under the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or any later version. Copyright ©1998 Gerald Combs.

This product includes software licensed from Rémi Denis-Courmont under the GNU Library General Public License. Copyright ©2006 - 2011.

This product includes software developed by jQuery Foundation and other contributors, distributed under the MIT License. Copyright ©2014 jQuery Foundation and other contributors (<http://jquery.com/>).

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software developed by Trent Richardson, distributed under the MIT License. Copyright ©2012 jQuery Foundation and other contributors (<http://jquery.com/>).

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software developed by Allan Jardine, distributed under the MIT License. Copyright ©2008 - 2012, Allan Jardine, all rights reserved, jQuery Foundation and other contributors (<http://jquery.com/>).

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,

OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software developed by Douglas Gilbert. Copyright ©1992 - 2012 The FreeBSD Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE FREEBSD PROJECT ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FREEBSD PROJECT OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The views and conclusions contained in the software and documentation are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of the FreeBSD Project.

This product includes software developed as open source software. Copyright ©1994 - 2012 The FreeBSD Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). Copyright ©1998 - 2011 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes software licensed from William Ferrell, Selene Scriven and many other contributors under the GNU General Public License, copyright ©1998 - 2006.

This product includes software developed by Thomas Williams and Colin Kelley. Copyright ©1986 - 1993, 1998, 2004, 2007

Permission to use, copy, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Permission to modify the software is granted, but not the right to distribute the complete modified source code. Modifications are to be distributed as patches to the released version. Permission to distribute binaries produced by compiling modified sources is granted, provided you

1. distribute the corresponding source modifications from the released version in the form of a patch file along with the binaries,
2. add special version identification to distinguish your version in addition to the base release version number,
3. provide your name and address as the primary contact for the support of your modified version, and
4. retain our contact information in regard to use of the base software.

Permission to distribute the released version of the source code along with corresponding source modifications in the form of a patch file is granted with same provisions 2 through 4 for binary distributions. This software is provided "as is" without express or implied warranty to the extent permitted by applicable law.

This product includes software developed by the Computer Systems Engineering Group at Lawrence Berkeley Laboratory. Copyright ©1990-1994 Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: This product includes software developed by the Computer Systems Engineering Group at Lawrence Berkeley Laboratory.

Acknowledgments

4. Neither the name of the University nor of the Laboratory may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes software developed by Sony Computer Science Laboratories Inc. Copyright © 1997-2003 Sony Computer Science Laboratories Inc. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY SONY CSL AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL SONY CSL OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product contains software developed by Google, Inc. Copyright ©2011 Google, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software developed by Jeremy Ashkenas and DocumentCloud, and distributed under the MIT license. Copyright © 2010-2013 Jeremy Ashkenas, DocumentCloud.

This product includes gson software, distributed under the Apache License version 2.0. Copyright © 2008-2011 Google Inc.

This product includes the ixgbevf Intel Gigabit Linux driver, Copyright © 1999 - 2012 Intel Corporation, and distributed under the GPLv2 license, as published by the Free Software Foundation.

This product includes libwebp software. Copyright © 2010, Google Inc. All rights reserved.

This product includes Angular software developed by Google, Inc., <http://angularjs.org>, copyright © 2010-2012 Google, Inc., and distributed under the MIT license.

This product includes node.js software, copyright © Joyent, Inc. and other Node contributors. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product MAY include Intel SSD software subject to the following license; check your hardware specification for details.

1. **LICENSE.** This Software is licensed for use only in conjunction with Intel solid state drive (SSD) products. Use of the Software in conjunction with non-Intel SSD products is not licensed hereunder. Subject to the terms of this Agreement, Intel grants to You a nonexclusive, nontransferable, worldwide, fully paid-up license under Intel's copyrights to:

- copy the Software onto a single computer or multiple computers for Your personal, noncommercial use; and
- make appropriate back-up copies of the Software, for use in accordance with Section 1a) above.

The Software may contain the software or other property of third party suppliers, some of which may be identified in, and licensed in accordance with, any enclosed "license.txt" file or other text or file.

Except as expressly stated in this Agreement, no license or right is granted to You directly or by implication, inducement, estoppel or otherwise. Intel will have the right to inspect or have an independent auditor inspect Your relevant records to verify Your compliance with the terms and conditions of this Agreement.

2. **RESTRICTIONS.** You will not:

- a. copy, modify, rent, sell, distribute or transfer any part of the Software, and You agree to prevent unauthorized copying of the Software; and,
- b. reverse engineer, decompile, or disassemble the Software; and,
- c. sublicense or permit simultaneous use of the Software by more than one user; and,
- d. otherwise assign, sublicense, lease, or in any other way transfer or disclose Software to any third party, except as set forth herein; and,
- e. subject the Software, in whole or in part, to any license obligations of Open Source Software including without limitation combining or distributing the Software with Open Source Software in a manner that subjects the Software or any portion of the Software provided by Intel hereunder to any license obligations of such Open Source Software. "Open Source Software" means any software that requires as a condition of use, modification and/or distribution of such software that such software or other software incorporated into, derived from or distributed with such software:

Acknowledgments

- a. be disclosed or distributed in source code form; or
- b. be licensed by the user to third parties for the purpose of making and/or distributing derivative works; or
- c. be redistributable at no charge.

Open Source Software includes, without limitation, software licensed or distributed under any of the following licenses or distribution models, or licenses or distribution models substantially similar to any of the following:

- a. GNU's General Public License (GPL) or Lesser/Library GPL (LGPL),
- b. the Artistic License (e.g., PERL),
- c. the Mozilla Public License,
- d. the Netscape Public License,
- e. the Sun Community Source License (SCSL),
- f. vi) the Sun Industry Source License (SISL),
- g. vii) the Apache Software license, and
- h. viii) the Common Public License (CPL).

3. **OWNERSHIP OF SOFTWARE AND COPYRIGHTS.** Title to all copies of the Software remains with Intel or its suppliers. The Software is copyrighted and protected by the laws of the United States and other countries, and international treaty provisions. You may not remove any copyright notices from the Software. Intel may make changes to the Software, or to materials referenced therein, at any time and without notice, but is not obligated to support or update the Software. Except as otherwise expressly provided, Intel grants no express or implied right or license under Intel patents, copyrights, trademarks, or other intellectual property rights.
4. **Entire Agreement.** This Agreement contains the complete and exclusive statement of the agreement between You and Intel and supersedes all proposals, oral or written, and all other communications relating to the subject matter of this Agreement. Only a written instrument duly executed by authorized representatives of Intel and You may modify this Agreement.
5. **LIMITED MEDIA WARRANTY.** If the Software has been delivered by Intel on physical media, Intel warrants the media to be free from material physical defects for a period of ninety (90) days after delivery by Intel. If such a defect is found, return the media to Intel for replacement or alternate delivery of the Software as Intel may select.
6. **EXCLUSION OF OTHER WARRANTIES.** EXCEPT AS PROVIDED ABOVE, THE SOFTWARE IS PROVIDED "AS IS" WITHOUT ANY EXPRESS OR IMPLIED WARRANTY OF ANY KIND, INCLUDING WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT, OR FITNESS FOR A PARTICULAR PURPOSE. Intel does not warrant or assume responsibility for any errors, the accuracy or completeness of any information, text, graphics, links or other materials contained within the Software.
7. **LIMITATION OF LIABILITY.** IN NO EVENT WILL INTEL OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, LOST PROFITS, BUSINESS INTERRUPTION OR LOST INFORMATION) ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF INTEL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. SOME JURISDICTIONS PROHIBIT EXCLUSION OR LIMITATION OF LIABILITY FOR IMPLIED WARRANTIES OR CONSEQUENTIAL OR INCIDENTAL DAMAGES, SO THE ABOVE LIMITATION MAY NOT APPLY TO YOU. YOU MAY ALSO HAVE OTHER LEGAL RIGHTS THAT VARY FROM JURISDICTION TO JURISDICTION.
8. **TERMINATION OF THIS AGREEMENT.** Intel may terminate this Agreement at any time if You violate its terms. Upon termination, You will immediately destroy the Software or return all copies of the Software to Intel.
9. **APPLICABLE LAWS.** Claims arising under this Agreement will be governed by the laws of Delaware, excluding its principles of conflict of laws and the United Nations Convention on Contracts for the Sale

of Goods. You may not export the Software in violation of applicable export laws and regulations. Intel is not obligated under any other agreements unless they are in writing and signed by an authorized representative of Intel.

- 10. GOVERNMENT RESTRICTED RIGHTS.** The Software is provided with "RESTRICTED RIGHTS." Use, duplication, or disclosure by the Government is subject to restrictions as set forth in FAR52.227-14 and DFAR252.227-7013 et seq. or their successors. Use of the Software by the Government constitutes acknowledgment of Intel's proprietary rights therein. Contractor or Manufacturer is Intel Corporation, 2200 Mission College Blvd., Santa Clara, CA 95054.

Chapter 1

Introducing BIG-IP Device Service Clustering

- *What is BIG-IP device service clustering?*
- *DSC components*
- *Best practices*

What is BIG-IP device service clustering?

Device service clustering, or DSC[®], is an underlying architecture within BIG-IP[®] Traffic Management Operation System[®] (TMOS[®]). DSC provides synchronization and failover of BIG-IP configuration data at user-defined levels of granularity, among multiple BIG-IP devices on a network. More specifically, you can configure a BIG-IP device on a network to:

- Synchronize some or all of its configuration data among several BIG-IP devices
- Fail over to one of many available devices
- Mirror connections to a peer device to prevent interruption in service during failover

If you have two BIG-IP devices only, you can create either an active-standby or an active-active configuration. With more than two devices, you can create a configuration in which multiple devices are active and can fail over to one of many, if necessary.

By setting up DSC, you ensure that BIG-IP configuration objects are synchronized and can fail over at useful levels of granularity to the most-available BIG-IP devices on the network. You also ensure that failover from one device to another, when enabled, occurs seamlessly, with minimal to no interruption in application delivery.

The BIG-IP system supports either homogeneous or heterogeneous hardware platforms within a device group.

Important: *If you use the Setup utility to create a DSC configuration, you can re-enter the utility at any time to adjust the configuration. Simply click the F5 logo in the upper-left corner of the BIG-IP Configuration utility, and on the Welcome screen, click **Run Config Sync/HA Utility**.*

DSC components

Device service clustering (DSC[®]) is based on a few key components.

Devices

A *device* is a physical or virtual BIG-IP® system, as well as a member of a local trust domain and a device group. Each device member has a set of unique identification properties that the BIG-IP system generates. For device groups configured for failover, it is important that the device with the smallest capacity has the capacity to process all traffic groups. This ensures application availability in the event that all but one device in the device group become unavailable for any reason.

Device groups

A *device group* is a collection of BIG-IP devices that trust each other and can synchronize, and sometimes fail over, their BIG-IP configuration data. You can create two types of device groups: A *Sync-Failover* device group contains devices that synchronize configuration data and support traffic groups for failover purposes when a device becomes unavailable. A *Sync-Only* device group contains devices that synchronize configuration data, such as policy data, but do not synchronize failover objects. The BIG-IP system supports either homogeneous or heterogeneous hardware platforms within a device group.

Important: *BIG-IP module provisioning must be equivalent on all devices within a device group. For example, module provisioning is equivalent when all device group members are provisioned to run BIG-IP® Local Traffic Manager™ (LTM®) and BIG-IP® Application Security Manager™ (ASM™) only. Maintaining equivalent module provisioning on all devices ensures that any device in the device group can process module-specific application traffic in the event of failover from another device.*

Traffic groups

A *traffic group* is a collection of related configuration objects (such as a virtual IP address and a self IP address) that run on a BIG-IP device and process a particular type of application traffic. When a BIG-IP device becomes unavailable, a traffic group can float to another device in a device group to ensure that application traffic continues to be processed with little to no interruption in service.

Device trust and trust domains

Underlying the success of device groups and traffic groups is a feature known as device trust. *Device trust* establishes trust relationships between BIG-IP devices on the network, through mutual certificate-based authentication. A *trust domain* is a collection of BIG-IP devices that trust one another and can therefore synchronize and fail over their BIG-IP configuration data, as well as exchange status and failover messages on a regular basis. A *local trust domain* is a trust domain that includes the local device, that is, the device you are currently logged in to.

Folders

Folders are containers for the configuration objects on a BIG-IP device. For every administrative partition on the BIG-IP system, there is a high-level folder. At the highest level of the folder hierarchy is a folder named `root`. The BIG-IP system uses folders to affect the level of granularity to which it synchronizes configuration data to other devices in the device group.

Best practices

F5 recommends that you follow these best practices when deploying a DSC configuration:

Licensing and module provisioning

Verify that the licensing and the module provisioning is equivalent on all devices in the device group.

Traffic group load and device capacity

For a Sync-Failover device group, verify that the device with the smallest capacity has the capacity to process all traffic groups. This ensures application availability in the event that all but one device in the device group become unavailable for any reason.

Chapter 2

Creating an Active-Standby Configuration using the Configuration Utility

- *Overview: Creating an active-standby DSC configuration*
- *DSC prerequisite worksheet*
- *Task summary*
- *Implementation result*

Overview: Creating an active-standby DSC configuration

The most common TMOS[®] device service clustering (DSC[™]) implementation is an *active-standby* configuration, where a single traffic group is active on one of the devices in the device group and is in a standby state on a peer device. If failover occurs, the standby traffic group on the peer device becomes active and begins processing the application traffic.

To implement this DSC implementation, you can create a Sync-Failover device group. A Sync-Failover device group with two or more members and one traffic group provides configuration synchronization and device failover, and optionally, connection mirroring.

If the device with the active traffic group goes offline, the traffic group becomes active on a peer device, and application processing is handled by that device.

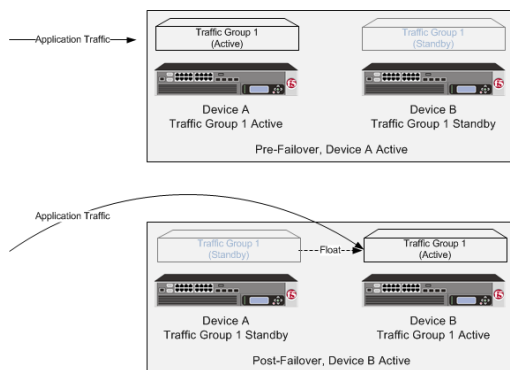


Figure 1: A two-member Sync-Failover device group for an active-standby configuration

About DSC configuration on a VIPRION system

The way you configure device service clustering (DSC[™]) (also known as redundancy) on a VIPRION[®] system varies depending on whether the system is provisioned to run the vCMP[®] feature.

For non-vCMP systems

For a device group that consists of VIPRION systems that are not licensed and provisioned for vCMP, each VIPRION cluster constitutes an individual device group member. The following table describes the IP addresses that you must specify when configuring redundancy.

Table 1: Required IP addresses for DSC configuration on a non-vCMP system

Feature	IP addresses required
Device trust	The primary floating management IP address for the VIPRION cluster.
ConfigSync	The unicast non-floating self IP address assigned to VLAN <code>internal</code> .
Failover	<ul style="list-style-type: none"> Recommended: The unicast non-floating self IP address that you assigned to an internal VLAN (preferably VLAN <code>HA</code>), as well as a multicast address. Alternative: All unicast management IP addresses that correspond to the slots in the VIPRION cluster.
Connection mirroring	For the primary address, the non-floating self IP address that you assigned to VLAN <code>HA</code> . The secondary address is not required, but you can specify any non-floating self IP address for an internal VLAN..

For vCMP systems

On a vCMP system, the devices in a device group are virtual devices, known as *vCMP guests*. You configure device trust, config sync, failover, and mirroring to occur between equivalent vCMP guests in separate chassis.

For example, if you have a pair of VIPRION systems running vCMP, and each system has three vCMP guests, you can create a separate device group for each pair of equivalent guests. Table 4.2 shows an example.

Table 2: Sample device groups for two VIPRION systems with vCMP

Device groups for vCMP	Device group members
Device-Group-A	<ul style="list-style-type: none"> Guest1 on chassis1 Guest1 on chassis2
Device-Group-B	<ul style="list-style-type: none"> Guest2 on chassis1 Guest2 on chassis2
Device-Group-C	<ul style="list-style-type: none"> Guest3 on chassis1 Guest3 on chassis2

By isolating guests into separate device groups, you ensure that each guest synchronizes and fails over to its equivalent guest. The following table describes the IP addresses that you must specify when configuring redundancy:

Table 3: Required IP addresses for DSC configuration on a VIPRION system with vCMP

Feature	IP addresses required
Device trust	The cluster management IP address of the guest.
ConfigSync	The non-floating self IP address on the guest that is associated with VLAN <code>internal</code> on the host.

Feature	IP addresses required
Failover	<ul style="list-style-type: none"> Recommended: The unicast non-floating self IP address on the guest that is associated with an internal VLAN on the host (preferably VLAN <code>HA</code>), as well as a multicast address. Alternative: The unicast management IP addresses for all slots configured for the guest.
Connection mirroring	For the primary address, the non-floating self IP address on the guest that is associated with VLAN <code>internal</code> on the host. The secondary address is not required, but you can specify any non-floating self IP address on the guest that is associated with an internal VLAN on the host.

DSC prerequisite worksheet

Before you set up device service clustering (DSC™), you must configure these BIG-IP® components on each device that you intend to include in the device group.

Table 4: DSC deployment worksheet

Configuration component	Considerations
Hardware, licensing, and provisioning	Devices in a device group must match with respect to product licensing and module provisioning. Heterogeneous hardware platforms within a device group are supported.
BIG-IP software version	Each device must be running BIG-IP version 11.x. This ensures successful configuration synchronization.
Management IP addresses	Each device must have a management IP address, a network mask, and a management route defined.
FQDN	Each device must have a fully-qualified domain name (FQDN) as its host name.
User name and password	Each device must have a user name and password defined on it that you will use when logging in to the BIG-IP Configuration utility.
<code>root</code> folder properties	The platform properties for the <code>root</code> folder must be set correctly (<code>Sync-Failover</code> and <code>traffic-group-1</code>).
VLANs	You must create these VLANs on each device, if you have not already done so: <ul style="list-style-type: none"> A VLAN for the internal network, named <code>internal</code> A VLAN for the external network, named <code>external</code> A VLAN for failover communications, named <code>HA</code>
Self IP addresses	You must create these self IP addresses on each device, if you have not already done so: <ul style="list-style-type: none"> Two self IP addresses (floating and non-floating) on the same subnet for VLAN <code>internal</code>. Two self IP addresses (floating and non-floating) on the same subnet for VLAN <code>external</code>. A non-floating self IP address on the internal subnet for VLAN <code>HA</code>.

Configuration component	Considerations
	<p>Note: When you create floating self IP addresses, the BIG-IP system automatically adds them to the default floating traffic group, <code>traffic-group-1</code>. To add a self IP address to a different traffic group, you must modify the value of the self IP address Traffic Group property.</p> <hr/> <p>Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the IP address you specify must be the floating IP address for high availability fast failover that you configured for the EC2 instance.</p> <hr/>
Port lockdown	For self IP addresses that you create on each device, you should verify that the Port Lockdown setting is set to Allow All , All Default , or Allow Custom . Do not specify None .
Application-related objects	You must create any virtual IP addresses and optionally, SNAT translation addresses, as part of the local traffic configuration. You must also configure any iApp™ application services if they are required for your application. When you create these addresses or services, the objects automatically become members of the default traffic group, <code>traffic-group-1</code> .
Time synchronization	The times set by the NTP service on all devices must be synchronized. This is a requirement for configuration synchronization to operate successfully.
Device certificates	Verify that each device includes an x509 device certificate. Devices with device certificates can authenticate and therefore trust one another, which is a prerequisite for device-to-device communication and data exchange.

