

53-1001760-01
30 March 2010



Access Gateway

Administrator's Guide

Supporting Fabric OS v6.4.0

BROCADE

Copyright © 2007-2010 Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, IronPoint, IronShield, IronView, IronWare, JetCore, NetIron, SecureIron, ServerIron, StorageX, and Turbolron are registered trademarks, and DCFM, Extraordinary Networks, and SAN Health are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. shall have no liability or responsibility to any person or entity with respect to any loss, cost, liability, or damages arising from the information contained in this book or the computer programs that accompany it.

The product described by this document may contain “open source” software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Brocade Communications Systems, Incorporated

Corporate and Latin American Headquarters
Brocade Communications Systems, Inc.
1745 Technology Drive
San Jose, CA 95110
Tel: 1-408-333-8000
Fax: 1-408-333-8101
E-mail: info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems China HK, Ltd.
No. 1 Guanghua Road
Chao Yang District
Units 2718 and 2818
Beijing 100020, China
Tel: +8610 6588 8888
Fax: +8610 6588 9999
E-mail: china-info@brocade.com

European Headquarters
Brocade Communications Switzerland Sàrl
Centre Swissair
Tour B - 4ème étage
29, Route de l'Aéroport
Case Postale 105
CH-1215 Genève 15
Switzerland
Tel: +41 22 799 5640
Fax: +41 22 799 5641
E-mail: emea-info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems Co., Ltd. (Shenzhen WFOE)
Citic Plaza
No. 233 Tian He Road North
Unit 1308 – 13th Floor
Guangzhou, China
Tel: +8620 3891 2000
Fax: +8620 3891 2111
E-mail: china-info@brocade.com

Document History

The following table lists all versions of the *Access Gateway Administrator's Guide*.

Document Title	Publication Number	Summary of Changes	Publication Date
<i>Access Gateway Administrator's Guide</i>	53-1000430-01	First version	January 2007
<i>Access Gateway Administrator's Guide</i>	53-1000633-01	Added support for the 200E	June 2007
<i>Access Gateway Administrator's Guide</i>	53-1000605-01	Added support for new policies and changes to N_Port mappings.	October 2007
<i>Access Gateway Administrator's Guide</i>	53-1000605-02	Added support for new platforms: 300 and the 4424. Added support for new features: - Masterless Trunking - Direct Target Connectivity - Advance Device Security policy - 16- bit routing	March 2008
<i>Access Gateway Administrator's Guide</i>	53-1000605-03	Added support for: - Cascading Access Gateway.	July 2008
<i>Access Gateway Administrator's Guide</i>	53-1000605-04	Updated to fix the table of contents	July 2008
<i>Access Gateway Administrator's Guide</i>	53-1001189-01	Updated for Fabric OS 6.2.0	November 2008
<i>Access Gateway Administrator's Guide</i>	53-1001345-01	Updated for Fabric OS 6.3.0	July 2009
<i>Access Gateway Administrator's Guide</i>	53-1001760-01	Updated for Fabric OS 6.4.0	March 2010

Contents

About This Document

How this document is organized	xiii
Supported hardware and software	xiii
What's new in this document	xiv
Document conventions	xv
Text formatting	xv
Command syntax conventions	xv
Notes, cautions, and warnings	xv
Notice to the reader	xvi
Key terms	xvi
Additional information	xvii
Brocade resources	xvii
Other industry resources	xvii
Optional Brocade features	xviii
Getting technical help	xviii
Document feedback	xix

Chapter 1

Access Gateway Basic Concepts

In this chapter	1
Brocade Access Gateway overview	1
Comparing Native Fabric and Access Gateway modes	1
Fabric OS features in Access Gateway mode	3
Access Gateway port types	4
Comparison of Access Gateway ports to standard switch ports	4
Access Gateway hardware considerations	5

Chapter 2

Configuring Ports in Access Gateway mode

In this chapter	7
Enabling and disabling Access Gateway mode	7
Port state description	9
Access Gateway mapping	10
Port-based mapping	10
Device-based mapping	15
Considerations for Access Gateway mapping	22

N_Port configurations	24
Displaying N_Port configurations	25
Unlocking N_Ports	25