Task summary

Use the tasks in this implementation to create a two-member device group, with one active traffic group, that syncs the BIG-IP® configuration to the peer device and provides failover capability if the peer device goes offline. Note that on a vCMP® system, the devices in a specific device group are vCMP guests, one per chassis.

Important: When you use this implementation, F5 Networks recommends that you synchronize the BIG-IP configuration twice, once after you create the device group, and again after you specify the IP addresses for failover.

Task list

- Specifying an IP address for config sync
- Specifying an IP address for connection mirroring
- Specifying the HA capacity of a device
- Establishing device trust
- Creating a Sync-Failover device group
- Syncing the BIG-IP configuration to the device group
- Specifying IP addresses for failover communication
- Syncing the BIG-IP configuration to the device group

Specifying an IP address for config sync

Before configuring the config sync address, verify that all devices in the device group are running the same version of BIG-IP® system software.

You perform this task to specify the IP address on the local device that other devices in the device group will use to synchronize their configuration objects to the local device.

Note: You must perform this task locally on each device in the device group.

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose ConfigSync.
5. For the **Local Address** setting, retain the displayed IP address or select another address from the list.
F5 Networks recommends that you use the default value, which is the self IP address for VLAN `internal`. This address must be a non-floating self IP address and not a management IP address.

Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the internal self IP address that you specify must be the internal private IP addresses that you configured for this EC2 instance as the **Local Address**.

6. Click **Update**.

After performing this task, the other devices in the device group can sync their configurations to the local device.

Specifying an IP address for connection mirroring

You can specify the local self IP address that you want other devices in a device group to use when mirroring their connections to this device. Connection mirroring ensures that in-process connections for an active traffic group are not dropped when failover occurs. You typically perform this task when you initially set up device service clustering (DSC®).

Note: You must perform this task locally on each device in the device group.

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Mirroring.
5. For the **Primary Local Mirror Address** setting, retain the displayed IP address or select another address from the list.

The recommended IP address is the self IP address for either VLAN `HA` or VLAN `internal`.

Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the self IP address you specify must be one of the private IP addresses that you configured for this EC2 instance as the **Primary Local Mirror Address**.

6. For the **Secondary Local Mirror Address** setting, retain the default value of **None**, or select an address from the list.
This setting is optional. The system uses the selected IP address in the event that the primary mirroring address becomes unavailable.
7. Click **Update**.

In addition to specifying an IP address for mirroring, you must also enable connection mirroring on the relevant virtual servers on this device.

Specifying the HA capacity of a device

Before you perform this task, verify that this device is a member of a device group and that the device group contains three or more devices.

You perform this task when you have more than one type of hardware platform in a device group and you want to configure load-aware failover. *Load-aware failover* ensures that the BIG-IP® system can intelligently select the next-active device for each active traffic group in the device group when failover occurs. As part of configuring load-aware failover, you define an HA capacity to establish the amount of computing resource that the device provides relative to other devices in the device group.

Note: If all devices in the device group are the same hardware platform, you can skip this task.

1. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
2. In the Name column, click the name of the device for which you want to view properties.
This displays a table of properties for the device.
3. In the **HA Capacity** field, type a relative numeric value.
You need to configure this setting only when you have varying types of hardware platforms in a device group and you want to configure load-aware failover. The value you specify represents the relative capacity of the device to process application traffic compared to the other devices in the device group.

Important: If you configure this setting, you must configure the setting on every device in the device group.

If this device has half the capacity of a second device and a third of the capacity of a third device in the device group, you can specify a value of 100 for this device, 200 for the second device, and 300 for the third device.

When choosing the next active device for a traffic group, the system considers the capacity that you specified for this device.

4. Click **Update**.

After you perform this task, the BIG-IP system uses the **HA Capacity** value to calculate the current utilization of the local device, to determine the next-active device for failover of other traffic groups in the device group.

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices A, B, and C each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device A and add devices B and C to the local trust domain. Note that there is no need to repeat this process on devices B and C.

1. On the Main tab, click **Device Management > Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.
3. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
 - If the BIG-IP device is a non-VIPRION® device, type the management IP address for the device.
 - If the BIG-IP device is a VIPRION device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
 - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
 - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
4. Click **Retrieve Device Information**.
5. Verify that the certificate of the remote device is correct.
6. Verify that the name of the remote device is correct.
7. Verify that the management IP address and name of the remote device are correct.
8. Click **Finished**.

The device you added is now a member of the local trust domain.

Repeat this task for each device that you want to add to the local trust domain.

Creating a Sync-Failover device group

This task establishes failover capability between two or more BIG-IP® devices. If an active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

Repeat this task for each Sync-Failover device group that you want to create for your network configuration.

1. On the Main tab, click **Device Management > Device Groups**.
2. On the Device Groups list screen, click **Create**.
The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. From the **Configuration** list, select **Advanced**.
5. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Includes** list.

The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only.

6. For the **Network Failover** setting, select or clear the check box:

- Select the check box if you want device group members to handle failover communications by way of network connectivity. This choice is required for active-active configurations.
- Clear the check box if you want device group members to handle failover communications by way of serial cable (hard-wired) connectivity.

For active-active configurations, you must select network failover, as opposed to serial-cable (hard-wired) connectivity.

7. For the **Automatic Sync** setting, select or clear the check box:

- Select the check box when you want the BIG-IP system to automatically sync the BIG-IP configuration data whenever a config sync operation is required. In this case, the BIG-IP system syncs the configuration data whenever the data changes on any device in the device group.
- Clear the check box when you want to manually initiate each config sync operation. In this case, F5 networks recommends that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.

8. For the **Full Sync** setting, select or clear the check box:

- Select the check box when you want all sync operations to be full syncs. In this case, the BIG-IP system syncs the entire set of BIG-IP configuration data whenever a config sync operation is required.
- Clear the check box when you want all sync operations to be incremental (the default setting). In this case, the BIG-IP system syncs only the changes that are more recent than those on the target device. When you select this option, the BIG-IP system compares the configuration data on each target device with the configuration data on the source device and then syncs the delta of each target-source pair.

If you enable incremental synchronization, the BIG-IP system might occasionally perform a full sync for internal reasons. This is a rare occurrence and no user intervention is required.

9. In the **Maximum Incremental Sync Size (KB)** field, retain the default value of 1024, or type a different value.

This value specifies the total size of configuration changes that can reside in the incremental sync cache. If the total size of the configuration changes in the cache exceeds the specified value, the BIG-IP system performs a full sync whenever the next config sync operation occurs.

10. Click **Finished**.

You now have a Sync-Failover type of device group containing BIG-IP devices as members.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust is established.

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: *You perform this task on either of the two devices, but not both.*

1. On the Main tab, click **Device Management > Overview**.

2. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

Except for non-floating self IP addresses, the entire set of BIG-IP configuration data is replicated on each device in the device group.

Specifying IP addresses for failover communication

You typically perform this task during initial Device Service Clustering (DSC®) configuration, to specify the local IP addresses that you want other devices in the device group to use for continuous health-assessment communication with the local device or guest. You must perform this task locally on each device in the device group.

Important: *If the system is running vCMP, you must log in to each guest to perform this task.*

Note: *The IP addresses that you specify must belong to route domain 0.*

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Failover.
5. For the Failover Unicast Configuration settings, click **Add** for each IP address on this device that other devices in the device group can use to exchange failover messages with this device. The unicast IP addresses you specify depend on the type of device:

Platform	Action
Appliance without vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}) and the static management IP address currently assigned to the device.
Appliance with vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}) and the unique management IP address currently assigned to the guest.
VIPRION without vCMP®	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}). If you choose to specify unicast addresses only (and not a multicast address), you must also type the existing, static management IP addresses that you previously configured for all slots in the cluster. If you choose to specify one or more unicast addresses and a multicast address, then you do not need to specify the existing, per-slot static management IP addresses when configuring addresses for failover communication.
VIPRION with vCMP	Type a self IP address that is defined on the guest and associated with an internal VLAN on the host (preferably VLAN _{HA}). If you choose to specify unicast failover addresses only (and not a multicast address), you must also type the existing, virtual static management IP addresses that you previously configured for all slots

Platform	Action
	in the guest's virtual cluster. If you choose to specify one or more unicast addresses and a multicast address, you do not need to specify the existing, per-slot static and virtual management IP addresses when configuring addresses for failover communication.

Important: Failover addresses should always be static, not floating, IP addresses.

6. To enable the use of a failover multicast address on a VIPRION[®] platform (recommended), then for the **Use Failover Multicast Address** setting, select the **Enabled** check box.
7. If you enabled **Use Failover Multicast Address**, either accept the default **Address** and **Port** values, or specify values appropriate for the device.
If you revise the default **Address** and **Port** values, but then decide to revert to the default values, click **Reset Defaults**.
8. Click **Update**.

After you perform this task, other devices in the device group can send failover messages to the local device using the specified IP addresses.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust is established.

This task synchronizes the BIG-IP[®] configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: You perform this task on either of the two devices, but not both.

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

Except for non-floating self IP addresses, the entire set of BIG-IP configuration data is replicated on each device in the device group.

Implementation result

You now have a Sync-Failover device group set up with an active-standby DSC™ configuration. This configuration uses the default floating traffic group (named `traffic-group-1`), which contains the application-specific floating self IP and virtual IP addresses, and is initially configured to be active on one of the two devices. If the device with the active traffic group goes offline, the traffic group becomes active on the other device in the group, and application processing continues.

Chapter 3

Creating an Active-Active Configuration using the Configuration Utility

- *Overview: Creating an active-active DSC configuration*
- *DSC prerequisite worksheet*
- *Task summary*
- *Implementation result*

Overview: Creating an active-active DSC configuration

A common TMOS[®] device service clustering (DSC[™]) implementation is an active-standby configuration, where a single traffic group is active on one of the devices in the device group, and is in a standby state on a peer device. Alternatively however, you can create a second traffic group and activate that traffic group on a peer device. In this *active-active* configuration, the devices each process traffic for a different application simultaneously. If one of the devices in the device group goes offline, the traffic group that was active on that device fails over to a peer device. The result is that two traffic groups can become active on one device.

To implement this DSC implementation, you create a Sync-Failover device group. A Sync-Failover device group with two or more members provides configuration synchronization and device failover, and optionally, connection mirroring.

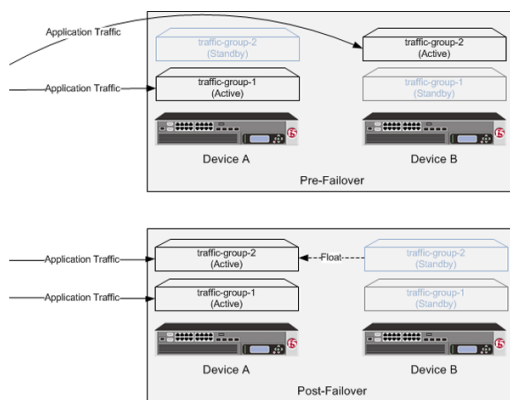


Figure 2: A two-member Sync-Failover group for an active-active configuration

About DSC configuration on a VIPRION system

The way you configure device service clustering (DSC™) (also known as redundancy) on a VIPRION® system varies depending on whether the system is provisioned to run the vCMP® feature.

For non-vCMP systems

For a device group that consists of VIPRION systems that are not licensed and provisioned for vCMP, each VIPRION cluster constitutes an individual device group member. The following table describes the IP addresses that you must specify when configuring redundancy.

Table 5: Required IP addresses for DSC configuration on a non-vCMP system

Feature	IP addresses required
Device trust	The primary floating management IP address for the VIPRION cluster.
ConfigSync	The unicast non-floating self IP address assigned to VLAN <i>internal</i> .
Failover	<ul style="list-style-type: none"> Recommended: The unicast non-floating self IP address that you assigned to an internal VLAN (preferably VLAN <i>HA</i>), as well as a multicast address. Alternative: All unicast management IP addresses that correspond to the slots in the VIPRION cluster.
Connection mirroring	For the primary address, the non-floating self IP address that you assigned to VLAN <i>HA</i> . The secondary address is not required, but you can specify any non-floating self IP address for an internal VLAN..

For vCMP systems

On a vCMP system, the devices in a device group are virtual devices, known as *vCMP guests*. You configure device trust, config sync, failover, and mirroring to occur between equivalent vCMP guests in separate chassis.

For example, if you have a pair of VIPRION systems running vCMP, and each system has three vCMP guests, you can create a separate device group for each pair of equivalent guests. Table 4.2 shows an example.

Table 6: Sample device groups for two VIPRION systems with vCMP

Device groups for vCMP	Device group members
Device-Group-A	<ul style="list-style-type: none"> Guest1 on chassis1 Guest1 on chassis2
Device-Group-B	<ul style="list-style-type: none"> Guest2 on chassis1 Guest2 on chassis2
Device-Group-C	<ul style="list-style-type: none"> Guest3 on chassis1 Guest3 on chassis2

By isolating guests into separate device groups, you ensure that each guest synchronizes and fails over to its equivalent guest. The following table describes the IP addresses that you must specify when configuring redundancy:

Table 7: Required IP addresses for DSC configuration on a VIPRION system with vCMP

Feature	IP addresses required
Device trust	The cluster management IP address of the guest.
ConfigSync	The non-floating self IP address on the guest that is associated with VLAN <code>internal</code> on the host.
Failover	<ul style="list-style-type: none"> Recommended: The unicast non-floating self IP address on the guest that is associated with an internal VLAN on the host (preferably VLAN <code>HA</code>), as well as a multicast address. Alternative: The unicast management IP addresses for all slots configured for the guest.
Connection mirroring	For the primary address, the non-floating self IP address on the guest that is associated with VLAN <code>internal</code> on the host. The secondary address is not required, but you can specify any non-floating self IP address on the guest that is associated with an internal VLAN on the host.

DSC prerequisite worksheet

Before you set up device service clustering (DSC™), you must configure these BIG-IP® components on each device that you intend to include in the device group.

Table 8: DSC deployment worksheet

Configuration component	Considerations
Hardware, licensing, and provisioning	Devices in a device group must match with respect to product licensing and module provisioning. Heterogeneous hardware platforms within a device group are supported.
BIG-IP software version	Each device must be running BIG-IP version 11.x. This ensures successful configuration synchronization.
Management IP addresses	Each device must have a management IP address, a network mask, and a management route defined.
FQDN	Each device must have a fully-qualified domain name (FQDN) as its host name.
User name and password	Each device must have a user name and password defined on it that you will use when logging in to the BIG-IP Configuration utility.
<code>root</code> folder properties	The platform properties for the <code>root</code> folder must be set correctly (<code>Sync-Failover</code> and <code>traffic-group-1</code>).
VLANs	<p>You must create these VLANs on each device, if you have not already done so:</p> <ul style="list-style-type: none"> A VLAN for the internal network, named <code>internal</code> A VLAN for the external network, named <code>external</code> A VLAN for failover communications, named <code>HA</code>

Configuration component	Considerations
Self IP addresses	<p>You must create these self IP addresses on each device, if you have not already done so:</p> <ul style="list-style-type: none"> • Two self IP addresses (floating and non-floating) on the same subnet for VLAN <code>internal</code>. • Two self IP addresses (floating and non-floating) on the same subnet for VLAN <code>external</code>. • A non-floating self IP address on the internal subnet for VLAN <code>HA</code>. <hr/> <p><i>Note:</i> When you create floating self IP addresses, the BIG-IP system automatically adds them to the default floating traffic group, <code>traffic-group-1</code>. To add a self IP address to a different traffic group, you must modify the value of the self IP address Traffic Group property.</p> <hr/> <p>Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the IP address you specify must be the floating IP address for high availability fast failover that you configured for the EC2 instance.</p>
Port lockdown	For self IP addresses that you create on each device, you should verify that the Port Lockdown setting is set to Allow All , All Default , or Allow Custom . Do not specify None .
Application-related objects	You must create any virtual IP addresses and optionally, SNAT translation addresses, as part of the local traffic configuration. You must also configure any iApp™ application services if they are required for your application. When you create these addresses or services, the objects automatically become members of the default traffic group, <code>traffic-group-1</code> .
Time synchronization	The times set by the NTP service on all devices must be synchronized. This is a requirement for configuration synchronization to operate successfully.
Device certificates	Verify that each device includes an x509 device certificate. Devices with device certificates can authenticate and therefore trust one another, which is a prerequisite for device-to-device communication and data exchange.

Configurations using Sync-Failover device groups

This illustration shows two separate Sync-Failover device groups. In the first device group, only **LTM1** processes application traffic, and the two BIG-IP devices are configured to provide active-standby high availability. This means that **LTM1** and **LTM2** synchronize their configurations, and the failover objects on **LTM1** float to **LTM2** if **LTM1** becomes unavailable.

In the second device group, both **LTM1** and **LTM2** process application traffic, and the BIG-IP devices are configured to provide active-active high availability. This means that **LTM1** and **LTM2** synchronize their configurations, the failover objects on **LTM1** float to **LTM2** if **LTM1** becomes unavailable, and the failover objects on **LTM2** float to **LTM1** if **LTM2** becomes unavailable.




Figure 3: Comparison of Active-Standby and Active-Active device groups

Task summary

Use the tasks in this implementation to create a two-member device group, with two active traffic groups, that syncs the BIG-IP® configuration to the peer device and provides failover capability if the peer device goes offline. Note that on a vCMP® system, the devices in a specific device group are vCMP guests, one per chassis.

Important: When you use this implementation, F5 Networks recommends that you synchronize the BIG-IP configuration twice, once after you create the device group, and again after you specify the IP addresses for failover.

Task list

- Specifying an IP address for config sync
- Specifying an IP address for connection mirroring
- Specifying the HA capacity of a device
- Establishing device trust
- Creating a Sync-Failover device group
- Syncing the BIG-IP configuration to the device group
- Specifying IP addresses for failover communication
- Creating a second traffic group for the device group
- Assigning traffic-group-2 to a floating virtual IP address
- Assigning traffic-group-2 to a floating self IP address
- Syncing the BIG-IP configuration to the device group
- Forcing a traffic group to a standby state

Specifying an IP address for config sync

Before configuring the config sync address, verify that all devices in the device group are running the same version of BIG-IP® system software.

You perform this task to specify the IP address on the local device that other devices in the device group will use to synchronize their configuration objects to the local device.

Note: You must perform this task locally on each device in the device group.

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose ConfigSync.
5. For the **Local Address** setting, retain the displayed IP address or select another address from the list.
F5 Networks recommends that you use the default value, which is the self IP address for VLAN `internal`. This address must be a non-floating self IP address and not a management IP address.

Important: *If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the internal self IP address that you specify must be the internal private IP addresses that you configured for this EC2 instance as the **Local Address**.*

6. Click **Update**.

After performing this task, the other devices in the device group can sync their configurations to the local device.

Specifying an IP address for connection mirroring

You can specify the local self IP address that you want other devices in a device group to use when mirroring their connections to this device. Connection mirroring ensures that in-process connections for an active traffic group are not dropped when failover occurs. You typically perform this task when you initially set up device service clustering (DSC®).

Note: *You must perform this task locally on each device in the device group.*

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Mirroring.
5. For the **Primary Local Mirror Address** setting, retain the displayed IP address or select another address from the list.

The recommended IP address is the self IP address for either VLAN `HA` or VLAN `internal`.

Important: *If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the self IP address you specify must be one of the private IP addresses that you configured for this EC2 instance as the **Primary Local Mirror Address**.*

6. For the **Secondary Local Mirror Address** setting, retain the default value of **None**, or select an address from the list.
This setting is optional. The system uses the selected IP address in the event that the primary mirroring address becomes unavailable.
7. Click **Update**.

In addition to specifying an IP address for mirroring, you must also enable connection mirroring on the relevant virtual servers on this device.

Specifying the HA capacity of a device

Before you perform this task, verify that this device is a member of a device group and that the device group contains three or more devices.

You perform this task when you have more than one type of hardware platform in a device group and you want to configure load-aware failover. *Load-aware failover* ensures that the BIG-IP® system can intelligently select the next-active device for each active traffic group in the device group when failover occurs. As part of configuring load-aware failover, you define an HA capacity to establish the amount of computing resource that the device provides relative to other devices in the device group.

Note: *If all devices in the device group are the same hardware platform, you can skip this task.*

1. On the Main tab, click **Device Management > Devices**.

This displays a list of device objects discovered by the local device.

2. In the Name column, click the name of the device for which you want to view properties.

This displays a table of properties for the device.

3. In the **HA Capacity** field, type a relative numeric value.

You need to configure this setting only when you have varying types of hardware platforms in a device group and you want to configure load-aware failover. The value you specify represents the relative capacity of the device to process application traffic compared to the other devices in the device group.

Important: *If you configure this setting, you must configure the setting on every device in the device group.*

If this device has half the capacity of a second device and a third of the capacity of a third device in the device group, you can specify a value of 100 for this device, 200 for the second device, and 300 for the third device.

When choosing the next active device for a traffic group, the system considers the capacity that you specified for this device.

4. Click **Update**.

After you perform this task, the BIG-IP system uses the **HA Capacity** value to calculate the current utilization of the local device, to determine the next-active device for failover of other traffic groups in the device group.

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices A, B, and C each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device A and add devices B and C to the local trust domain. Note that there is no need to repeat this process on devices B and C.

1. On the Main tab, click **Device Management > Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.
3. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
 - If the BIG-IP device is a non-VIPRION® device, type the management IP address for the device.
 - If the BIG-IP device is a VIPRION device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
 - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
 - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
4. Click **Retrieve Device Information**.
5. Verify that the certificate of the remote device is correct.
6. Verify that the name of the remote device is correct.
7. Verify that the management IP address and name of the remote device are correct.
8. Click **Finished**.

The device you added is now a member of the local trust domain.

Repeat this task for each device that you want to add to the local trust domain.

Creating a Sync-Failover device group

This task establishes failover capability between two or more BIG-IP devices that you intend to run in an active-active configuration. If an active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

Repeat this task for each Sync-Failover device group that you want to create for your network configuration.

1. On the Main tab, click **Device Management > Device Groups**.
2. On the Device Groups list screen, click **Create**.
The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Includes** list.
The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only.
5. For the **Network Failover** setting, verify that network failover is enabled.
Network failover must be enabled for active-active configurations (that is, device groups that will contain two or more active traffic groups).
6. Click **Finished**.

You now have a Sync-Failover type of device group containing BIG-IP devices as members. This device group is configured for environments that require the use of two or more active traffic groups to process application traffic.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust is established.

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: *You perform this task on either of the two devices, but not both.*

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

Except for non-floating self IP addresses, the entire set of BIG-IP configuration data is replicated on each device in the device group.

Specifying IP addresses for failover communication

You typically perform this task during initial Device Service Clustering (DSC®) configuration, to specify the local IP addresses that you want other devices in the device group to use for continuous health-assessment communication with the local device or guest. You must perform this task locally on each device in the device group.

Important: *If the system is running vCMP, you must log in to each guest to perform this task.*

Note: *The IP addresses that you specify must belong to route domain 0.*

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Failover.
5. For the Failover Unicast Configuration settings, click **Add** for each IP address on this device that other devices in the device group can use to exchange failover messages with this device. The unicast IP addresses you specify depend on the type of device:

Platform	Action
Appliance without vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN HA) and the static management IP address currently assigned to the device.

Platform	Action
Appliance with vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}) and the unique management IP address currently assigned to the guest.
VIPRION without vCMP[®]	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}). If you choose to specify unicast addresses only (and not a multicast address), you must also type the existing, static management IP addresses that you previously configured for all slots in the cluster. If you choose to specify one or more unicast addresses and a multicast address, then you do not need to specify the existing, per-slot static management IP addresses when configuring addresses for failover communication.
VIPRION with vCMP	Type a self IP address that is defined on the guest and associated with an internal VLAN on the host (preferably VLAN _{HA}). If you choose to specify unicast failover addresses only (and not a multicast address), you must also type the existing, virtual static management IP addresses that you previously configured for all slots in the guest's virtual cluster. If you choose to specify one or more unicast addresses and a multicast address, you do not need to specify the existing, per-slot static and virtual management IP addresses when configuring addresses for failover communication.

Important: Failover addresses should always be static, not floating, IP addresses.

- To enable the use of a failover multicast address on a VIPRION[®] platform (recommended), then for the **Use Failover Multicast Address** setting, select the **Enabled** check box.
- If you enabled **Use Failover Multicast Address**, either accept the default **Address** and **Port** values, or specify values appropriate for the device.
If you revise the default **Address** and **Port** values, but then decide to revert to the default values, click **Reset Defaults**.
- Click **Update**.

After you perform this task, other devices in the device group can send failover messages to the local device using the specified IP addresses.

Creating a second traffic group for the device group

This task creates a second active floating traffic group to process application traffic. The default floating traffic group (traffic-group-1) processes application traffic for the local device.

Note: For this implementation, name this traffic group **traffic-group-2**.

- On the Main tab, click **Device Management > Traffic Groups**.
- On the Traffic Group List screen, click **Create**.
- Type the name `traffic-group-2` for the new traffic group.
- In the **HA Load Factor** field, specify a value that represents the application load for this traffic group relative to other active traffic groups on the local device.

Important: If you configure this setting, you must configure the setting on every traffic group in the device group.

- In the **MAC Masquerade Address** field, type a MAC masquerade address.

When you specify a MAC masquerade address, you reduce the risk of dropped connections when failover occurs. This setting is optional.

6. Select or clear the check box for the **Auto Failback** option:
 - Select the check box to cause the traffic group, after failover, to fail over again to the first device in the traffic group's ordered list when that device (and only that device) is available.
 - Clear the check box to cause the traffic group, after failover, to remain active on its current device until failover occurs again.
7. For the **Failover Order** setting, in the **Available** box, select a device name and using the Move button, move the device name to the **Enabled** box. Repeat for each device that you want to include in the ordered list.

This setting is optional. Only devices that are members of the relevant Sync-Failover device group are available for inclusion in the ordered list. If you have enabled the auto-failback feature on the traffic group, ensure that the first device in the ordered list is the device to which you want this traffic group to fail back to when that first device becomes available.

If auto-failback is enabled and the first device in the **Failover Order** list is unavailable, no auto-failback occurs and the traffic group continues to run on the current device. Also, if none of the devices in the **Failover Order** list is currently available when failover occurs, the BIG-IP system ignores the **Failover Order** setting and performs load-aware failover instead, using the **HA Load Factor** setting.
8. Click **Finished**.

You now have a second floating traffic group on the local device (in addition to the default floating traffic group) so that once the traffic group is activated on the remote devices, devices in the device group can process traffic for different applications.

Assigning traffic-group-2 to a floating virtual IP address

This task assigns a floating traffic group to a virtual IP address on a device.

1. On the Main tab, click **Local Traffic > Virtual Servers > Virtual Address List**.
The Virtual Address List screen opens.
2. In the Name column, click the virtual address that you want to assign to the traffic group.
This displays the properties of that virtual address.
3. From the **Traffic Group** list, select **traffic-group-2 (floating)**.
4. Click **Update**.

The device's floating virtual IP address is now a member of your second traffic group. The virtual IP address can now fail over to other devices in the device group.

Assigning traffic-group-2 to a floating self IP address

This task assigns your floating self IP address to traffic-group-2.

1. On the Main tab, click **Network > Self IPs**.
The Self IPs screen opens.
2. In the Name column, click the floating self IP address assigned to VLAN `internal`.
This displays the properties of that self IP address.
3. From the **Traffic Group** list, select **traffic-group-2 (floating)**.
4. Click **Update**.

The device's floating self IP address is now a member of your second traffic group. The self IP address can now fail over to other devices in the traffic group.

Syncing the BIG-IP configuration to the device group

Before you sync the configuration, verify that the devices targeted for config sync are members of a device group and that device trust is established.

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: *You perform this task on either of the two devices, but not both.*

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of `Changes Pending`.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

Except for non-floating self IP addresses, the entire set of BIG-IP configuration data is replicated on each device in the device group.

Forcing a traffic group to a standby state

You perform this task when you want the selected traffic group on the local device to fail over to another device (that is, switch to a `Standby` state). Users typically perform this task when no automated method is configured for a traffic group, such as auto-failback or an HA group. By forcing the traffic group into a `Standby` state, the traffic group becomes active on another device in the device group. For device groups with more than two members, you can choose the specific device to which the traffic group fails over.

1. Log in to the device on which the traffic group is currently active.
2. On the Main tab, click **Device Management > Traffic Groups**.
3. In the Name column, locate the name of the traffic group that you want to run on the peer device.
4. Select the check box to the left of the traffic group name.
If the check box is unavailable, the traffic group is not active on the device to which you are currently logged in. Perform this task on the device on which the traffic group is active.
5. Click **Force to Standby**.
This displays target device options.
6. Choose one of these actions:
 - If the device group has two members only, click **Force to Standby**. This displays the list of traffic groups for the device group and causes the local device to appear in the Next Active Device column.

- If the device group has more than two members, then from the **Target Device** list, select a value and click **Force to Standby**.

The selected traffic group is now in a standby state on the local device and active on another device in the device group.

Implementation result

You now have a Sync-Failover device group set up with an active-active DSC™ configuration. In this configuration, each device has a different active traffic group running on it. That is, the active traffic group on one device is the default traffic group (named `traffic-group-1`), while the active traffic group on the peer device is a traffic group that you create. Each traffic group contains the floating self IP and virtual IP addresses specific to the relevant application.

If one device goes offline, the traffic group that was active on that device becomes active on the other device in the group, and processing for both applications continues on one device.

Chapter 4

Working with DSC Devices

- *About IP addresses for config sync, failover, and mirroring*
- *About device properties*
- *About device status*

About IP addresses for config sync, failover, and mirroring

Each trust domain member contains *device connectivity* information, that is, the IP addresses that you define on a device for configuration synchronization (config sync), failover, and connection mirroring.

***Note:** You specify a config sync address, as well as failover and mirroring addresses, for the local device only. You do not need to specify the addresses of peer devices because devices in a device group exchange their addresses automatically during device discovery.*

Config sync IP address

This is the IP address that you want the BIG-IP[®] system to use when synchronizing configuration objects to the local device.

By default, the system uses the self IP address of `VLAN internal`. This is the recommended IP address to use for config sync. You can, however, use a different self IP address for config sync.

***Important:** A self IP address is the only type of BIG-IP system address that encrypts the data during synchronization. For this reason, you cannot use a management IP address for config sync.*

Failover IP addresses

These are the IP addresses that you want the BIG-IP system to use when another device in the device group fails over to the local device. You can specify two types of addresses: unicast and multicast.

For appliance platforms, specifying two unicast addresses should suffice. For VIPRION[®] platforms, you should also retain the default multicast address that the BIG-IP system provides.

The recommended unicast addresses for failover are:

- The self IP address that you configured for either `VLAN HA` or `VLAN internal`. If you created `VLAN HA` when you initially ran the Setup utility on the local device, F5 recommends that you use the self IP address for that VLAN. Otherwise, use the self IP address for `VLAN internal`.
- The IP address for the local management port.

Mirroring IP addresses

These are the IP addresses that you want the BIG-IP system to use for connection mirroring. You specify both a primary address, as well as a secondary address for the system to use if the primary address is unavailable. If you configured VLAN `HA`, the system uses the associated self IP address as the default address for mirroring. If you did not configure VLAN `HA`, the system uses the self IP address of VLAN `internal`.

Note: On a VIPRION® system, you can mirror connections between blades within the cluster (intra-cluster mirroring) or between the clusters in a redundant system configuration (inter-cluster mirroring).

About device properties

From the local BIG-IP® device, you can view or configure the properties of any device within the local trust domain, including the local device.

Viewing device properties

On each member of the local trust domain, the BIG-IP® system generates a set of information. This information consists of properties such as the device name, serial number, and management IP address. By default, every BIG-IP device in the local trust domain has a set of device properties. You can use the BIG-IP Configuration utility to view these properties.

1. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
2. In the Name column, click the name of the device for which you want to view properties.
This displays a table of properties for the device.

Specifying values for device properties

Using the BIG-IP® Configuration utility, you can specify values for a few of the properties for a device. The device properties that you can specify provide information about the device for you to refer to when needed. For the HA Capacity property in particular, you can specify a relative capacity of the device compared to other BIG-IP devices in a Sync-Failover device group, as a way to affect the device that the BIG-IP system chooses as the next-active device. All of these property values are optional.

1. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
2. In the Name column, click the name of the device for which you want to specify properties.
This displays a table of properties for the device.
3. In the **Description** field, type a description of the device.
4. In the **Location** field, type a location for the device.
5. In the **Contact** field, type contact information for the device.
6. In the **Comment** field, type a comment about the device.
7. In the **HA Capacity** field, type a relative numeric value.

You need to configure this setting only when you have varying types of hardware platforms in a device group and you want to configure load-aware failover. The value you specify represents the relative capacity of the device to process application traffic compared to the other devices in the device group.

Important: *If you configure this setting, you must configure the setting on every device in the device group.*

If this device has half the capacity of a second device and a third of the capacity of a third device in the device group, you can specify a value of 100 for this device, 200 for the second device, and 300 for the third device.

When choosing the next active device for a traffic group, the system considers the capacity that you specified for this device.

8. Click **Update**.

Device properties

The following table lists and describes the properties of a device.

Property	Description
Device name	The name of the device, such as <code>siterequest</code> .
Host name	The host name of the device, such as <code>www.siterequest.com</code>
Device address	The IP address for the management port.
Serial number	The serial number of the device.
Platform MAC address	The MAC address for the management port.
Description	A user-created description of the device.
Location	The location of the device, such as <code>Seattle, Bldg. 1</code>
Contact	The name of the person responsible for this device.
Comment	Any user-specified remarks about the device.
HA Capacity	An arbitrary, user-specified value that represents the capacity of the device relative to other device group members.
Status	The status of the device, such as <code>Device is active</code>
Time zone	The time zone in which the device resides.
Platform ID	An identification for the platform.
Platform name	The platform name, such as <code>BIG-IP 8900</code> .
Software version	The BIG-IP version number, such as <code>BIG-IP 11.0.0</code> .
Active modules	The complete list of active modules, that is, the modules for which the device is licensed.

About device status

A BIG-IP® device can have any status shown in the following table.

Status	Description
Active	A minimum of one floating traffic group is currently active on the device. This status applies to Sync-Failover device groups only.
Forced offline	An administrator has intentionally made the device unavailable for processing traffic.
Offline	The device is unavailable for processing traffic.
Standby	The device is available for processing traffic, but all traffic groups on the device are in a standby state. This status applies to Sync-Failover device groups only.
Unknown	The status of the device is unknown.

Viewing possible status types for a device

You can view a list of possible status types for a device.

1. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
2. In the status column, click **Status**.
This displays a list of all possible status types for a device.

Viewing the status of a device

You can view the status of a device in a device group. Viewing the status of a device can help with troubleshooting or to verify that the devices in the device group are working properly.

1. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
2. In the Name column, locate the name of the device for which you want to view status.
3. In the Status column, view the status of the device.

Device status

At all times, the BIG-IP® system displays a specific status for each device in a device group.

Table 9: Possible statuses of a DSC™ device

Device status	Description
Active	The device is available and is processing traffic on the network. If the device is a member of a Sync-Failover device group, this status indicates that at least one traffic group is active on the device.
Forced Offline	An authorized user has intentionally taken the device offline, usually for maintenance purposes.
Offline	The device is offline for a reason other than being forced offline by an authorized user.

Device status	Description
Standby	The device is available but is not processing traffic on the network. This applies to devices in a Sync-Failover device group only, and all traffic groups on the device are Standby traffic groups only.
Unknown/Not Watched	The BIG-IP system cannot determine the status of the device. This status usually occurs when the device has not yet joined a device group.

Chapter 5

Managing Device Trust

- *What is device trust?*
- *Types of trust authority*
- *Device identity*
- *Device discovery in a local trust domain*
- *Establishing device trust*
- *Adding a device to the local trust domain*
- *Managing trust authority for a device*

What is device trust?

Before any BIG-IP® devices on a local network can synchronize configuration data or fail over to one another, they must establish a trust relationship known as device trust. *Device trust* between any two BIG-IP devices on the network is based on mutual authentication through the signing and exchange of x509 certificates.

Devices on a local network that trust one another constitute a trust domain. A *trust domain* is a collection of BIG-IP devices that trust one another and can therefore synchronize and possibly fail over their BIG-IP configuration data, as well as exchange status and failover messages on a regular basis. A *local trust domain* is a trust domain that includes the local device, that is, the device you are currently logged in to. You can synchronize a device's configuration data with either all of the devices in the local trust domain, or to a subset of devices in the local trust domain.

Note: *You can add devices to a local trust domain from a single device on the network. You can also view the identities of all devices in the local trust domain from a single device in the domain. However, to maintain or change the authority of each trust domain member, you must log in locally to each device.*

Types of trust authority

Within a local trust domain, in order to establish device trust, you designate each BIG-IP® device as either a certificate signing authority or a subordinate non-authority. For each device, you also specify peer authorities.

Certificate signing authorities

A *certificate signing authority* can sign x509 certificates for another BIG-IP device that is in the local trust domain. For each authority device, you specify another device as a peer authority device that can also sign

certificates. In a standard redundant system configuration of two BIG-IP devices, both devices are typically certificate signing authority devices.

Important: For security reasons, F5 Networks recommends you limit the number of authority devices in a local trust domain to as few as possible.

Subordinate non-authorities

A *subordinate non-authority device* is a device for which a certificate signing authority device signs its certificate. A subordinate device cannot sign a certificate for another device. Subordinate devices provide an additional level of security because in the case where the security of an authority device in a trust domain is compromised, the risk of compromise is minimized for any subordinate device. Designating devices as subordinate devices is recommended for device groups with a large number of member devices, where the risk of compromise is high.

Peer authorities

A *peer authority* is another device in the local trust domain that can sign certificates if the certificate signing authority is not available. In a standard redundant system configuration of two BIG-IP devices, each device is typically a peer authority for the other.

Device identity

The devices in a BIG-IP[®] device group use x509 certificates for mutual authentication. Each device in a device group has an x509 certificate installed on it that the device uses to authenticate itself to the other devices in the group.

Device identity is a set of information that uniquely identifies that device in the device group, for the purpose of authentication. Device identity consists of the x509 certificate, plus this information:

- Device name
- Host name
- Platform serial number
- Platform MAC address
- Certificate name
- Subjects
- Expiration
- Certificate serial number
- Signature status

Tip: From the Device Trust: Identity screen in the BIG-IP Configuration utility, you can view the x509 certificate installed on the local device.

Device discovery in a local trust domain

When a BIG-IP[®] device joins the local trust domain and establishes a trust relationship with peer devices, the device and its peers exchange their device properties and device connectivity information. This exchange of device properties and IP addresses is known as *device discovery*.

For example, if a device joins a trust domain that already contains three trust domain members, the device exchanges device properties with the three other domain members. The device then has a total of four sets of device properties defined on it: its own device properties, plus the device properties of each peer. In this exchange, the device also learns the relevant device connectivity information for each of the other devices.

Establishing device trust

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

You perform this task to establish trust among devices on one or more network segments. Devices that trust each other constitute the *local trust domain*. A device must be a member of the local trust domain prior to joining a device group.

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain. For example, devices A, B, and C each initially shows only itself as a member of the local trust domain. To configure the local trust domain to include all three devices, you can simply log into device A and add devices B and C to the local trust domain. Note that there is no need to repeat this process on devices B and C.

1. On the Main tab, click **Device Management > Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.
3. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
 - If the BIG-IP device is a non-VIPRION® device, type the management IP address for the device.
 - If the BIG-IP device is a VIPRION device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
 - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
 - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
4. Click **Retrieve Device Information**.
5. Verify that the certificate of the remote device is correct.
6. Verify that the name of the remote device is correct.
7. Verify that the management IP address and name of the remote device are correct.
8. Click **Finished**.