Chapter 3

Managing Policies and Features in Access Gateway Mode

In this chapter	27
Access Gateway policies overview	27
Displaying current policies	27
Access Gateway policy enforcement matrix	28
Advanced Device Security policy	28
How the ADS policy works	28
Enabling and disabling the Advanced Device Security policy	29
Setting the list of devices allowed to log in	29
Setting the list of devices not allowed to log in	30
Removing devices from the list of allowed devices	30
Adding new devices to the list of allowed devices	30
Displaying the list of allowed devices on the switch	31
ADS policy considerations	31
Upgrade and downgrade considerations for the ADS policy	31
Automatic Port Configuration policy	31
How the APC policy works	31
Enabling and disabling the APC policy	32
Automatic Port Configuration policy considerations	32
Upgrade and downgrade considerations for the APC policy	33
Port Grouping policy	33
How port groups work	33
Adding an N_Port to a port group	34
Deleting an N_Port from a port group	35
Removing a port group	35
Renaming a port group	35
Disabling the Port Grouping policy	35
Port Grouping policy modes	36
Creating a port group and enabling Automatic Login	
Balancing mode	36
Rebalancing F_Ports	37
Enabling Managed Fabric Name Monitoring mode	38
Disabling Managed Fabric Name Monitoring mode	38
Displaying the current fabric name monitoring timeout value	38
Setting the current fabric name monitoring timeout value	39
Port Grouping policy considerations	39
Upgrade and downgrade considerations for the	
Port Grouping policy	40
Device Load Balancing Policy	40
Enabling WWN Load Balancing	40
Disabling Device Load Balancing	40
Device Load Balancing considerations	41

Persistent ALPA Policy	41
Enabling Persistent ALPA	42
Disabling Persistent ALPA	42
Persistent ALPA device data	42
Removing device data from the database	42
Displaying device data	43
Clearing ALPA values	43
Persistent ALPA policy considerations	43
Upgrade and downgrade considerations for Persistent ALPA	43
Failover	44
Failover with port-based mapping	44
Failover with device-based mapping	46
Enabling and disabling Failover on a N_Port	47
Enabling and disabling Failover for a port group	47
Upgrade and downgrade considerations for Failover	48
Failback	48
Failback configurations in Access Gateway	48
Enabling and disabling Failback on an N_Port	49
Enabling and disabling Failback for a port group	50
Upgrade and downgrade considerations for Failback	50
Trunking in Access Gateway mode	50
How Trunking works	50
Configuring Trunking on the Edge switch	51
Configuration management for trunk areas	52
Enabling trunking	53
Disabling F_Port trunking	54
Trunking monitoring	54
Trunking considerations for the Edge switch	54
Trunking considerations for Access Gateway module	57
Upgrade and downgrade considerations for Trunking in Access Gateway mode	57
Adaptive Networking on Access Gateway	58
Upgrade and downgrade considerations with Adaptive Networking in AG mode enabled	59
Adaptive Networking on Access Gateway considerations	59
Per Port NPIV login limit	60
Setting the login limit	60
Considerations for the Brocade 8000	60

Chapter 4

SAN Configuration with Access Gateway

In this chapter	63
Connectivity of multiple devices overview	63
Direct target attachment	63
Considerations	63
Target aggregation	64
Access Gateway cascading	64

Fabric and Edge switch configuration	65
Verifying the switch mode	65
Enabling NPIV on M-EOS switches	66
Connectivity to Cisco Fabrics	67
Enabling NPIV on a Cisco switch.	67
Rejoining Fabric OS switches to a fabric	67
Reverting to a previous configuration.	67

Appendix A Troubleshooting

Index

Figures

Figure 1	Switch function in Native mode	2
Figure 2	Switch function in Access Gateway mode	2
Figure 3	Port usage comparison	5
Figure 4	Example port-based mapping.	11
Figure 5	Example of device mapping to N_Port groups.	17
Figure 6	Example device mapping to an N_Port	18
Figure 7	Example of adding an external F_Port (F9) on an embedded switch	24
Figure 8	Port grouping behavior	34
Figure 9	Port group 1 (pg1) setup	34
Figure 10	Example 1 and 2 Failover behavior	45
Figure 11	Failback behavior.	49
Figure 12	Starting point for QoS	59
Figure 13	Access Gateway cascading.	64

Tables

Table 1	Fabric OS components supported on Access Gateway	3
Table 2	Port configurations	5
Table 3	Port state description	9
Table 4	Description of port mapping	11
Table 5	Access Gateway default port mapping	12
Table 6	Policy enforcement matrix	28
Table 7	Address identifier	52
Table 8	Access Gateway trunking considerations for the Edge switch	54
Table 9	PWWN format for F_Port and N_Port trunk ports.	57
Table 10	Troubleshooting	69

About This Document

• How this document is organized	xiii
• Supported hardware and software	xiii
• What's new in this document	xiv
• Document conventions	xv
• Notice to the reader	xvi
• Key terms	xvi
• Additional information	xvii
• Getting technical help	xviii
• Document feedback	xix

How this document is organized

This document is a procedural guide to help SAN administrators configure and manage Brocade Access Gateway.