The device you added is now a member of the local trust domain.

Repeat this task for each device that you want to add to the local trust domain.

Adding a device to the local trust domain

Verify that each BIG-IP® device that is to be part of a local trust domain has a device certificate installed on it.

Follow these steps to log in to any BIG-IP® device on the network and add one or more devices to the local system's local trust domain.

***Note:** Any BIG-IP devices that you intend to add to a device group at a later point must be members of the same local trust domain.*

1. On the Main tab, click **Device Management > Device Trust**, and then either **Peer List** or **Subordinate List**.
2. In the Peer Authority Devices or the Subordinate Non-Authority Devices area of the screen, click **Add**.
3. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
 - If the BIG-IP device is a non-VIPRION® device, type the management IP address for the device.
 - If the BIG-IP device is a VIPRION device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
 - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
 - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
4. Click **Retrieve Device Information**.
5. Verify that the displayed information is correct.
6. Click **Finished**.

After you perform this task, the local device and the device that you specified in this procedure have a trust relationship and, therefore, are qualified to join a device group.

Managing trust authority for a device

You can use a Reset Device Trust wizard in the BIG-IP® Configuration utility to manage the certificate authority of a BIG-IP device in a local trust domain. Specifically, you can:

- Retain the current authority (for certificate signing authorities only).
- Regenerate the self-signed certificate for a device.
- Import a user-defined certificate authority.

***Warning:** If you reset trust authority on a certificate signing authority by retaining the authority of the device, you must subsequently recreate the local trust domain and the device group. If you reset trust authority on a subordinate non-authority, the BIG system removes the non-authority device from the local trust domain. You can then re-add the device as an authority or non-authority device.*

1. On the Main tab, click **Device Management > Device Trust > Local Domain**.
2. In the Trust Information area of the screen, click **Reset Device Trust**.

3. Choose a certificate signing authority option, and then click **Update**.
The system prompts you to confirm your choice.

When you confirm your choice, the system changes the **Authority Type**.

Chapter 6

Working with Device Groups

- *About Sync-Failover device groups*
- *About Sync-Only device groups*
- *Viewing a list of device groups*
- *Viewing the members of a device group*
- *Adding a device to a device group*
- *A note about folders and overlapping device groups*

About Sync-Failover device groups

One of the types of device groups that you can create is a Sync-Failover type of device group. A *Sync-Failover* device group contains devices that synchronize their configuration data and fail over to one another when a device becomes unavailable. A Sync-Failover device group supports a maximum of eight devices.




Figure 4: traffic_group_1 is active on a device in a Sync-Failover device group




Figure 5: On failover, traffic_group_1 becomes active on another device in the Sync-Failover device group

A device in the trust domain can be a member of both a Sync-Failover group and a Sync-Only group simultaneously.

For devices in a Sync-Failover group, the BIG-IP® system uses both the device group and the traffic group attributes of a folder to make decisions about which devices to target for synchronizing the contents of the folder, and which application-related configuration objects to include in failover.

You can control the way that the BIG-IP chooses a target failover device. This control is especially useful when a device group contains heterogeneous hardware platforms that differ in load capacity, because you can ensure that when failover occurs, the system will choose the device with the most available resource to process the application traffic.

Sample Sync-Failover configuration

You can use a Sync-Failover device group in a variety of ways. This sample configuration shows two separate Sync-Failover device groups in the local trust domain. Device group A is a standard active-standby configuration. Prior to failover, only Bigip1 processes traffic for application A. This means that Bigip1 and Bigip2 synchronize their configurations, and Bigip1 fails over to Bigip2 if Bigip1 becomes unavailable. Bigip1 cannot fail over to Bigip3 or Bigip4 because those devices are in a separate device group.

Device group B is also a standard active-standby configuration, in which Bigip3 normally processes traffic for application B. This means that Bigip3 and Bigip4 synchronize their configurations, and Bigip3 fails over to Bigip4 if Bigip3 becomes unavailable. Bigip3 cannot fail over to Bigip1 or Bigip2 because those devices are in a separate device group.




Figure 6: Sample Sync-Failover device groups in a trust domain

Sync-Failover device group considerations

The following configuration restrictions apply to Sync-Failover device groups:

- A specific BIG-IP device in a trust domain can belong to one Sync-Failover device group only.
- On each device in a Sync-Failover device group, the BIG-IP® system automatically assigns the device group name to the `root` and `/Common` folders. This ensures that the system synchronizes any traffic groups for that device to the correct devices in the local trust domain.

- The BIG-IP system creates all device groups and traffic-groups in the /Common folder, regardless of the partition to which the system is currently set.
- If no Sync-Failover device group is defined on a device, then the system sets the device group value that is assigned to the root and /Common folders to None.
- By default, on each device, the BIG-IP system assigns a Sync-Failover device group to any sub-folders of the root or /Common folders that inherit the device group attribute.
- You can configure a maximum of 15 floating traffic groups for a Sync-Failover device group.

Creating a Sync-Failover device group

This task establishes failover capability between two or more BIG-IP® devices. If an active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

Repeat this task for each Sync-Failover device group that you want to create for your network configuration.

1. On the Main tab, click **Device Management > Device Groups**.
2. On the Device Groups list screen, click **Create**.
The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. From the **Configuration** list, select **Advanced**.
5. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Includes** list.

The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only.

6. For the **Network Failover** setting, select or clear the check box:
 - Select the check box if you want device group members to handle failover communications by way of network connectivity. This choice is required for active-active configurations.
 - Clear the check box if you want device group members to handle failover communications by way of serial cable (hard-wired) connectivity.

For active-active configurations, you must select network failover, as opposed to serial-cable (hard-wired) connectivity.

7. For the **Automatic Sync** setting, select or clear the check box:
 - Select the check box when you want the BIG-IP system to automatically sync the BIG-IP configuration data whenever a config sync operation is required. In this case, the BIG-IP system syncs the configuration data whenever the data changes on any device in the device group.
 - Clear the check box when you want to manually initiate each config sync operation. In this case, F5 networks recommends that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.
8. For the **Full Sync** setting, select or clear the check box:
 - Select the check box when you want all sync operations to be full syncs. In this case, the BIG-IP system syncs the entire set of BIG-IP configuration data whenever a config sync operation is required.
 - Clear the check box when you want all sync operations to be incremental (the default setting). In this case, the BIG-IP system syncs only the changes that are more recent than those on the target device. When you select this option, the BIG-IP system compares the configuration data on each

target device with the configuration data on the source device and then syncs the delta of each target-source pair.

If you enable incremental synchronization, the BIG-IP system might occasionally perform a full sync for internal reasons. This is a rare occurrence and no user intervention is required.

- In the **Maximum Incremental Sync Size (KB)** field, retain the default value of 1024, or type a different value.

This value specifies the total size of configuration changes that can reside in the incremental sync cache. If the total size of the configuration changes in the cache exceeds the specified value, the BIG-IP system performs a full sync whenever the next config sync operation occurs.

- Click **Finished**.

You now have a Sync-Failover type of device group containing BIG-IP devices as members.

About Sync-Only device groups

One of the types of device groups that you can create is a Sync-Only device group. A *Sync-Only* device group contains devices that synchronize configuration data with one another, but their configuration data does not fail over to other members of the device group. A Sync-Only device group supports a maximum of 32 devices.

A device in a trust domain can be a member of more than one Sync-Only device group. A device can also be a member of both a Sync-Failover group and a Sync-Only group simultaneously.

A typical use of a Sync-Only device group is one in which you configure a device to synchronize the contents of a specific folder to a different device group than to the device group to which the other folders are synchronized.




Figure 7: Sync-only device group

Sample Sync-Only configuration

The most common reason to use a Sync-Only device group is to synchronize a specific folder containing policy data that you want to share across all BIG-IP® devices in a local trust domain, while setting up a Sync-Failover device group to fail over the remaining configuration objects to a subset of devices in the

domain. In this configuration, you are using a Sync-Only device group attribute on the policy folder to override the inherited Sync-Failover device group attribute. Note that in this configuration, `Bigip1` and `Bigip2` are members of both the Sync-Only and the Sync-Failover groups.

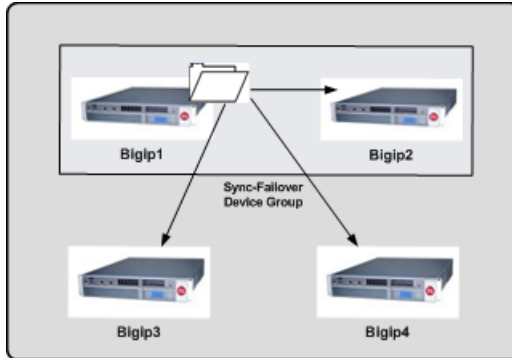


Figure 8: Sync-Only Device Group

To implement this configuration, you can follow this process:

1. Create a Sync-Only device group on the local device, adding all devices in the local trust domain as members.
2. Create a Sync-Failover device group on the local device, adding a subset of devices as members.
3. On the folder containing the policy data, use `tmsh` to set the value of the device group attribute to the name of the Sync-Only device group.
4. On the `root` folder, retain the default Sync-Failover device group assignment.

Creating a Sync-Only device group

You perform this task to create a Sync-Only type of device group. When you create a Sync-Only device group, the BIG-IP® system can then automatically synchronize certain types of data such as security policies and acceleration applications and policies to the other devices in the group, even when some of those devices reside in another network. You can perform this task on any BIG-IP device within the local trust domain.

1. On the Main tab, click **Device Management > Device Groups**.
2. On the Device Groups list screen, click **Create**.
The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Only**, and type a description for the device group.
4. From the **Configuration** list, select **Advanced**.
5. For the **Members** setting, select an IP address and host name from the **Available** list for each BIG-IP device that you want to include in the device group. Use the Move button to move the host name to the **Includes** list.
The list shows any devices that are members of the device's local trust domain.
6. For the **Automatic Sync** setting, select or clear the check box:
 - Select the check box when you want the BIG-IP system to automatically sync the BIG-IP configuration data whenever a config sync operation is required. In this case, the BIG-IP system syncs the configuration data whenever the data changes on any device in the device group.
 - Clear the check box when you want to manually initiate each config sync operation. In this case, F5 networks recommends that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.

- For the **Full Sync** setting, select or clear the check box:
 - Select the check box when you want all sync operations to be full syncs. In this case, the BIG-IP system syncs the entire set of BIG-IP configuration data whenever a config sync operation is required.
 - Clear the check box when you want all sync operations to be incremental (the default setting). In this case, the BIG-IP system syncs only the changes that are more recent than those on the target device. When you select this option, the BIG-IP system compares the configuration data on each target device with the configuration data on the source device and then syncs the delta of each target-source pair.

If you enable incremental synchronization, the BIG-IP system might occasionally perform a full sync for internal reasons. This is a rare occurrence and no user intervention is required.

- In the **Maximum Incremental Sync Size (KB)** field, retain the default value of 1024, or type a different value.

This value specifies the total size of configuration changes that can reside in the incremental sync cache. If the total size of the configuration changes in the cache exceeds the specified value, the BIG-IP system performs a full sync whenever the next config sync operation occurs.

- Click **Finished**.

You now have a Sync-Only type of device group containing BIG-IP devices as members.

Viewing a list of device groups

You can perform this task when you want to display a list of the device groups of which the local device is a member.

- On the Main tab, click **Device Management > Overview**.
- In the Device Groups area of the screen, in the Name column, view the list of device groups.

The list shows all device groups that include the local device as a member, as well as the sync status of each group.

Viewing the members of a device group

You can list the members of a device group and view information about them, such as their management IP addresses and host names.

- On the Main tab, click **Device Management > Device Groups**.
- In the Group Name column, click the name of the relevant device group.

The screen shows a list of the device group members.

Adding a device to a device group

You must ensure that the device you are adding is a member of the local trust domain.

You can use this procedure to add a member to an existing device group.

1. On the Main tab, click **Device Management > Device Groups**.
2. In the Group Name column, click the name of the relevant device group.
3. In the Members area of the screen, select a host name from the **Available** list for each BIG-IP® device that you want to include in the device group. Use the Move button to move the host name to the **Selected** list.

The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. If you are attempting to add a member to a Sync-Failover group and you do not see the member name in the list, it is possible that the device is already a member of another Sync-Failover device group. A device can be a member of one Sync-Failover group only.

4. Click **Update**.

A note about folders and overlapping device groups

Sometimes when one BIG-IP® object references another, one of the objects gets synchronized to a particular device, but the other object does not. This can result in an invalid device group configuration.

For example, suppose you create two device groups that share some devices but not all. In the following illustration, Device A is a member of both Device Group 1 and Device Group 2.




Figure 9: One device with membership in two device groups

Device Group 1 is associated with folder `/Common`, and Device Group 2 is associated with the folder `/Common/my_app`. This configuration causes Device A to synchronize all of the data in folder `/Common` to Device B in Device Group 1. The only data that Device A can synchronize to Device C in Device Group 2 is the data in the folder `/Common/my_app`, because this folder is associated with Device Group 2 instead of Device Group 1.

Now suppose that you create a pool in the `/Common/my_app` folder, which is associated with Device Group 2. When you create the pool members in that folder, the BIG-IP system automatically creates the associated node addresses and puts them in folder `/Common`. This results in an invalid configuration, because the node objects in folder `/Common` do not get synchronized to the device on which the nodes' pool members reside, Device C. When an object is not synchronized to the device on which its referenced objects reside, an invalid configuration results.

Chapter 7

Managing Configuration Synchronization

- *About configuration synchronization*
- *About automatic and manual sync*
- *About full and incremental sync*
- *Specifying an IP address for config sync*
- *Viewing config sync status for the local device*
- *Viewing config sync status for all device groups and members*
- *Troubleshooting the config sync process*

About configuration synchronization

Configuration synchronization (also known as *config sync*) is the operation that the BIG-IP® system performs to propagate BIG-IP configuration changes to all devices in a device group. BIG-IP devices that contain the same configuration data can work in tandem to more efficiently process application traffic on the network.

If you want to exclude certain devices from config sync, you simply exclude them from membership in that particular device group.

You can sync some types of data on a global level across all BIG-IP devices, while syncing other data in a more granular way, on an individual application level to a subset of devices. For example, you can set up a large device group to sync resource and policy data (such as iRules® and profiles) among all BIG-IP devices in a data center, while setting up a smaller device group for syncing application-specific data (such as virtual IP addresses) between the specific devices that are delivering those applications.

Whenever synchronization occurs, either automatically or manually, the BIG-IP system attempts to optimize performance by syncing only the data that changed since the last config sync operation.

Important: *To synchronize configuration data among device group members, all members must be running the same version of the BIG-IP system software.*

About automatic and manual sync

You can configure the BIG-IP® system to synchronization configuration data automatically, or you can manually initiate synchronization:

Automatic

Automatic synchronization (also known as *auto sync*) ensures that the BIG-IP system automatically synchronizes the configuration among device group members whenever you make a change to any one of those devices.

Manual

If you do not enable auto sync, you must manually synchronize the BIG-IP configuration among device group members to ensure that the devices remain in sync. With manual synchronization, the BIG-IP system notifies you whenever configuration data within the group has changed and therefore needs to be synchronized.

Enabling and disabling automatic sync

You can use the BIG-IP® Configuration utility to enable or disable automatic synchronization for Sync-Failover and Sync-Only device groups. When you enable automatic synchronization, a BIG-IP device in the device group automatically synchronizes its configuration data to the other members of the device group whenever its configuration data changes.

By default, the BIG-IP system syncs only the data that changed since the previous sync, rather than the entire set of configuration data.

1. On the Main tab, click **Device Management > Device Groups**.
2. In the Group Name column, click the name of the relevant device group.
3. For the **Automatic Sync** setting, select or clear the check box:

Action	Result
Select (Enable)	Select the check box when you want the BIG-IP system to automatically sync configuration data to device group members whenever a change occurs. When you enable this setting, the BIG-IP system automatically syncs, but does not save, the configuration change on each device (this is the default behavior). To save the updated configuration on each device, you can log into each device and, at the <code>tmsh</code> prompt, type <code>save sys config</code> . Alternatively, you can change the default behavior so that the system automatically saves configuration changes on target devices after an automatic config sync. You make this change by logging into one of the devices in the device group and, at the <code>tmsh</code> prompt, typing <code>modify cm device-group name save-on-auto-sync true</code> .

Warning: Enabling the `save-on-auto-sync` option can unexpectedly impact system performance when the BIG-IP system automatically saves a large configuration change to each device.

Clear (Disable)	Clear the check box when you want to disable automatic sync. When this setting is disabled, you must manually initiate each config sync operation. F5 Networks® recommends that you perform a config sync whenever configuration data changes on one of the devices in the device group. After you perform a manual config sync, the BIG-IP system automatically saves the configuration change on each device group member.
------------------------	--

4. Click **Update**.

After you enable automatic sync, the BIG-IP system automatically syncs future configuration changes to each device group member.

Whenever an automatic sync operation occurs, you must log in to each device group member and use the Traffic Management shell to save the configuration on the device. An alternative is to configure the `tmsh save-on-auto-sync` option for the device group.

Manually synchronizing the BIG-IP configuration

Before you perform this task, verify that device trust has been established and that all devices that you want to synchronize are members of a device group.

You perform this task when the automatic sync feature is disabled and you want to manually synchronize BIG-IP® configuration data among the devices in the device group. This synchronization ensures that any device in the device group can process application traffic successfully. You can determine the need to perform this task by viewing sync status in the upper left corner of any BIG-IP Configuration utility screen. A status of `Changes Pending` indicates that you need to perform a config sync within the device group.

Important: *You can log into any device in the device group to perform this task.*

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select a device.
4. From the **Sync** options list, select an option:

Option	Description
Sync Device to Group	Select this option when you want to sync the configuration of the selected device to the other device group members.
Sync Group to Device	Select this option when you want to sync the most recent configurations of one or more device group members to the selected device.

5. Click **Sync**.

After you initiate a manual config sync, the BIG-IP system compares the configuration data on the local device with the data on each device in the device group, and synchronizes the most recently-changed configuration data from one or more source devices to one or more target devices. The system then automatically saves the synced configuration changes on each device. Note that the system does not synchronize non-floating self IP addresses.

About full and incremental sync

You can configure the BIG-IP® system to perform either full or incremental synchronization operations whenever a config sync is required:

Full

When you enable *full sync*, the BIG-IP system syncs the entire set of BIG-IP configuration data whenever a config sync operation occurs.

Incremental

When you enable *incremental sync*, the BIG-IP system syncs only the changes that are more recent than those on the target device. The BIG-IP system accomplishes this by comparing the configuration data on each target device with the configuration data on the source device and then syncs the delta of each target-source pair. F5 networks recommends that you use incremental sync, for optimal performance. The incremental sync feature is a performance improvement feature and is the default value.

You can also configure the cache size for any configuration changes slated for incremental sync. (This applies to incremental sync only.) For example, using the default cache size value of 1024, if you make more than 1024 KB worth of incremental changes, the system performs a full synchronization operation. Using incremental synchronization operations can reduce the per-device sync/load time for configuration changes.

Enabling and disabling full sync

You can enable or disable full synchronization for Sync-Failover and Sync-Only device groups. When you enable *full sync*, the BIG-IP[®] system syncs the entire set of configuration data whenever a sync operation occurs. When you disable full synchronization, the BIG-IP system performs *incremental synchronization*, which causes the system to sync only the changes that are more recent than the changes on the target device. The incremental sync feature is a performance improvement feature.

1. On the Main tab, click **Device Management > Device Groups**.
2. In the Group Name column, click the name of the relevant device group.
3. For the **Full Sync** setting, select or clear the check box:
 - Select the check box when you want all sync operations to be full syncs. In this case, the BIG-IP system syncs the entire set of BIG-IP configuration data whenever a config sync operation is required.
 - Clear the check box when you want all sync operations to be incremental (the default setting). In this case, the BIG-IP system syncs only the changes that are more recent than those on the target device. When you select this option, the BIG-IP system compares the configuration data on each target device with the configuration data on the source device and then syncs the delta of each target-source pair.

If you enable incremental synchronization, the BIG-IP system might occasionally perform a full sync for internal reasons. This is a rare occurrence and no user intervention is required.

4. Click **Update**.

After you configure this feature, the BIG-IP system performs either a full or an incremental sync whenever a sync operation occurs.

Specifying an IP address for config sync

Before configuring the config sync address, verify that all devices in the device group are running the same version of BIG-IP[®] system software.

You perform this task to specify the IP address on the local device that other devices in the device group will use to synchronize their configuration objects to the local device.

Note: You must perform this task locally on each device in the device group.

1. Confirm that you are logged in to the actual device you want to configure.

2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose ConfigSync.
5. For the **Local Address** setting, retain the displayed IP address or select another address from the list.
F5 Networks recommends that you use the default value, which is the self IP address for VLAN `internal`. This address must be a non-floating self IP address and not a management IP address.

Important: *If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the internal self IP address that you specify must be the internal private IP addresses that you configured for this EC2 instance as the **Local Address**.*

6. Click **Update**.

After performing this task, the other devices in the device group can sync their configurations to the local device.

Viewing config sync status for the local device

You can use the BIG-IP® Configuration utility to view the config sync status of the local device relative to the other members of the device group. If you have configured the device group for manual synchronization, you can use the config sync status information to determine whether you need to perform a manual sync operation.

1. Display any BIG-IP Configuration utility screen.
2. In the upper left corner of the screen, view the status of the device group:
 - If the sync status is green (`In Sync`), the local device is synchronized with all device group members, and you do not need to perform a config sync operation.
 - If the sync status is yellow (`Changes Pending`), the local device is out of sync with one or more device group members. You must therefore ensure that a config sync operation occurs for the local device. If the **Automatic Sync** setting is enabled for the device group, the BIG-IP system synchronizes the configuration automatically, and no user action is required.

For more details, you can click the status, which displays more information:

- If the status pertains to config sync specifically, you can click on the status displays the Overview screen. Using this screen, you can view a detailed message about the status, as well as the status of each device group member.
- If the status pertains to an issue with device trust, you can click on the status displays the Device Trust screen. Using this screen, you can re-establish trust among all device group members or add devices to the trust domain.

Viewing config sync status for all device groups and members

You can use the BIG-IP® Configuration utility to view the config sync status of any device group and each of its members. If the **Automatic Sync** setting is disabled for a device group, you can use the config sync status information to determine whether a manual sync operation is needed.

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, view the sync status of each device:
 - If all devices show a sync status of green, the configurations of all device members are synchronized, and you do not need to perform a config sync operation.
 - If any device shows a sync status of *Changes Pending*, you must synchronize the configuration on that device to the other members of the device group.

A status of *Changes Pending* for a device indicates that the device contains recent configuration changes that have not yet been synchronized to the other members of the device group.

Troubleshooting the config sync process

The BIG-IP® Configuration utility displays a number of different statuses and messages to help you diagnose and correct a config sync problem. These statuses and messages pertain to both device groups and individual device group members.

Sync status for device groups

At all times, the BIG-IP® system displays a specific sync status for each device group.

Table 10: Possible sync status for device groups

Color	Sync Status	Summary Message	Explanation and Recommended Action
Green	In Sync	All devices in the device group are in sync	All devices in the device group contain the current configuration. Recommended action: None.
Green	Standalone	None.	The local trust domain contains one member only, which is the local device. Recommended action: None. You can optionally add other devices to the local trust domain.
Blue	Awaiting Initial Sync	None.	All devices have been recently added to the device group and are awaiting an initial config sync. Recommended action: Sync any one of the devices to the device group.
Blue	Awaiting Initial Sync	<i>Device_name1, device_name2, etc.</i> awaiting the initial config sync	One or more of the devices in the device group has either not yet synchronized its data to the device group members or has not yet received a sync from another member.

Color	Sync Status	Summary Message	Explanation and Recommended Action
Green	Syncing	None.	<p>Recommended action: View the individual sync status of each device group member, and then sync the device with the most current configuration to the other devices.</p> <p>A sync operation is in progress.</p> <p>Recommended action: None.</p>
Yellow	Changes Pending	Changes Pending	<p>One or more devices in the device group has recent configuration changes that have not yet been synchronized to the other members of the device group.</p> <p>Recommended action: View the individual sync status of each device group member, and then sync the device with the most current configuration to the device group.</p>
Yellow	Changes Pending	There is a possible change conflict between <i>device_name1</i> , <i>device_name2</i> , etc.	<p>There is a possible conflict among two or more devices because more than one device contains changes that have not been synchronized to the device group.</p> <p>Recommended action: View the individual sync status of each device group member, and then sync the device with the most current configuration to the device group.</p>
Red	Not All Devices Synced	<i>Device_name1</i> , <i>device_name2</i> , etc. did not receive last sync successfully.	<p>One or more of the devices in the device group does not contain the most current configuration.</p> <p>Recommended action: View the individual sync status of each device group member, and then sync the device with the most current configuration to the device group.</p>
Red	Sync Failure	A validation error occurred while syncing to a remote device	<p>Because of a validation error, the named device was unable to accept a sync successfully.</p> <p>Recommended action: Review the error message and determine corrective action on the device.</p>
Blue	Unknown	The local device is not a member of the selected device group	<p>The device that you are logged into is not a member of the selected device group.</p> <p>Recommended action: Add the local device to the device group to view sync status for the device group.</p>
Blue	Unknown	Not logged into the primary cluster member	<p>The system cannot determine the sync status of the device group because you are logged in to a secondary cluster member instead of the primary cluster member. Pertains to VIPRION® systems only.</p>

Color	Sync Status	Summary Message	Explanation and Recommended Action
Red	Unknown	Error in trust domain	<p>Recommended action: Log out and then log in to the primary cluster member, using the primary cluster IP address.</p> <p>The trust relationships among devices in the device group are not properly established.</p> <p>Recommended action: On the local device, reset device trust and then re-add all relevant devices to the local trust domain.</p>
None.	None.	X devices with Y different configurations	<p>The configuration time for two or more devices in the device group differs from the configuration time of the other device group members. This condition causes one of these status messages to appear for each relevant device:</p> <ul style="list-style-type: none"> • <i>Device_name</i> awaiting initial config sync • <i>Device_name</i> made last configuration change on <i>date_time</i> <p>Recommended action: Identify a device with the most current configuration and sync the device to the device group.</p>

Sync status for device group members

At all times, the BIG-IP® system displays a specific sync status for each device within a device group.

Table 11: Possible sync status for individual devices

Color	Sync Status	Explanation and Recommended Action
Green	None.	<p>This status indicates one of the following conditions:</p> <ul style="list-style-type: none"> • The device has the most recent set of configuration data within the device group. • The device is a standalone device or is the only device group member. <p>Recommended action: None.</p>
Blue	Awaiting Initial Sync	<p>This status indicates one of the following conditions:</p> <ul style="list-style-type: none"> • The device has no configuration changes to be synced since joining the device group. • The device is member of a device group with autosync enabled, and no changes have been made on any device in the device group. • The device has not yet received a sync from another device and has no configuration changes to be synced to other members of the device group. <p>Recommended action: Perform the appropriate type of config sync.</p>

Color	Sync Status	Explanation and Recommended Action
Yellow	Changes Pending	<p>The device has recent configuration changes since the last sync that have not yet been synchronized to the other members of the device group.</p> <p>Recommended action: Sync the device with the most recent configuration to the other members of the device group.</p>
Red	Changes Pending	<p>A device group member reported in its most recent status update that it had changes pending but is now disconnected from the local device.</p> <p>Recommended action: When the disconnected device comes back online, sync the device with the most recent changes to the device group.</p>
Yellow	Awaiting Initial Sync with Changes Pending	<p>This status indicates one of the following conditions:</p> <ul style="list-style-type: none"> The configuration on the device has changed since the device joined the device group. <p>Recommended action: Sync the device to the device group.</p> <ul style="list-style-type: none"> The device has not yet received a sync from another device but has configuration changes to be synced to other members of the device group. <p>Recommended action: Sync the device with the most recent configuration to this device.</p>
Red	Does not have the last synced configuration, and has changes pending	<p>The device received at least one sync previously but did not receive the last synced configuration, and the configuration on the device has changed since the last sync.</p> <p>Recommended action: Sync the device with the most recent configuration to this device.</p>
Red	Disconnected	<p>The local device does not recognize the disconnected device.</p> <p>Recommended actions:</p> <ul style="list-style-type: none"> View the screens at Device Management > Device Trust to see if the disconnected device is a member of the local trust domain, and if not, add the device to the domain. Use the screen at Device Management>Devices to view the specified config sync address of the disconnected device and determine whether the local device has a route to that address.
Red	Device does not recognize membership in this group	<p>The device does not recognize that it is a member of the device group.</p> <p>Recommended action: Log into the relevant device and view the screens at Device Management > Device Groups to see if the device is a member of the device group. If not, add the device to the device group.</p>
Red	No config sync address has been specified for this device.	<p>The device does not have a config sync address.</p> <p>Recommended action: Log into the relevant device, and using the screen at Device Management > Devices, specify the IP address that you want remote devices to use to sync configuration data to the device. As a best practice, this address should be a non-floating self IP address associated with an internal VLAN. The address must</p>

Color	Sync Status	Explanation and Recommended Action
Red	Does not have the last synced configuration	<p>either be on the same subnet as the other devices in the device group or have a route to that address defined on the other devices.</p> <p>The device previously received the configuration from other members of the device group but did not receive the last synced configuration.</p> <p>Recommended action: Sync the device group to the device.</p>

Advanced config sync properties for a device

A device in a device group has several advanced properties.

Property	Description
CID Originator	<p>Commit ID originator. This indicates the source of the most recent change to the configuration on the relevant device. More specifically, the CID originator is either:</p> <ul style="list-style-type: none"> The relevant device itself (due to locally-made changes) Another device in the device group that synchronized a change to the relevant device
CID Time	<p>Commit ID time. This indicates either the last time that a user updated the configuration locally, or, if the configuration on the device was synced from a remote device group member, the actual time that the synced configuration change was made on that remote device.</p>
Last Sync Time	<p>This is the last time that a sync was initiated or forced to or from the relevant device.</p>
Last Sync Type	<p>This is the type of sync. Possible values are: <code>Manual Full Load</code>, <code>Manual Incremental</code>, and <code>Automatic</code>.</p>
LSS Originator	<p>Last Successful Sync originator. This is the device that most recently performed a successful sync operation to the relevant device.</p>
LSS Time	<p>This is the actual time that the synced configuration change was made on a remote device group member. Whenever a device in the device group syncs its configuration to the other device group members, the LSS time on each device is updated to reflect the Commit ID time of the configuration change on the device that initiated the sync operation.</p>

Chapter

8

Managing Failover

- *Introduction to failover*
- *About traffic groups*
- *Active and standby states*
- *About active-standby vs. active-active configurations*
- *Description of current and next-active devices*
- *About the next-active device*
- *About MAC masquerade addresses*

Introduction to failover

When you configure a Sync-Failover device group as part of device service clustering (DSC[®]), you ensure that a user-defined set of application-specific IP addresses, known as a *floating traffic group*, can fail over to another device in that device group if necessary. DSC failover gives you granular control of the specific configuration objects that you want to include in failover operations.

If you want to exclude certain devices on the network from participating in failover operations, you simply exclude them from membership in that particular Sync-Failover device group.

What triggers failover?

The BIG-IP system initiates failover according to any of several events that you define. These events fall into these categories:

System fail-safe

With *system fail-safe*, the BIG-IP system monitors various hardware components, as well as the heartbeat of various system services. You can configure the system to initiate failover whenever it detects a heartbeat failure.

Gateway fail-safe

With *gateway fail-safe*, the BIG-IP system monitors traffic between an active BIG-IP[®] system in a device group and a pool containing a gateway router. You can configure the system to initiate failover whenever some number of gateway routers in a pool of routers becomes unreachable.

VLAN fail-safe

With *VLAN fail-safe*, the BIG-IP system monitors network traffic going through a specified VLAN. You can configure the system to initiate failover whenever the system detects a loss of traffic on the VLAN and the fail-safe timeout period has elapsed.

HA groups

With an *HA group*, the BIG-IP system monitors trunk, pool, or cluster health to create an HA health score for a device. You can configure the system to initiate failover whenever the health score falls below configurable levels.

Auto-failback

When you enable *auto-failback*, a traffic group that has failed over to another device fails back to a preferred device when that device is available. If you do not enable auto-failback for a traffic group, and the traffic group fails over to another device, the traffic group remains active on that device until that device becomes unavailable.

About IP addresses for failover

Part of configuring a Sync-Failover device group is configuring failover. Configuring failover requires you to specify certain types of IP addresses on each device. Some of these IP addresses enable continual, high availability (HA) communication among devices in the device group, while other addresses ensure that application traffic processing continues when failover occurs.

The types of IP addresses that you need to specify on each device are:

A local, static self IP address for VLAN HA

This unicast self IP address is the main address that other devices in the device group use to communicate continually with the local device to assess the health of that device. When a device in the device group fails to receive a response from the local device, the BIG-IP[®] system triggers failover.

A local management IP address

This unicast management IP address serves the same purpose as the static self IP address for VLAN HA , but is only used when the local device is unreachable through the HA static self IP address.

One or more floating IP addresses associated with a traffic group

These are the IP addresses that application traffic uses when passing through a BIG-IP system. Each traffic group on a device includes application-specific floating IP addresses as its members. Typical traffic group members are: floating self IP addresses, virtual addresses, NAT or SNAT translation addresses, and IP addresses associated with an iApp application service. When a device with active traffic groups becomes unavailable, each of the active traffic groups becomes active on another device in the device group. This ensures that application traffic processing continues with little to no interruption.

Specifying IP addresses for failover communication

You typically perform this task during initial Device Service Clustering (DSC[®]) configuration, to specify the local IP addresses that you want other devices in the device group to use for continuous health-assessment communication with the local device or guest. You must perform this task locally on each device in the device group.

Important: *If the system is running vCMP, you must log in to each guest to perform this task.*

Note: *The IP addresses that you specify must belong to route domain 0.*

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.

4. From the Device Connectivity menu, choose Failover.
5. For the Failover Unicast Configuration settings, click **Add** for each IP address on this device that other devices in the device group can use to exchange failover messages with this device. The unicast IP addresses you specify depend on the type of device:

Platform	Action
Appliance without vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}) and the static management IP address currently assigned to the device.
Appliance with vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}) and the unique management IP address currently assigned to the guest.
VIPRION without vCMP®	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}). If you choose to specify unicast addresses only (and not a multicast address), you must also type the existing, static management IP addresses that you previously configured for all slots in the cluster. If you choose to specify one or more unicast addresses and a multicast address, then you do not need to specify the existing, per-slot static management IP addresses when configuring addresses for failover communication.
VIPRION with vCMP	Type a self IP address that is defined on the guest and associated with an internal VLAN on the host (preferably VLAN _{HA}). If you choose to specify unicast failover addresses only (and not a multicast address), you must also type the existing, virtual static management IP addresses that you previously configured for all slots in the guest's virtual cluster. If you choose to specify one or more unicast addresses and a multicast address, you do not need to specify the existing, per-slot static and virtual management IP addresses when configuring addresses for failover communication.

Important: Failover addresses should always be static, not floating, IP addresses.

6. To enable the use of a failover multicast address on a VIPRION® platform (recommended), then for the **Use Failover Multicast Address** setting, select the **Enabled** check box.
7. If you enabled **Use Failover Multicast Address**, either accept the default **Address** and **Port** values, or specify values appropriate for the device.
If you revise the default **Address** and **Port** values, but then decide to revert to the default values, click **Reset Defaults**.
8. Click **Update**.

After you perform this task, other devices in the device group can send failover messages to the local device using the specified IP addresses.

Failover addresses for appliance models

Each device or guest in a DSC device group needs specific failover IP addresses configured on it that other devices in the device group can use for failover communication with that device. The type of failover addresses required on each device varies depending on whether the system is provisioned for vCMP.

Table 12: Required and optional unicast IP addresses for failover communication (appliance models only)

vCMP provisioned?	Required failover IP addresses	Optional failover IP addresses
No	<ul style="list-style-type: none"> A unicast static self IP address associated with an internal VLAN (VLAN_{HA} recommended) The unicast static management IP address for the appliance 	Additional unicast static self IP addresses associated with internal VLANs
Yes	<ul style="list-style-type: none"> A unicast static self IP address associated with an internal VLAN (VLAN_{HA} recommended) The unicast virtual management IP address assigned to the guest 	Additional unicast static self IP addresses associated with internal VLANs

Failover addresses for VIPRION models

Each device or guest in a DSC[®] device group needs specific failover IP addresses configured on it that other devices in the device group can use for failover communication with that device. The type of failover addresses required on each device varies depending on whether the system is provisioned for vCMP[®] and whether a multicast failover address is being used.

Table 13: Required and optional unicast and multicast IP addresses for failover communication (VIPRION models only)

vCMP provisioned?	Required failover IP addresses (unicast only, no multicast)	Required failover IP addresses (unicast and multicast)	Optional IP addresses
No	<ul style="list-style-type: none"> A unicast static self IP address associated with an internal VLAN (preferably VLAN_{HA}) The existing, unicast static management IP addresses that you previously configured for all slots in the cluster 	<ul style="list-style-type: none"> A unicast static self IP address associated with an internal VLAN (preferably VLAN_{HA}) A multicast IP address 	Additional unicast static self IP addresses associated with one or more internal VLANs
Yes	<ul style="list-style-type: none"> A unicast static self IP address associated with an internal VLAN (preferably VLAN_{HA}) The existing, unicast virtual static management IP addresses that you previously configured for all slots in the guest's virtual cluster 	<ul style="list-style-type: none"> A unicast static self IP address configured within the guest and associated with a host-based internal VLAN (preferably VLAN_{HA}) A multicast IP address 	Additional unicast static self IP addresses associated with one or more host-based internal VLANs

About traffic groups

A *traffic group* is a collection of related configuration objects, such as a floating self IP address, a virtual IP address, and a SNAT translation address, that run on a BIG-IP® device. Together, these objects process a particular type of application traffic on that device. When a BIG-IP device becomes unavailable, a traffic group floats (that is, fails over) to another device in a device group to ensure that application traffic continues to be processed with little to no interruption in service. In general, a traffic group ensures that when a device becomes unavailable, all of the failover objects in the traffic group fail over to any one of the available devices in the device group.

A traffic group is initially active on the device on which you create it, until the traffic group fails over to another device. For example, if you initially create three traffic groups on Device A, these traffic groups remain active on Device A until one or more traffic groups fail over to another device. If you want an active traffic group to become active on a different device in the device group when failover has not occurred, you can intentionally force the traffic group to switch to a standby state, thereby causing failover to another device.

Only objects with floating IP addresses can be members of a floating traffic group.

An example of a set of objects in a traffic group is an iApps® application service. If a device with this traffic group is a member of a device group, and the device becomes unavailable, the traffic group floats to another member of the device group, and that member becomes the device that processes the application traffic.

Note: A Sync-Failover device group can support a maximum of 15 floating traffic groups.

Failover objects and traffic group association

Any traffic group that you explicitly create on the BIG-IP® system is a floating traffic group. The types of configuration objects that you can associate with a floating traffic group are:

- Virtual IP addresses
- NATs
- SNAT translation addresses
- Self IP addresses
- Folders (such as an iApps® folder)

You can associate configuration objects with a traffic group in these ways:

- You can rely on the folders in which the objects reside to inherit the traffic group that you assign to the `root` folder.
- You can use the BIG-IP Configuration utility to directly associate a traffic group with an iApp application service, a virtual IP address, a NAT or SNAT translation address, or a floating self IP address.
- You can use the BIG-IP® Configuration utility to directly associate a traffic group with a folder.