This preface contains the following components:

- [Chapter 1, “Access Gateway Basic Concepts”](#) describes the Brocade Access Gateway and provides an overview of its key features.
- [Chapter 2, “Configuring Ports in Access Gateway mode”](#) describes how to configure ports in Access Gateway mode.
- [Chapter 3, “Managing Policies and Features in Access Gateway Mode”](#) describes how to enable policies on a switch in Access Gateway mode. It also provides information on how to set up Failover and Failback, and discusses how Trunking and Adaptive Networking works in AG.
- [Chapter 4, “SAN Configuration with Access Gateway”](#) describes how to connect multiple devices using Access Gateway.
- [Appendix A, “Troubleshooting”](#) provides symptoms and troubleshooting tips to resolve issues.

Supported hardware and software

In those instances in which procedures or parts of procedures documented here apply to some switches but not to others, this guide identifies exactly which switches are supported and which are not.

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc. For Fabric OS v6.4.0, documenting all possible configurations and scenarios is beyond the scope of this document.

All Fabric OS switches must be running v6.1.0 or later; all M-EOS switches must be running M-EOSc 9.1 or later, M-EOSn must be running 9.6.2 or later, and Cisco switches with SAN OS must be running 3.0 (1) and 3.1 (1) or later.

Fabric OS v6.4.0 supports the following Brocade hardware platforms for Access Gateway:

- Brocade 300
- Brocade 5100
- Brocade M5424
- Brocade 5450
- Brocade 5460
- Brocade 5470
- Brocade 5480
- Brocade VA40-FC
- Brocade 8000

What's new in this document

The following changes have been made since this document was last released:

Information on the following subjects was added:

- Device mapping
- Mapping priority support.
- Support for the device login balancing policy.
- AG support for the Brocade 8000.
- Setting per port NPIV login limits
- Recommendations for connecting host and target ports to N_Ports.

For further information, refer to the release notes.

Document conventions

This section describes text formatting conventions and important notices formats.

Text formatting

The narrative-text formatting conventions that are used in this document are as follows:

bold text	Identifies command names Identifies the names of user-manipulated GUI elements Identifies keywords and operands Identifies text to enter at the GUI or CLI
<i>italic text</i>	Provides emphasis Identifies variables Identifies paths and Internet addresses Identifies document titles
<code>code text</code>	Identifies CLI output Identifies syntax examples

For readability, command names in the narrative portions of this guide are presented in mixed lettercase: for example, **switchShow**. In actual examples, command lettercase is often all lowercase.

Command syntax conventions

Command syntax in this manual follows these conventions:

command	Commands are printed in bold.
-- option, option	Command options are printed in bold.
- argument , arg	Arguments.
[]	Optional element.
<i>variable</i>	Variables are printed in italics. In the help pages, values are <u>underlined</u> or enclosed in angled brackets < >.
...	Repeat the previous element, for example "member[;member...]"
value	Fixed values following arguments are printed in plain font. For example, -- show WWN
	Boolean. Elements are exclusive. Example: -- show -mode egress ingress

Notes, cautions, and warnings

The following notices appear in this document.

NOTE

A note provides a tip, emphasizes important information, or provides a reference to related information.

ATTENTION

An Attention statement indicates potential damage to hardware or data.

**CAUTION**

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.

**DANGER**

A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Notice to the reader

This document may contain references to the trademarks of the following corporations. These trademarks are the properties of their respective companies and corporations.

These references are made for informational purposes only.

Corporation	Referenced Trademarks and Products
Cisco Systems, Inc.	Cisco
Oracle Corporation.	Sun, Solaris
Netscape Communications Corporation	Netscape
Red Hat, Inc.	Red Hat, Red Hat Network, Maximum RPM, Linux Undercover
Emulex Corporation	Emulex
QLogic Corporation	QLogic

Key terms

For definitions of SAN-specific terms, visit the Storage Networking Industry Association online dictionary at: <http://www.snia.org/education/dictionary>.