Important: By default, floating objects that you create with the BIG-IP Configuration utility are associated with `traffic-group-1`. Non-floating objects are associated with `traffic-group-local-only`. You can change these associations by using the BIG-IP Configuration utility to change the traffic group that is associated with each floating IP address on the system.

Note: The only non-floating traffic group that resides on the system is the default non-floating traffic group named `traffic-group-local-only`.

Pre-configured traffic groups

Each BIG-IP® device contains two pre-configured traffic groups:

- A floating traffic group named `traffic-group-1` initially contains the floating self IP addresses that you configured for VLANs `internal` and `external`, as well as any iApps® application services, virtual IP addresses, NATs, or SNAT translation addresses that you have configured on the device.
- A non-floating traffic group named `traffic-group-local-only` contains the static self IP addresses that you configured for VLANs `internal` and `external`. This traffic group never fails over to another device.

Before you configure a traffic group

The following considerations apply to traffic groups:

- On each device in a Sync-Failover device group, the BIG-IP® system automatically assigns the default floating traffic group name to the `root` and `/Common` folders.
- The BIG-IP system creates all traffic groups in the `/Common` folder, regardless of the partition to which the system is currently set.
- Any traffic group named other than `traffic-group-local-only` is a floating traffic group.
- You can specify a floating traffic group on a folder only when the device group that is set on the folder is a Sync-Failover type of device-group.
- You can set a floating traffic group on only those objects that reside in a folder with a device group of type Sync-Failover.
- Setting the traffic group on a failover object to `traffic-group-local-only` prevents the system from synchronizing that object to other devices in the device group.

Creating a traffic group

If you intend to specify a MAC masquerade address when creating a traffic group, you must first create the address, using an industry-standard method for creating a locally administered MAC address.

Perform this task when you want to create a traffic group for a BIG-IP® device. You can perform this task on any BIG-IP device within the device group, and the traffic group becomes active on the local device.

Important: *This procedure creates a traffic group but does not automatically associate it with failover objects. You associate a traffic group with specific failover objects when you create or modify each object.*

1. On the Main tab, click **Device Management > Traffic Groups**.
2. On the Traffic Group List screen, click **Create**.
3. In the **Name** field, type a name for the new traffic group.
4. In the **Description** field, type a description for the new traffic group.
5. In the **MAC Masquerade Address** field, type a MAC masquerade address.

When you specify a MAC masquerade address, you reduce the risk of dropped connections when failover occurs. This setting is optional.

6. From the **Failover Methods** list, select a failover method:
 - Select **Load Aware** when the device group contains heterogeneous platforms and you want to ensure that a traffic group fails over to the device with the most capacity at the moment that failover occurs.

- Select **HA Order** to cause the traffic group to fail over to the first available device in the **Failover Order** list.
 - Select **HA Group** to cause the BIG-IP system to trigger failover based on an HA health score for the device.
7. If you configured the **Failover Methods** setting with a value of **Load Aware**, then in the **HA Load Factor** field, specify a value that represents the application load for this traffic group relative to other active traffic groups on the local device. Otherwise, skip this step.

***Important:** If you configure this setting, you must configure the setting on every traffic group in the device group.*

8. If you configured the **Failover Methods** setting with a value of **HA Order**, then configure the following settings. Otherwise, skip this step.
- a) Select or clear the check box for the **Auto Failback** option.
 - b) If **Auto Failback** is selected, then in the **Auto Failback Timeout** field, type the number of seconds that you want the system to wait before failing back to the specified device. The range in seconds is from 0 to 300. The default is 60. A value of 40 to 60 seconds allows for state mirroring information to be re-mirrored for traffic groups.
 - c) For the **Failover Order** setting, in the **Available** box, select a device name and using the Move button, move the device name to the **Enabled** box. Repeat for each device that you want to include in the ordered list.

This setting is optional. Only devices that are members of the relevant Sync-Failover device group are available for inclusion in the ordered list. If you have enabled the auto-failback feature on the traffic group, ensure that the first device in the ordered list is the device to which you want this traffic group to fail back to when that first device becomes available.

If auto-failback is enabled and the first device in the **Failover Order** list is unavailable, no auto-failback occurs and the traffic group continues to run on the current device. Also, if none of the devices in the list is currently available when failover occurs, the BIG-IP system ignores the **Failover Order** setting and performs load-aware failover instead, using the **HA Load Factor** setting.

9. If you configured the **Failover Methods** setting with a value of **HA Group**, then from the **HA Groups** list, select the HA group that you want to associate with this traffic group.
10. Confirm that the displayed traffic group settings are correct.
11. Click **Finished**.

You now have a floating traffic group with zero members.

After creating the traffic group, you must add members to it. Possible members are floating IP addresses such as self IP addresses, virtual addresses, NAT or SNAT translation addresses, and iApp application services. Also, if you want the traffic group to become active on a device other than this local device, you can use the **Force to Standby** button. By forcing the traffic group into a standby state on the local device, you cause the traffic group to become active on another device.

Adding members to a traffic group

Before performing this task, verify that the traffic group exists on the BIG-IP system.

You perform this task to add members to a newly-created or existing traffic group. Traffic group members are the floating IP addresses associated with application traffic passing through the BIG-IP® system. Typical members of a traffic group are: a floating self IP address, a floating virtual address, and a floating SNAT translation address.

1. From the Main tab, display the properties page for an existing BIG-IP object, such as a self IP address or a virtual address.
For example, from the Main tab, click **Network > Self IPs**, and then from the Self IPs list, click a self IP address.
2. From the **Traffic Group** list, select the floating traffic group that you want the BIG-IP object to join.
3. Click **Update**.

After performing this task, the BIG-IP object belongs to the selected traffic group.

Repeat this task for each BIG-IP object that you want to be a member of the traffic group.

Viewing a list of traffic groups for a device

You can view a list of traffic groups for the device group. Using this list, you can add floating IP addresses to a traffic group, force a traffic group into a Standby state, and view information such as the current and next-active devices for a traffic group and its HA load factor.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Name column, view the names of the traffic groups on the local device.

Viewing the members of a traffic group

You can use the BIG-IP[®] Configuration utility to view a list of all failover objects associated with a specific traffic group. For each failover object, the list shows the name of the object, the type of object, and the folder in which the object resides.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Name column, click the name of the traffic group for which you want to view the associated objects.

This displays a list of all failover objects for the traffic group.

Traffic group properties

This table lists and describes the properties of a traffic group.

Property	Description
Name	The name of the traffic group, such as <code>Traffic-Group-1</code> .
Partition	The name of the folder or sub-folder in which the traffic group resides.
Description	A user-defined description of the traffic group.
MAC Masquerade Address	A user-created MAC address that floats on failover, to minimize ARP communications and dropped connections.
Current Device	The device on which a traffic group is currently running.
Next Active Device	The device currently most available to accept a traffic group if failover of that traffic group should occur.
Floating	An indicator of whether the traffic group includes floating IP addresses as members.

Property	Description
Failover Method	A list of the possible failover methods to configure: Load-Aware, HA Order, and HA Group.
HA Load Factor	A numeric value pertaining to load-aware failover that represents the application traffic load of this traffic group relative to other active traffic groups on the same device.
Auto Failback	The condition where the traffic group tries to fail back to the first device in the ordered failover list, when that device (and that device only) is available.
Auto Failback Timeout	The number of seconds before auto failback expires. This setting appear only when you enable the Auto Failback setting.
Failover Order	An ordered list of devices that the BIG-IP® system uses to determine the next-active device for the traffic group.
HA Group	A list of existing HA groups that you can choose from to enable HA group failover for this traffic group. Note that the HA Group property does not sync to other devices during a config sync operation and pertains only to this local instance of the traffic group.

Active and standby states

During any config sync operation, each traffic group within a device group is synchronized to the other device group members. Therefore, on each device, a particular traffic group is in either an active state or a standby state. In an *active* state, a traffic group on a device processes application traffic. In a *standby* state, a traffic group on a device is idle.

For example, on Device A, traffic-group-1 might be active, and on Device B, traffic-group-1 might be standby. Similarly, on Device B, traffic-group-2 might be active, and on Device A, traffic-group-2 might be standby.

When a device with an active traffic group becomes unavailable, the traffic group floats to (that is, becomes active on) another device. The BIG-IP® system chooses the target device based on how you initially configured the traffic group when you created it. Note that the term *floats* means that on the target device, the traffic group switches from a standby state to an active state.

The following illustration shows a typical device group configuration with two devices and one traffic group (named `my_traffic_group`). In this illustration, the traffic group is active on Device A and standby on Device B prior to failover.




Figure 10: Traffic group states before failover

If failover occurs, the traffic group becomes active on the other device. In the following illustration, Device A has become unavailable, causing the traffic group to become active on Device B and process traffic on that device.




Figure 11: Traffic group states after failover

When `Device A` comes back online, the traffic group becomes standby on `Device A`.

Viewing the failover state of a device

You can use the BIG-IP® Configuration utility to view the current failover state of a device in a device group. An `Active` failover state indicates that at least one traffic group is currently active on the device. A `Standby` failover state indicates that all traffic groups on the device are in a `Standby` state.

1. Display any screen of the BIG-IP Configuration utility.
2. In the upper left corner of the screen, view the failover state of the device.

Viewing the state of a traffic group

You can use the BIG-IP® Configuration utility to view the current state of all traffic groups on the device.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Failover Status area of the screen, view the state of all traffic groups on the device.

Forcing a traffic group to a standby state

You perform this task when you want the selected traffic group on the local device to fail over to another device (that is, switch to a `Standby` state). Users typically perform this task when no automated method is configured for a traffic group, such as auto-failback or an HA group. By forcing the traffic group into a `Standby` state, the traffic group becomes active on another device in the device group. For device groups with more than two members, you can choose the specific device to which the traffic group fails over.

1. Log in to the device on which the traffic group is currently active.
2. On the Main tab, click **Device Management > Traffic Groups**.
3. In the Name column, locate the name of the traffic group that you want to run on the peer device.
4. Select the check box to the left of the traffic group name.
If the check box is unavailable, the traffic group is not active on the device to which you are currently logged in. Perform this task on the device on which the traffic group is active.
5. Click **Force to Standby**.
This displays target device options.
6. Choose one of these actions:
 - If the device group has two members only, click **Force to Standby**. This displays the list of traffic groups for the device group and causes the local device to appear in the Next Active Device column.

- If the device group has more than two members, then from the **Target Device** list, select a value and click **Force to Standby**.

The selected traffic group is now in a standby state on the local device and active on another device in the device group.

About active-standby vs. active-active configurations

A device group that contains only one traffic group is known as an *active-standby* configuration.

A device group that contains two or more traffic groups is known as an *active-active* configuration. For example, if you configure multiple virtual IP addresses on the BIG-IP® system to process traffic for different applications, you might want to create separate traffic groups that each contains a virtual IP address and its relevant floating self IP address. You can then choose to make all of the traffic groups active on one device in the device group, or you can balance the traffic group load by making some of the traffic groups active on other devices in the device group.

Description of current and next-active devices

Within a Sync-Failover type of device group, a BIG-IP® device sometimes has a specific designation with respect to a traffic group: either a current device or a next-active device.

Table 14: Current and next-active devices described

Designation	Description
Current Device	A <i>current</i> device is the device on which a traffic group is currently active. For example, if Device A is currently processing traffic using the objects in Traffic-Group-1, then Device A is the current device. If Device A becomes unavailable and Traffic-Group-1 fails over to Device C, then Device C becomes the current device.
Next-Active Device	A <i>next-active</i> device is the device on which a traffic group will become active if that traffic group eventually fails over to another device. For every active traffic group, the BIG-IP system assigns a corresponding next-active device. The next-active device for a traffic group is system-selected, based on criteria you specify when you configure the traffic group.

About the next-active device

For every active traffic group on a device, the BIG-IP® system identifies the device that is to be the next-active device if failover of that active traffic group occurs. A *next-active* device is the device on which a traffic group will become active if that traffic group eventually fails over to another device. This next-active designation changes continually depending on which devices are currently available in the device group.

There are various configuration options for you to choose from to affect the BIG-IP system's selection of the next-active device for failover:

- Load-aware failover

- An ordered list with auto-failback
- HA groups

What is load-aware failover?

Load-aware failover is a BIG-IP® feature designed for use in a Sync-Failover device group. Configuring *load-aware failover* ensures that the traffic load on all devices in a device group is as equivalent as possible, factoring in any differences in device capacity and the amount of application traffic that traffic groups process on a device. The load-aware configuration option is most useful for device groups with heterogeneous hardware platforms or varying application traffic loads (or both).

For example, suppose you have a heterogeneous three-member device group in which one device (`Bigip_C`) has twice the hardware capacity of the other two devices (`Bigip_A` and `Bigip_B`).

If the device group has four active traffic groups that each process the same amount of application traffic, then the load on all devices is equivalent when devices `Bigip_A` and `Bigip_B` each contain one active traffic group, while device `Bigip_C` contains two active traffic groups.

The BIG-IP system implements load-aware failover by calculating a numeric, current utilization score for each device, based on numeric values that you specify for each device and traffic group relative to the other devices and traffic groups in the device group. The system then uses this current utilization score to determine which device is the best device in the group to become the next-active device when failover occurs for a traffic group.

The overall result is that the traffic load on each device is as equivalent as possible in a relative way, that is, factoring in individual device capacity and application traffic load per traffic group.

About device utilization calculation

The BIG-IP® system on each device performs a calculation to determine the device's current level of utilization. This utilization level indicates the ability for the device to be the next-active device in the event that an active traffic group on another device must fail over within a heterogeneous device group.

The calculation that the BIG-IP performs to determine the current utilization of a device is based on these factors:

Device capacity

A local device capacity relative to other device group members.

Active local traffic groups

The number of active traffic groups on the local device.

Active remote traffic groups

The number of remote active traffic groups for which the local device is the next-active device.

A multiplying load factor for each active traffic group

A multiplier value for each traffic group. The system uses this value to weight each active traffic group's traffic load compared to the traffic load of each of the other active traffic groups in the device group.

The BIG-IP system uses all of these factors to perform a calculation to determine, at any particular moment, a score for each device that represents the current utilization of that device. This utilization score indicates whether the BIG-IP system should, in its attempt to equalize traffic load on all devices, designate the device as a next-active device for an active traffic group on another device in the device group.

The calculation that the BIG-IP performs for each device is:

$$\frac{(\text{The sum of all local active traffic group HA load factors} + \text{The sum of all remote active traffic group HA load factors})}{\text{device capacity}}$$

About HA capacity

For each device in a BIG-IP® device group, you can assign a high availability (HA) capacity value. An *HA capacity* value is a number that represents the relative processing capacity of that device compared to the other devices in a device group. Assigning different HA capacity values to the devices in the device group is useful when the device group contains heterogeneous hardware platforms.

For example, if the device group has two devices with equal capacity and a third device that has twice the capacity of each of the other two devices, then you can assign values of 2, 2, and 4, respectively. You can assign any number to represent the HA capacity, as long as the number reflects the device's relative capacity compared to the other devices in the device group.

Specifying the HA capacity of a device

Before you perform this task, verify that this device is a member of a device group and that the device group contains three or more devices.

You perform this task when you have more than one type of hardware platform in a device group and you want to configure load-aware failover. *Load-aware failover* ensures that the BIG-IP® system can intelligently select the next-active device for each active traffic group in the device group when failover occurs. As part of configuring load-aware failover, you define an HA capacity to establish the amount of computing resource that the device provides relative to other devices in the device group.

Note: *If all devices in the device group are the same hardware platform, you can skip this task.*

1. On the Main tab, click **Device Management > Devices**.

This displays a list of device objects discovered by the local device.

2. In the Name column, click the name of the device for which you want to view properties.

This displays a table of properties for the device.

3. In the **HA Capacity** field, type a relative numeric value.

You need to configure this setting only when you have varying types of hardware platforms in a device group and you want to configure load-aware failover. The value you specify represents the relative capacity of the device to process application traffic compared to the other devices in the device group.

Important: *If you configure this setting, you must configure the setting on every device in the device group.*

If this device has half the capacity of a second device and a third of the capacity of a third device in the device group, you can specify a value of 100 for this device, 200 for the second device, and 300 for the third device.

When choosing the next active device for a traffic group, the system considers the capacity that you specified for this device.

4. Click **Update**.

After you perform this task, the BIG-IP system uses the **HA Capacity** value to calculate the current utilization of the local device, to determine the next-active device for failover of other traffic groups in the device group.

About the HA load factor

For each traffic group on a BIG-IP® device, you can assign an high availability (HA) load factor. An *HA load factor* is a number that represents the relative application traffic load that an active traffic group processes compared to other active traffic groups in the device group.

For example, if the device group has two active traffic groups, and one traffic group processes twice the amount of application traffic as the other, then you can assign values of 4 and 2, respectively. You can assign any number for the HA load factor, as long as the number reflects the traffic group's relative load compared to the other active traffic groups.

About metrics for the HA load factor

User-specified values for the HA load factor can be based on different metrics. For example, suppose you have the three devices `Bigip_A`, `Bigip_B`, and `Bigip_C`, and each device has one active traffic group with an HA load factor of 2, 4, or 8 respectively. These values could indicate either of the following:

- If each traffic group contains one virtual address, then the sample factor values could indicate that the virtual server for `Bigip_B` processes twice the amount of traffic as that of `Bigip_A`, and the virtual server for `Bigip_C` processes twice the amount of traffic as that of `Bigip_B`.
- If the traffic group on `Bigip_A` contains one virtual address, the traffic group on `Bigip_B` contains two virtual addresses, and the traffic group on `Bigip_C` contains four virtual addresses, this could indicate that the virtual servers corresponding to those virtual addresses each process the same amount of traffic compared to the others.

Specifying an HA load factor for a traffic group

You perform this task when you want to specify the relative application load for an existing traffic group, for the purpose of configuring load-aware failover. *Load-aware failover* ensures that the BIG-IP® system can intelligently select the next-active device for each active traffic group in the device group when failover occurs. When you configure load-aware failover, you define an application traffic load (known as an *HA load factor*) for a traffic group to establish the amount of computing resource that an active traffic group uses relative to other active traffic groups.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Name column, click the name of a traffic group.
This displays the properties of the traffic group.
3. From the **Failover Methods** list, select **Load Aware**.
This displays the **HA Load Factor** setting.
4. In the **HA Load Factor** field, specify a value that represents the application load for this traffic group relative to other active traffic groups on the local device.

Important: *If you configure this setting, you must configure the setting on every traffic group in the device group.*

5. Click **Update**.

After performing this task, the BIG-IP system uses the **HA Load Factor** value as a factor in calculating the current utilization of the local device, to determine whether this device should be the next-active device for failover of other traffic groups in the device group.

Viewing an HA traffic factor summary

Before performing this task, verify that you have configured load-aware failover for the Sync-Failover device group.

You can perform this task to view statistics about load-aware failover. More specifically, you can view the specified traffic load (HA traffic factor) for each traffic group on a current device and for each traffic group on a next-active device.

1. Open a console window on a device in the Sync-Failover device group.
2. At the `tmsh` command line prompt, type this command:

```
tmsh show cm traffic-group all-properties
```

Name	Device	Status	Next Active	Load	Next Active Load	Times Became Active	Last Became Active
tg-1	bigip_B	standby	true	-	1	0	-
tg-1	bigip_A	active	false	1	-	4	15:57:24
tg-2	bigip_B	active	false	4	-	1	11:11:45
tg-2	bigip_A	standby	true	-	4	1	11:08:49
tg-local-only							

This `tmsh` output shows a sample device group with two members and a total of two traffic groups. The user has configured a traffic load value (HA traffic factor) of 1 for `tg-1` and 4 for `tg-2`. Both devices are the same hardware platform and so have a relative HA capacity value set to 0 (not shown).

For each device, the BIG-IP® system determines the overall utilization score by adding together the loads for the active traffic groups. Thus:

- For `bigip_A`, the system adds together traffic loads of 1 (for its active traffic group) and 4 (if `tg-2` fails over to `bigip_A`).
- For `bigip_B`, the system adds together traffic loads of 4 (for its active traffic group) and 1 (if `tg-1` fails over to `bigip_B`).

In this example, because there are only two devices with one active traffic group each, the two devices have the same utilization score (not shown), which is 5.

Examples of device utilization scores

The utilization scores that the BIG-IP system calculates for the devices in a device group vary depending on:

- The differences in hardware capacity of the device group members
- The application load on each traffic group
- The number of active traffic groups on each device

Homogeneous device group with equivalent traffic group loads

In this example, all devices are the same hardware platform, and all three active traffic groups process equivalent application traffic load. Because the load is equivalent for all three traffic groups, the configured HA load factor for each traffic group is the same (in this case, 1).

The device utilization that the BIG-IP® system calculates in this example is the sum of the two traffic load values (one per active traffic group).

Table 15: Calculating the utilization score for Bigip_A

HA Capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
0	Traffic-group-1	1	Traffic-group-2	1	1 + 1 = 2

Table 16: Calculating the utilization score for Bigip_B

HA Capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
0	Traffic-group-2	1	Traffic-group-3	1	$1 + 1 = 2$

Table 17: Calculating the utilization score for Bigip_C

HA Capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
0	Traffic-group-3	1	Traffic-group-1	1	$1 + 1 = 2$

Homogeneous device group with disparate traffic group loads

In this example, all devices are the same hardware platform. Also, each user-specified traffic group load factor is defined simply as 1, 4, or 8, to indicate that `traffic-group-2` processes four times the application load of `traffic-group-1`, and that `traffic-group-3` processes twice the application load of `traffic-group-2`.

The device utilization that the BIG-IP® system calculates in this example is the sum of the two traffic load values (one per active traffic group).

Table 18: Calculating the utilization score for Bigip_A

HA capacity	Active traffic group	Traffic load	Potential traffic group	Traffic load	Device utilization score
0	Traffic-group-1	1	Traffic-group-2	4	$1 + 4 = 5$

Table 19: Calculating the utilization score for Bigip_B

HA capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
0	Traffic-group-2	4	Traffic-group-3	8	$4 + 8 = 12$

Table 20: Calculating the utilization score for Bigip_C

HA capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
0	Traffic-group-3	8	Traffic-group-1	1	$8 + 1 = 9$

This example shows that device `Bigip_A` is currently the least-used device, with a score of 5, while `Bigip_B` is the most used, with a score of 12. Therefore, the BIG-IP system would currently choose `Bigip_A` to receive failover traffic, to ensure that the application traffic load remains as equivalent as possible on all devices.

Heterogeneous device group with disparate traffic group loads

In this example, the load-aware configuration consists of a user-specified relative capacity for each device and a relative load for each active traffic group. The device group contains three heterogeneous devices, each with one active traffic group. Being different hardware platforms, the three devices each have a different user-specified relative device capacity, and each traffic group on a device has a different application traffic load.

The device utilization score that the BIG-IP® system calculates in this example is the sum of two traffic load values on a device divided by the device capacity.

Table 21: Calculating the utilization score for Bigip_A

HA capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
10	Traffic-group-1	1	Traffic-group-2	4	5/10 = .50

Table 22: Calculating the utilization score for Bigip_B

HA capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
80	Traffic-group-2	4	Traffic-group-3	8	12/80 = .15

Table 23: Calculating the utilization score for Bigip_C

HA capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
20	Traffic-group-3	8	Traffic-group-1	1	9/20 = .45

This example shows the results of the calculations that the BIG-IP system performs for each device in the device group. For each device, the BIG-IP system factors in the device capacity, the load of the device's active traffic group, and the load of the next-active traffic group, to determine the current utilization of that device. The example shows that device `Bigip_B`, with a utilization score of .15, has the most available resource, despite having the heaviest traffic load. This is due to the large device capacity of 80 that the user specified relative to the other devices. `Bigip_B` is therefore most able to accept failover traffic from another device.

Heterogeneous device group with multiple active traffic groups on a single device

In this example, the load-aware configuration consists of a user-specified relative high availability (HA) capacity for each device and relative load for each active traffic group. The device group contains three heterogeneous devices, where `Bigip_A` and `Bigip_B` currently have one active traffic group each, while `Bigip_C` has two active traffic groups. Being different hardware platforms, the three devices each have a different user-specified relative device capacity, and each traffic group has a different relative application traffic load.

The device utilization score that the BIG-IP® system calculates in this example is the sum of all traffic load values on a device divided by the device capacity.

Table 24: Calculating the utilization score for Bigip_A

HA capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
10	Traffic-group-1	1	Traffic-group-2	4	5/10 = .5

Table 25: Calculating the utilization score for Bigip_B

HA capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
80	Traffic-group-2	4	Traffic-group-3	8	12/80 = .15

Table 26: Calculating the utilization score for Bigip_C

HA capacity	Active traffic group	Traffic load	Potential active traffic group	Traffic load	Device utilization score
20	Traffic-group-3 and Traffic-group-4	8 and 6	Traffic-group-1	1	15/20 = .75

This example shows the results of the calculations that the BIG-IP system performs for each device in the device group. The example shows that device `Bigip_B` has the most available resource due to its low utilization score of .15. Conversely, `Bigip_C` has the highest utilization score (.75), due to having an additional active traffic group (`Traffic-group-4`) on the device with a relatively high traffic load value (6). In this case, `Bigip_C` is unlikely to become the next-active device on failover.

About matching device utilization values

In rare cases, the BIG-IP® system might calculate that two or more devices in a device group have the same lowest device utilization score. In this case, the BIG-IP system needs an additional method for choosing the next-active device for an active traffic group.

The way that the BIG-IP system chooses the next-active device when device utilization scores match is by determining the management IP address of each matching device and then calculating a score based on the highest management IP address of those devices.

For example, if `Bigip_A` has an IP address of `192.168.20.11` and `Bigip_B` has an IP address of `192.168.20.12`, and their utilization scores match, the BIG-IP system calculates a score based on the address `192.168.20.12`.

What is an HA ordered list?

If you do not want to use the load-aware feature to determine the next-active device for a traffic group, you can configure a traffic group to use a static, ordered list of devices instead. This list of devices specifies the order in which you want those devices to become active for the traffic group if the traffic group must fail over.

If failover occurs and the first device in the list is unavailable, the BIG-IP® system tries to make the traffic group active on the second device in the list. If the second device is also unavailable, the BIG-IP system tries to make the traffic group active on the third device, and so on.

If you do not specify an ordered list or if none of the devices in the list is available, the BIG-IP system attempts to use load-aware failover to choose the next-active device.

***Note:** When you enable the auto-failback feature for a traffic group, the BIG-IP system tries to ensure that the traffic group is always active on the first device in the list. If the first device in the list is unavailable, no fail-back occurs.*

About auto-failback

The failover feature includes an option known as auto-failback. When you enable *auto-failback*, a traffic group that has failed over to another device fails back to a preferred device when that device is available. If you do not enable auto-failback for a traffic group, and the traffic group fails over to another device, the traffic group remains active on that device until that device becomes unavailable.

You can enable auto-failback on a traffic group only when you have configured an ordered list with at least one entry, for that traffic group. In this case, if auto-failback is enabled and the traffic group has failed over to another device, then the traffic group fails back to the first device in the traffic group's ordered list (the preferred device) when that device becomes available.

Important: *If the first device in the ordered list is unavailable, no fail-back occurs. The traffic group does not fail back to the next available device in the list and instead remains on its current device.*

Creating an HA ordered list

You perform this task to create a prioritized, ordered list on an existing traffic group. The BIG-IP® system uses this list to determine the next-active device for this traffic group. This configuration option is most useful for device groups with homogeneous hardware platforms and similar application traffic loads, or for applications that require a specific target failover device, such as those that use connection mirroring. When failover occurs, the traffic group fails over to the first available device in the list.

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Name column, click the name of a traffic group on the local device.
This displays the properties of the traffic group.
3. From the **Failover Methods** list, select **HA Order**.
4. Select or clear the check box for the **Auto Failback** option:
 - Select the check box to cause the traffic group, after failover, to fail over again to the first device in the traffic group's ordered list when that device (and only that device) is available.
 - Clear the check box to cause the traffic group, after failover, to remain active on its current device until failover occurs again.
5. If auto-failback is enabled, in the **Auto Failback Timeout** field, type the number of seconds that you want the system to wait before failing back to the specified device. The range is from 0 to 300 seconds. The default is 60. A value of 40 to 60 allows for state mirroring information to be re-mirrored for traffic groups.
6. For the **Failover Order** setting, in the **Available** box, select a device name and using the Move button, move the device name to the **Enabled** box. Repeat for each device that you want to include in the ordered list.
This setting is optional. Only devices that are members of the relevant Sync-Failover device group are available for inclusion in the ordered list. If you have enabled the auto-failback feature on the traffic group, ensure that the first device in the ordered list is the device to which you want this traffic group to fail back to when that first device becomes available.
If auto-failback is enabled and the first device in the **Failover Order** list is unavailable, no auto-failback occurs and the traffic group continues to run on the current device. Also, if none of the devices in the **Failover Order** list is currently available when failover occurs, the BIG-IP system ignores the **Failover Order** setting and performs load-aware failover instead, using the **HA Load Factor** setting.
7. Click **Update**.

After you perform this task, the BIG-IP system designates the first available device that is highest in the ordered list as the next-active device for the traffic group.

What is an HA group?

An *HA group* is a BIG-IP configuration object that specifies a set of trunks, pools, and/or VIPRION® clusters for a device. When you associate an HA group with an instance of a traffic group on a specific device, the

BIG-IP system calculates an overall health score on that device for the associated traffic group, based on trunk, pool, or cluster availability. This *HA health score* determines the device that the associated traffic group should be active on at any given time and then triggers failover to that device if necessary. Trunk configurations are not synchronized between devices, which means that the number of trunk members on each device in the device group can vary whenever a trunk loses or gains members, thereby affecting the health score.

The ability of an HA group to choose a next-active device and trigger failover for a traffic group distinguishes the HA groups feature from the other failover methods (load-aware failover and HA order). Although the load-aware and HA order failover methods also identify the next-active device for a traffic group, these methods do not trigger failover; instead, failover is triggered by standard system, gateway, or VLAN fail-safe configuration.

The most common reason for configuring HA groups is to ensure that a traffic group fails over to another device when trunk availability for a specific device falls below a certain level. Also, failover triggered by an HA group is noticeably faster than failover triggered by system, gateway, or VLAN failure.

As an example, suppose that `Bigip_A` and `Bigip_B` each has the traffic group `traffic-group-1`. If you want `traffic-group-1` to always be active on the device with the highest HA score, you can create a unique HA group for that traffic group, on each device, that defines a minimum value for acceptable link availability. You can then associate the local HA group with `traffic-group-1` on each device. More specifically, on `Bigip_A`, you can create an HA group named `HA_GroupA` and associate it with `traffic-group-1`. On `Bigip_B`, you can create a different HA group, `HA_GroupB`, and also associate it with `traffic-group-1`. This ensures that `traffic-group-1` is always active on the device with the best link availability at any given time.

Important: Whenever you associate an HA group with a traffic group on a device, you must associate an HA group with that same traffic group on every device in the device group. For example, if the device group contains three devices and you associate an HA group with `traffic-group-1` on device `Bigip_A`, then you must also associate an HA group with `traffic-group-1` on devices `Bigip_B` and `Bigip_C`. In a typical configuration, the values of each HA group differ for the traffic group on each device.

Factors in HA health score calculation

The BIG-IP® system calculates an HA health score per traffic group on a device, based on the values that you specify when you configure the feature.

A weight value (required)

A *weight* is a health value that you assign to each member of the HA group (that is, a pool, trunk, and/or VIPRION® cluster). The weight that you assign to each HA group member must be in the range of 10 through 100.

The maximum overall score that the BIG-IP system can potentially calculate is the sum of the individual weights for the HA group members, plus the active bonus value. There is no limit to the sum of the member weights for the HA group as a whole.

A threshold value (optional)

For each member of an HA group, you can specify a setting known as a threshold. A *threshold* is a value that specifies the number of object members that must be available to prevent failover. If the number of available object members is less than the threshold, the BIG-IP system assigns a score of 0 to the HA group member so that the score of that HA group member no longer contributes to the overall score.

For example, if a trunk in the HA group has four trunk members and you specify a threshold value of 3, and the number of available trunk members falls to 2, then the trunk contributes a score of 0 to the total score for the traffic group or device.

If the number of available object members equals or exceeds the threshold value, or you do not specify a threshold, the BIG-IP system calculates the score as described previously, by multiplying the percentage of available object members by the weight for each HA group member and then adding the scores to determine the overall score for the traffic group or device.

The threshold that you define for pools can be less than or equal to the number of members in the pool. For clusters, the threshold can be less than or equal to the number of possible blades in the chassis, and for trunks, the threshold can be less than or equal to the number of possible members in a trunk for that platform.

Tip: Do not configure the `tmsh` attribute `min-up-members` on any pool that you intend to include in the HA group.

An active bonus value (optional)

An *active bonus* is an amount that the BIG-IP system automatically adds to the overall HA score of an active traffic group or device. An active bonus ensures that the traffic group or device remains active when its score would otherwise temporarily fall below the score of the standby traffic group on another device. The active bonus that you configure can be in the range of 0 to 100.

A common reason to specify an active bonus is to prevent failover due to *flapping*, the condition where failover occurs frequently as a trunk member switches rapidly between availability and unavailability. In this case, you might want to prevent the HA scoring feature from triggering failover each time a trunk member is lost. You might also want to prevent the HA scoring feature from triggering failover when you make minor changes to the BIG-IP system configuration, such as adding or removing a trunk member.

For example, suppose that the HA group for a traffic group on each device contains a trunk with four members, and you assign a weight of 30 to each trunk. Without an active bonus defined, if the trunk on one device loses some number of members, failover occurs because the overall calculated score for that traffic group becomes lower than that of a peer device. You can prevent this failover from occurring by specifying an active bonus value.

The BIG-IP system uses an active bonus to contribute to the HA score of an active traffic group only; the BIG-IP system never uses an active bonus to contribute to the score of a standby traffic group.

Note: An exception to this behavior is when the active traffic group score is 0. In this case, the system does not add the active bonus to the active traffic group or active device score.

To decide on an active bonus value, calculate the trunk score for some number of failed members (such as one of four members), and then specify an active bonus that results in a trunk score that is greater than or equal to the weight that you assigned to the trunk.

For example, if you assigned a weight of 30 to the trunk, and one of the four trunk members fails, the trunk score becomes 23 (75% of 30), putting the traffic group at risk for failover. However, if you specified an active bonus of 8 or higher, failover would not actually occur, because a score of 8 or higher, when added to the score of 23, is greater than 30.

Example of HA health score calculation

This example illustrates the way that HA group configuration results in the calculation of an HA health score for a traffic group on a specific device. Suppose that you previously created an HA group for `traffic-group-1` on all device group members and that `traffic-group-1` is currently active on device `Bigip_A`. Also suppose that on device `Bigip_B`, the HA group for `traffic-group-1` consists of two pools and a trunk, with weights that you assign:

Table 27: Sample HA group configuration for traffic-group-1 on Bigip_B

HA group object	Member count	User-specified weight
http_pool	8	50
ftp_pool	6	20
trunk1	4	30

Now suppose that on device Bigip_B, the current member availability of pool http_pool, pool ftp_pool, and trunk trunk1 is 5, 6, and 3, respectively. The resulting HA score that the BIG-IP system calculates for traffic-group-1 on Bigip_B is shown here:

Table 28: Sample health score calculation for traffic-group-1 on Bigip_B

HA group object	Member count	Available member count	User-specified weight	Current HA score
http_pool	8	5 (62.5%)	50	31 (60% x 50)
ftp_pool	6	6 (100%)	20	20 (100% x 20)
trunk1	4	3 (75%)	30	23 (75% x 30)
				Total score: 74

In this example, the total HA score for traffic-group-1 on Bigip_B is currently 74. If this score is currently the highest score in the device group for traffic-group-1, then traffic-group-1 will automatically failover and become active on Bigip_B.

Creating an HA group

You use this task to create an HA group for a traffic group on a device in a BIG-IP® device group. Also known as *fast failover*, an HA group is most useful when you want an active traffic group on a device to fail over to another device based on trunk and pool availability, and on VIPRION® systems, also cluster availability. You can create multiple HA groups on a single device, and you associate each HA group with the local instance of a traffic group.

Important: Once you create an HA group on a device and associate the HA group with a traffic group, you must create an HA group and associate it with that same traffic group on every device in the device group. For example, on Device_A, if you create HA_GroupA_TG1 and associate it with traffic-group-1, then on Device_B you can create HA_GroupB_TG1 and also associate it with traffic-group-1.

1. On the Main tab, click **System > High Availability > HA Groups**
2. In the **HA Group Name** field, type a name for the HA group, such as ha_group1.
3. Verify that the **Enable** check box is selected.
4. In the **Active Bonus** field, specify an integer that represents the amount by which you want the system to increase the overall score of the active device.
The purpose of the active bonus is to prevent failover when minor or frequent changes occur to the configuration of a pool, trunk, or cluster.
5. In the table displayed along the bottom of the screen, for the **Threshold** setting, for each pool, trunk, or VIPRION cluster in the HA group, optionally specify an integer for a threshold value.

6. For the **Weight** setting, for each pool, trunk, or VIPRION cluster in the HA group, specify an integer for the weight. The allowed weight for an HA group object ranges from 10 through 100.
This value is required.
7. Click **Create**.

You now have an HA group that the BIG-IP system can later use to calculate an HA score for fast failover. After creating an HA group on the local device, you must assign it to a traffic group on the local device.

Associating an HA group with a traffic group

You use this task to associate an HA group with an existing traffic group. Also known as *fast failover*, this configuration option is most useful when you want an active traffic group to fail over to another device due to trunk, pool, and/or VIPRION® cluster availability specifically. When you configure an HA group for a traffic group, you ensure that the traffic group, when active, fails over to the device on which the traffic group has the highest HA score.

Important: *HA groups are not included in config sync operations. For this reason, you must create a separate HA group on every device in the device group for this traffic group. For example, if the device group contains three devices and you want to create an HA group for `traffic-group-1`, you must configure the HA group property for `traffic-group-1` on each of the three devices separately. In a typical device group configuration, the values of the HA group settings on the traffic group will differ on each device.*

1. On the Main tab, click **Device Management > Traffic Groups**.
2. In the Name column, click the name of a traffic group on the local device.
This displays the properties of the traffic group.
3. From the **Failover Methods** list, select **HA Group**.
4. From the **HA Group** list, select an HA group.
5. Click **Update**.

After you perform this task for this traffic group on each device group member, the BIG-IP system ensures that this traffic group is always active on the device with the highest HA score.

About MAC masquerade addresses

A *MAC masquerade address* is a unique, floating Media Access Control (MAC) address that you create and control. You can assign one MAC masquerade address to each traffic group on a BIG-IP® device. By assigning a MAC masquerade address to a traffic group, you indirectly associate that address with any floating IP addresses (services) associated with that traffic group. With a MAC masquerade address per traffic group, a single VLAN can potentially carry traffic and services for multiple traffic groups, with each service having its own MAC masquerade address.