For definitions specific to Brocade and Fibre Channel, see the *Brocade Glossary*.

The following terms are used in this manual to describe Access Gateway mode and its components.

Access Gateway (AG)

Fabric OS mode for switches that reduces SAN (storage area network) deployment complexity by leveraging NPIV (N_Port ID Virtualization).

Device

Any host or target device with a distinct WWN. Devices may be physical or virtual.

E_Port	An ISL (Interswitch link) port. A switch port that connects switches together to form a fabric.
Edge switch	A fabric switch that connects host, storage, or other devices, such as Brocade Access Gateway, to the fabric.
F_Port	A fabric port. A switch port that connects a host, HBA (host bus adaptor), or storage device to the SAN. On Brocade Access Gateway, the F_Port connects to a host or a target.
Mapping	In Access Gateway mapping defines the routes between devices or F_Ports to the fabric facing ports (N_Ports).
N_Port	A node port. A Fibre Channel host or storage port in a fabric or point-to-point connection. On Brocade Access Gateway, the N_Port connects to the Edge switch.
NPIV	N_Port ID Virtualization. This is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port. This allows multiple Fibre Channel initiators to occupy a single physical port, easing hardware requirements in Storage Area Network design, especially for virtual SANs.

Additional information

This section lists additional Brocade and industry-specific documentation that you might find helpful.

Brocade resources

To get up-to-the-minute information, go to <http://my.brocade.com> and register at no cost for a user ID and password.

For additional Brocade documentation, visit the Brocade SAN Info Center and click the Resource Library location:

<http://www.brocade.com>

Release notes are available on the My Brocade website (<http://my.brocade.com>) and are also bundled with the Fabric OS firmware.

Other industry resources

- White papers, online demonstrations, and data sheets are available through the Brocade website at <http://www.brocade.com/products/software.jhtml>.
- Best practice guides, white papers, data sheets, and other documentation is available through the Brocade Partner website.

For additional resource information, visit the Technical Committee T11 website. This website provides interface standards for high-performance and mass storage applications for Fibre Channel, storage management, and other applications:

<http://www.t11.org>

For information about the Fibre Channel industry, visit the Fibre Channel Industry Association website:

<http://www.fibrechannel.org>

Optional Brocade features

For a list of optional Brocade features and descriptions, see the *Fabric OS Administrator's Guide*.

Getting technical help

Contact your switch support supplier for hardware, firmware, and software support, including product repairs and part ordering. To expedite your call, have the following information available:

1. General Information

- Technical Support contract number, if applicable
- Switch model
- Switch operating system version
- Error numbers and messages received
- **supportSave** command output
- Detailed description of the problem, including the switch or fabric behavior immediately following the problem, and specific questions
- Description of any troubleshooting steps already performed and the results
- Serial console and Telnet session logs
- Syslog message logs

2. Switch Serial Number

The switch serial number and corresponding bar code are provided on the serial number label, as shown here.



The serial number label is located as follows:

- Brocade 300, 4100, 4900, 5100, 5300, 7500, 7500E, 7800, 8000, VA-40FC, and Brocade Encryption Switch—On the switch ID pull-out tab located inside the chassis on the port side on the left
- Brocade 5000—On the switch ID pull-out tab located on the bottom of the port side of the switch
- Brocade 7600—On the bottom of the chassis
- Brocade 48000—Inside the chassis next to the power supply bays
- Brocade DCX—On the bottom right on the port side of the chassis
- Brocade DCX-4S—On the bottom right on the port side of the chassis, directly above the cable management comb

3. World Wide Name (WWN)

Use the **licenseIdShow** command to display the WWN of the chassis.

If you cannot use the **licenseIdShow** command because the switch is inoperable, you can get the WWN from the same place as the serial number, except for the Brocade DCX. For the Brocade DCX, access the numbers on the WWN cards by removing the Brocade logo plate at the top of the nonport side of the chassis.

Document feedback

Quality is our first concern at Brocade and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. Forward your feedback to:

documentation@brocade.com

Provide the title and version number of the document and as much detail as possible about your comment, including the topic heading and page number and your suggestions for improvement.

Access Gateway Basic Concepts

In this chapter

- [Brocade Access Gateway overview.](#) 1
- [Fabric OS features in Access Gateway mode.](#) 3
- [Access Gateway port types](#) 4
- [Access Gateway hardware considerations.](#) 5

Brocade Access Gateway overview

Brocade Access Gateway (AG) is a Fabric OS feature that lets you configure your Enterprise fabric to handle additional devices instead of domains. You do this by configuring F_Ports to connect to the fabric as N_Ports, which increases the number of device ports you can connect to a single fabric. Multiple AGs can connect to the DCX enterprise-class platform, directors, and switches.

Access Gateway is compatible with Fabric OS, M-EOS v9.1 or v9.6 and later, and Cisco-based fabrics v3.0 (1) or later and v3.1 (1) and later. Enabling and disabling AG mode and configuring AG features on a switch can be performed from the command line interface (CLI), Web Tools, or Fabric Manager. This document describes configurations using the CLI commands. Please see the **Web Tools Administrator's Guide**, the **Fabric Manager Administrator's Guide**, or the **Data Center Fabric Manager User Guide** for more information about AG support in those tools.

After you set a Fabric OS switch to AG mode, the F_Ports connect to the Enterprise fabric as N_Ports rather than as E_Ports. [Figure 1](#) shows a comparison of a configuration that connects eight hosts to a fabric using AG to the same configuration with Fabric OS switches in Native mode.

Switches in AG mode are logically transparent to the host and the fabric. Therefore, you can increase the number of hosts that have access to the fabric without increasing the number of switch domains. This simplifies configuration and management in a large fabric by reducing the number of domain IDs and ports.

Comparing Native Fabric and Access Gateway modes

The following points summarize the differences between a Fabric OS switch functioning in Native operating mode and a Fabric OS switch functioning in AG operating mode:

- The Fabric OS switch in Native mode is a part of the fabric; it requires two to four times as many physical ports, consumes fabric resources, and can connect to a Fabric OS fabric only.
- A switch in AG mode is outside of the fabric; it reduces the number of switches in the fabric and the number of required physical ports. You can connect an AG switch to either a Fabric OS, M-EOS, or Cisco-based fabric.

For comparison, [Figure 1](#) illustrates switch function in Native mode and [Figure 2](#) illustrates switch function in AG mode.

1 Brocade Access Gateway overview




FIGURE 1 Switch function in Native mode




FIGURE 2 Switch function in Access Gateway mode

Fabric OS features in Access Gateway mode

Table 1 lists Fabric OS components that are supported on a switch when AG mode is enabled. “No” indicates that the feature is not provided in AG mode. “NA” indicates this feature is not applicable in Access Gateway mode of operation. A single asterisk (*) indicates the feature is transparent to AG, that is AG forwards the request to the Enterprise fabric. Two asterisks (**) indicates that if the Enterprise fabric is not a Brocade fabric, the feature may not be available.

TABLE 1 Fabric OS components supported on Access Gateway

Feature	Support
Access Control	Yes (limited roles) ¹
Adaptive Networking	Yes
Admin Domains	No
Audit	Yes
Beaconing	Yes
Config Download/Upload	Yes
DHCP	Yes
Environmental Monitor	Yes
Error Event Management	Yes
Extended Fabrics	No
Fabric Device Management Interface (FDMI)	Yes*
Fabric Manager	Yes**
Fabric Watch	Yes (limited)
FICON (includes CUP)	No
High Availability	Hot Code Load
Native Interoperability Mode	NA
License	Yes**
Log Tracking	Yes
Management Server	NA
Manufacturing Diagnostics	Yes
N_Port ID Virtualization	Yes
Name Server	NA
Network Time Protocol (NTP)	No (no relevance from fabric perspective) ²
Open E_Port	NA
Performance Monitor	Yes (Basic PM only, no APM support)
Persistent ALPA	Yes
Port Mirroring	No
QuickLoop, QuickLoop Fabric Assist	No
Security	Yes (ADS/DCC Policy)
SNMP	Yes

TABLE 1 Fabric OS components supported on Access Gateway (Continued)

Feature	Support
Speed Negotiation	Yes
Syslog Daemon	Yes
Trunking	Yes**
ValueLineOptions (Static POD, DPOD)	Yes
Web Tools	Yes
Zoning	NA

1. When a switch is behaving as an AG, RBAC features in Fabric OS are available, but there are some limitations. For more information on the limitations, refer to ["Access Gateway hardware considerations"](#) on page 5.
2. In embedded switches, time should be updated by the server management utility.