A primary purpose of a MAC masquerade address is to minimize ARP communications or dropped packets as a result of a failover event. A MAC masquerade address ensures that any traffic destined for the relevant traffic group reaches an available device after failover has occurred, because the MAC masquerade address floats to the available device along with the traffic group. Without a MAC masquerade address, on failover the sending host must relearn the MAC address for the newly-active device, either by sending an ARP request for the IP address for the traffic or by relying on the gratuitous ARP from the newly-active device to refresh its stale ARP entry.

The assignment of a MAC masquerade address to a traffic group is optional. Also, there is no requirement for a MAC masquerade address to reside in the same MAC address space as that of the BIG-IP device.

***Note:** When you assign a MAC masquerade address to a traffic group, the BIG-IP system sends a gratuitous ARP to notify other hosts on the network of the new address.*

Chapter 9

Managing Connection Mirroring

- *About connection and persistence mirroring*
- *Connection mirroring considerations*
- *Configuration task summary*

About connection and persistence mirroring

BIG-IP® system redundancy includes the ability for a device to mirror connection and persistence information to another device, to prevent interruption in service during failover. The BIG-IP system mirrors connection and persistence data over TCP port 1028 with every packet or flow state update.

Important: *Connection mirroring only functions between devices that reside on identical hardware platforms.*

Connection mirroring operates at the traffic group level. That is, each device in a device group has a specific mirroring peer device for each traffic group. The mirroring peer device is the traffic group's next-active device.

For example, if device `Bigip_A` is active for traffic group `traffic-group-1`, and the next-active device for that traffic group is `Bigip_C`, then the traffic group on the active device mirrors its in-process connections to `traffic-group-1` on `Bigip_C`.

If `Bigip_A` becomes unavailable and failover occurs, `traffic-group-1` on `Bigip_C` becomes active and continues the processing of any current connections.

Note: *The BIG-IP system can mirror connections for as many as 15 active traffic groups simultaneously.*

Connection mirroring considerations

You should enable connection mirroring whenever failover would cause a user session to be lost or significantly disrupted.

For example, long-term connections such as FTP and Telnet are good candidates for mirroring. For this type of traffic, if failover occurs, an entire session can be lost if the connections are not being mirrored to a peer device.

Conversely, the mirroring of short-term connections such as HTTP and UDP is not recommended, because these protocols allow for failure of individual requests without loss of the entire session, and the mirroring of short-term connections can negatively impact system performance.

Connection mirroring only functions between devices that reside on identical hardware platforms.

Note: For vCMP guests on a VIPRION system to function as mirrored pairs, the guests must be configured identically. That is, they must have the same **Cores Per Slots** setting and reside in the same slots on their respective VIPRION chassis. Additionally, on a single guest, the values of the **Minimum Number of Slots** and **Number of Slots** settings must match, and these values must match those of the mirroring peer.

Configuration task summary

Configuring connection mirroring requires you to perform these specific tasks:

Specifying a local self IP address for connection mirroring (required)

This local self IP address is the address that you want other devices in a device group to use when other traffic groups mirror their connections to a traffic group on this device.

Enabling connection mirroring on a virtual server

The BIG-IP® can mirror TCP or UDP connections for a virtual server. When you enable connection mirroring on a virtual server, and you then make the relevant virtual address a member of an active floating traffic group, the traffic group can mirror its connections to its corresponding standby traffic group on another device.

Enabling connection mirroring on a SNAT

The BIG-IP system can mirror TCP or UDP connections for a SNAT.

Enabling persistence mirroring on a persistence profile

The BIG-IP system can mirror persistence information between peers for the following persistence profiles:

- Destination address affinity
- Hash
- Microsoft Remote Desktop (MSRDP)
- Session Initiation Protocol (SIP)
- Source address affinity
- SSL
- Universal

Specifying an IP address for connection mirroring

You can specify the local self IP address that you want other devices in a device group to use when mirroring their connections to this device. Connection mirroring ensures that in-process connections for an active traffic group are not dropped when failover occurs. You typically perform this task when you initially set up device service clustering (DSC®).

Note: You must perform this task locally on each device in the device group.

1. Confirm that you are logged in to the actual device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Mirroring.

- For the **Primary Local Mirror Address** setting, retain the displayed IP address or select another address from the list.

The recommended IP address is the self IP address for either VLAN `HA` or VLAN `internal`.

Important: *If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the self IP address you specify must be one of the private IP addresses that you configured for this EC2 instance as the **Primary Local Mirror Address**.*

- For the **Secondary Local Mirror Address** setting, retain the default value of **None**, or select an address from the list.

This setting is optional. The system uses the selected IP address in the event that the primary mirroring address becomes unavailable.

- Click **Update**.

In addition to specifying an IP address for mirroring, you must also enable connection mirroring on the relevant virtual servers on this device.

Enabling connection mirroring for TCP and UDP connections

Verify that you have specified primary and secondary mirroring IP addresses on this device. Other traffic groups in the device group use these addresses when mirroring connections to this device.

You can perform this task to enable TCP or UDP connections for a virtual server. *Connection mirroring* is an optional feature of the BIG-IP® system, designed to ensure that when failover occurs, in-process connections are not dropped. You enable mirroring for each virtual server that is associated with a floating virtual address.

- On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
- Click the name of the virtual server you want to modify.
- From the **Configuration** list, select **Advanced**.
- For the **Connection Mirroring** setting, select the check box.
- Click **Update** to save the changes.

Enabling connection mirroring for SNAT connections

You can perform this task to enable connection mirroring for source network address translation (SNAT). *Connection mirroring* is an optional feature of the BIG-IP® system, designed to ensure that when failover occurs, in-process SNAT connections are not dropped. You can enable mirroring on each SNAT that is associated with a floating virtual address.

- On the Main tab, click **Local Traffic > Address Translation**.
The **SNAT List** screen displays a list of existing SNATs.
- In the Name column, click the relevant SNAT name.
- For the **Stateful Failover Mirror** setting, select the check box.
- Click **Update**.

In addition to enabling connection mirroring on a SNAT, you must also specify a mirroring IP address on this device. Other traffic groups in the device group use this address when mirroring their connections to this device.

Enabling mirroring of persistence records

Verify that you have specified primary and secondary mirroring IP addresses on this device. Other traffic groups in the device group use these addresses when mirroring persistence records to this device.

You can perform this task to mirror persistence records to another device in a device group.

1. On the Main tab, click **Local Traffic > Profiles > Persistence**.
The Persistence profile list screen opens.
2. In the Name column, click the name of the relevant persistence profile.
3. For the **Mirror Persistence** setting, select the check box.
4. Click **Update**.

Chapter 10

Working with Folders

- *About folders on the BIG-IP system*
- *About folder attributes for redundancy*
- *About the root folder*
- *Viewing redundancy attributes for the root folder*
- *Configuring the traffic group attribute for the root folder*

About folders on the BIG-IP system

At the most basic level, a *folder* is a container for BIG-IP® configuration objects and files on a BIG-IP device. Virtual servers, pools, and self IP addresses are examples of objects that reside in folders on the system. Folders resemble standard directories, in that the system includes a root folder (represented by the / symbol) that is the parent for other folders on the system.

A folder can contain other folders.

One of the important ways that you can use folders is to set up full or granular synchronization and failover of BIG-IP configuration data in a device group. You can synchronize and fail over all configuration data on a BIG-IP device, or you can synchronize and fail over objects within a specific folder only.

About folder attributes for redundancy

Folders have two specific redundancy attributes that enable granular synchronization and failover of BIG-IP® system data within a device group. These two attributes are a device group name and a traffic group name.

Device group name

This attribute determines the scope of the synchronization, that is, the specific devices to which the system synchronizes the contents of the associated folder. When you create a Sync-Failover device group on a BIG-IP device, the system assigns that device group name as an attribute of folder `root`. Any other folders that you subsequently create on a device group member then inherit that same device group name, by default.

The result is that when you enable config sync for the local device, the contents of the `root` folder and any sub-folders are synchronized across the members of the specified device group.

If you want to synchronize a specific sub folder across only a subset of device group members, you can create a second, smaller Sync-Only device group in which the local device is also a member, and then change the sub folder's device group attribute to the new Sync-Only device group name. All objects within

that sub folder are then synchronized to the Sync-Only device group, while objects outside of that sub folder are still synchronized to the members of the larger Sync-Failover device group.

Note: *The device group assigned to a folder must contain the local BIG-IP device. Also, you cannot remove the local BIG-IP device from the Sync-Failover device group assigned to a folder.*

Traffic group name

This attribute determines the scope of a failover action, that is, the specific configuration objects that will fail over if the device becomes unavailable. If you enabled failover on a device (as part of running the Setup utility or upgrading from a previous BIG-IP version), the device contains the default traffic group named `traffic-group-1`. The system assigns this traffic group name by default as an attribute of folder `root`. Any other folders that you subsequently create on a device group member inherit that same traffic group name, by default. The result is that when the local device is a member of a Sync-Failover device group, all failover objects within the `root` folder and its hierarchy fail over based on the definition of the specified traffic group.

You can assign a different traffic group to a specific sub folder. For example, you can create an iApps™ application in a sub folder and change the inherited traffic group value of `traffic-group-1` to a traffic group that you create, such as `traffic-group-2`. You can then manually cause `traffic-group-2` to fail over to another device so that the iApp application runs on a separate device from `traffic-group-1`.

About the root folder

At the highest-level, the BIG-IP® system includes a `root` folder. The `root` folder contains all BIG-IP configuration objects on the system, by way of a hierarchical folder and sub-folder structure within it.

By default, the BIG-IP system assigns a Sync-Failover device group and a traffic group to the `root` folder. All folders and sub-folders under the `root` folder inherit these default assignments.

Viewing redundancy attributes for the root folder

You can view the device group and traffic group attributes assigned to the `root` folder. All eligible configuration objects in the `root` folder hierarchy synchronize to the named device group, and all failover objects in the hierarchy fail over with the named traffic group.

Note: *All folders and sub-folders in the root folder hierarchy inherit these attribute values, by default.*

1. On the Main tab, click **System > Platform**.
The Platform screen opens.
2. For the **Redundant Device Configuration** setting, view the device group and the traffic group attributes.

Configuring the traffic group attribute for the root folder

If you have two or more traffic groups defined on the BIG-IP® system, you can configure the traffic group attribute assigned to the `root` folder. By default, this value is `traffic-group-1`.

Note: All folders and sub folders in the `root` folder hierarchy inherit this attribute value, by default.

1. On the Main tab, click **System > Platform**.
The Platform screen opens.
2. If the system includes two or more traffic groups, then for the **Default traffic group** setting, select a traffic group from the list.
3. Click **Update**.

By default, all failover objects in the `root` folder hierarchy fail over with the named traffic group, when failover occurs.

Appendix

A

Summary of tmsh Troubleshooting Tools

- *Summary of tmsh troubleshooting tools*

Summary of tmsh troubleshooting tools

The `tmsh` utility includes a set of debugging commands for troubleshooting Sync-Only and Sync-Failover device group operations. For detailed reference material on `tmsh` commands, see the F5 Networks Technical Support web site <http://support.f5.com>.

Table 29: Summary of troubleshooting tools for device groups

Debugging Tool	Description
<code>sniff-updates</code>	Displays the commit ID updates that occur over the configuration management communications channel.
<code>watch-devicegroup-device</code>	Displays information about the devices in the device group to which the local device belongs.
<code>watch-sys-device</code>	Displays information about the local device.
<code>watch-trafficgroup-device</code>	Displays information about the traffic groups associated with devices in a device group.

Index

A

- active-active configuration
 - defined 89
- active bonus values 98
- active-standby configuration
 - defined 89
- active state
 - defined 87
- appliance models
 - failover addresses for 81
- application load
 - and failover 92
 - balancing 90
- ARP communications 101
- authentication
 - and device identity 56
 - and local trust domains 55
- authority
 - changing 55
- auto-failback feature
 - defined 96
- automatic synchronization
 - defined 69
 - enabling and disabling 65, 70
- availability
 - during failover 83
- AWS floating IP address 25, 37

C

- certificate authority
 - importing 58
 - managing and retaining 58
- certificates
 - for device trust 58
- certificate signing authorities
 - described 55
 - resetting trust on 58
- configsnc
 - configuring for VIPRION systems 23, 36
- config sync address
 - described 49
- config sync addresses
 - specifying 27, 39, 72
- config sync properties
 - advanced 78
- config sync status
 - determining 73
 - displaying 74, 76
 - troubleshooting 74
 - viewing 66
- config sync types
 - defined 71
- configuration objects
 - and traffic group associations 83
- configuration synchronization
 - about 69

- configuration synchronization (*continued*)
 - automating 70
 - preventing 84
 - scope of 107
 - syncing to group 30, 32, 43, 46, 71
- connection mirroring
 - about 49
 - and SNATs 105
 - configuring 27, 40, 104
 - considerations for 103
 - for TCP and UDP connections 105
- connections
 - preserving on failover 27, 40, 104
- current devices
 - defined 89

D

- default traffic groups
 - described 84
- device availability
 - 96
 - defined 83
- device capacity
 - and next-active device 91
- device discovery
 - defined 56
 - for device trust 28, 41, 57–58
- device group assignments
 - to root and /Common folders 62
- device group attribute
 - described 107
 - viewing on root folder 108
- device group members
 - adding and viewing 66
- device group membership 62
- device groups
 - and root folder 108
 - configuration restrictions for 62
 - configuring for VIPRION systems 23, 36
 - creating 29, 42, 63, 65
 - defined 21
 - sync status for 74
 - types of 64
 - viewing 66
- device group subset 107
- device identity
 - defined 56
- device objects
 - defined 21
- device properties 50–51
- devices
 - and mirroring limit 27, 40, 104
 - defined 21
 - discovering 56
 - excluding from config sync 69
 - running traffic groups on 89
 - selecting for failover 83, 89

- device service clustering
 - about 21
- device status types
 - described 51
 - viewing 52
- device trust
 - about 55
 - adding domain members 58
 - configuring for VIPRION systems 23, 36
 - defined 21
 - establishing 28, 41, 57
- device utilization
 - about 90, 93
 - examples of 93–95
- dropped connections 103
- DSC deployment worksheet 25, 37

F

- failback
 - defined 96
- failover
 - about 79
 - and default traffic groups 84
 - and dropped packets 101
 - and failback 96
 - and HA groups 101
 - and HA scores 98, 100
 - and next-active device 90
 - and ordered lists 96–97
 - and traffic groups 83
 - configuring for VIPRION systems 23, 36
 - scope of 107
- failover addresses
 - for appliances 81
 - for VIPRION 82
- failover devices
 - about 89
 - selecting 89
 - targeting 28, 41, 91
- failover IP addresses
 - about 49
 - specifying 31, 43, 80
 - types of 80
- failover objects
 - adding 85
 - associating with traffic groups 84
 - viewing 86
- failover states
 - viewing 88
- failover status
 - of traffic groups 86
- fast failover 97
- floating IP address
 - for AWS 25, 37
- floating IP addresses
 - and traffic groups 84
- floating traffic groups
 - and traffic group states 87
 - defined 79
- folder attributes
 - described 107

- folder hierarchy 108
- folder inheritance 61
- folders
 - and traffic group associations 83
 - associating device groups with 67
 - defined 21, 107
- Force to Standby option 84
- FTP connections
 - and mirroring 103
- full sync
 - defined 71

G

- granular synchronization
 - about 69
 - with folders 107
- gratuitous ARPs 101

H

- HA Capacity setting
 - about 91
- HA groups
 - about 79, 97
 - configuring 100–101
 - purpose of 100
- HA load factor
 - about 92
 - viewing 86
- HA load factors
 - examples of 93–95
- hardware platforms
 - and failover 28, 41, 91
 - heterogeneous 94–95
 - homogeneous 93–94
- HA scores
 - calculating 98, 100
 - purpose of 97
- HA traffic load
 - about 92
- health scores 97
- HTTP connections
 - and mirroring 103

I

- iApps applications
 - and traffic group associations 83
 - and traffic groups 84
- incremental sync
 - defined 71
- information exchange 56
- IP addresses
 - as traffic group members 85
 - for failover 80
 - for redundancy 49

L

- load-aware failover
 - about 28, 41, 79, 90–91
- local trust domain
 - and device group members 66
 - and device groups 29, 42, 63, 65
 - defined 28, 41, 55, 57–58
 - joining 56

M

- MAC masquerade addresses
 - defined 101
- management IP addresses
 - for next-active device 96
- manual synchronization 69
- mirroring
 - configuring for VIPRION systems 23, 36
 - considerations for 103
 - of connections 103
- mirroring IP addresses 49
- mirroring tasks 104

N

- network failover
 - configuring 29, 42, 63
- next-active devices
 - about 89
 - and failback 96
 - and failover 96–97, 101
 - choosing 79
 - controlling 28, 41, 91
 - defined 89

O

- object referencing 67
- ordered failover lists
 - about 96
 - configuring 97
- ordered lists
 - about 79

P

- peer authorities
 - described 55
- persistence mirroring
 - about 103
- persistence records
 - mirroring 106
- policy-sharing 64
- profiles
 - and persistence mirroring 106

R

- redundancy attributes
 - configuring 108

- redundant system configuration
 - described 21
- relative load value
 - viewing 86
- relative traffic load values 92
- root folder attributes
 - configuring 108
 - viewing 108
- root folder contents 108

S

- scores
 - for device utilization 96
- self IP addresses
 - and traffic group associations 83
 - assigning to traffic group 45
- self-signed certificates
 - regenerating 58
- service interruptions 103
- SNATs
 - and mirroring 105
 - and traffic group associations 83
- standby state
 - defined 87
 - forcing to 46, 88
- static self IP addresses
 - and traffic groups 84
- status
 - for config sync 73
- status types
 - for devices 51
 - viewing 52
- subordinate non-authorities
 - described 55
 - resetting trust on 58
- Sync-Failover configuration
 - example of 62
- Sync-Failover device groups
 - about 61
 - creating 29, 42, 63
 - illustrated 38
- synchronization 69
- Sync-Only device groups
 - about 64
 - and automatic synchronization 70
 - creating 65
 - example of 64
- sync status
 - determining 73
 - displaying 76
 - of configuration 76
 - troubleshooting 74
- sync types
 - defined 69, 71

T

- target failover devices
 - about 79, 89
- TCP connections
 - mirroring of 105

- Telnet connections
 - and mirroring *103*
 - threshold values *98*
 - traffic group attribute
 - described *107*
 - viewing on root folder *108*
 - traffic group load
 - viewing in tmsh *92*
 - traffic group members
 - adding *85*
 - traffic group properties *86*
 - traffic groups
 - activating *84*
 - and defaults *84*
 - and failover *84*
 - and failover objects *84, 86*
 - and root folder *108*
 - assigning MAC masquerade addresses to *101*
 - associating objects with *83*
 - balancing load of *89*
 - configuration restrictions for *84*
 - creating *23, 26, 35, 39, 44*
 - defined *21, 83*
 - forcing to standby state *46, 84, 88*
 - for remote devices *46, 88*
 - inheriting *83*
 - maximum number supported *83*
 - specifying load for *92*
 - viewing list of *86*
 - traffic group states
 - defined *87*
 - viewing *88*
 - traffic load
 - balancing *90*
 - traffic load values *92*
 - troubleshooting tools *111*
 - trust authority
 - managing and resetting *58*
 - trust domains
 - and local trust domain *28, 41, 55, 57–58*
 - trust relationships
 - between devices *55*
- ## U
- UDP connections
 - and mirroring *103*
 - mirroring of *105*
- ## V
- VIPRION
 - failover addresses for *82*
 - VIPRION systems
 - mirroring connections on *49*
 - virtual addresses
 - assigning to traffic group *45*
 - virtual IP addresses
 - and traffic group associations *83*
 - virtual servers
 - and mirroring *105*
 - VLANs
 - and traffic groups *84*
- ## W
- weight values *98*
- ## X
- x509 certificates
 - and device identity *56*
 - and device trust *55*
 - for device trust *28, 41, 57*