Access Gateway port types

Access Gateway differs from a typical fabric switch because it is not a switch; instead, it is a mode that you enable on a switch using the **ag** command. After a switch is set in **ag** mode, it can connect to the fabric using node ports (N_Ports). Typically fabric switches connect to the Enterprise fabric using ISL (InterSwitch Link) ports, such as E_Ports.

Following are the Fibre Channel (FC) ports that AG uses:

- **F_Port** - fabric port that connects a host, HBA, or storage device to a switch in AG mode.
- **N_Port** - node port that connects a switch in AG mode to the F_Port of the fabric switch.

Comparison of Access Gateway ports to standard switch ports

Access Gateway multiplexes host connections to the fabric. It presents an F_Port to the host and an N_Port to an Edge fabric switch. Using N_Port ID Virtualization (NPIV), AG allows multiple FC initiators to access the SAN on the same physical port. This reduces the hardware requirements and management overhead of hosts to the SAN connections.

A fabric switch presents F_Ports (or FL_Ports) and storage devices to the host and presents E_Ports, VE_Ports, or EX_Ports to other switches in the fabric. A fabric switch consumes SAN resources, such as domain IDs, and participates in fabric management and zoning distribution. A fabric switch requires more physical ports than AG to connect the same number of hosts.

[Figure 3](#) on page 5 shows a comparison of the types of ports a switch in AG mode uses to the type of ports that a switch uses in standard mode.

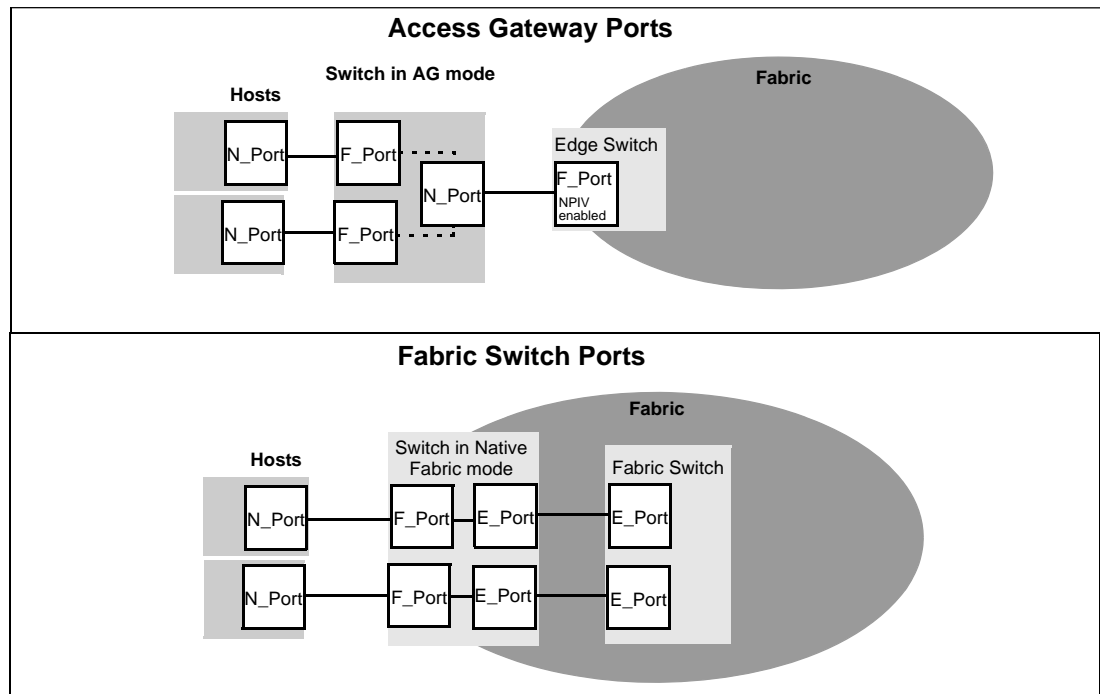


FIGURE 3 Port usage comparison

Table 2 shows a comparison of port configurations with AG to a standard fabric switch.

TABLE 2 Port configurations

Port Type	Access Gateway		Fabric switch	
F_Port	Yes	Connects hosts and targets to Access Gateway.	Yes	Connects devices, such as hosts, HBAs, and storage to the fabric.
N_Port	Yes	Connects Access Gateway to a fabric switch.	NA	N_Ports are not supported.
E_Port	NA	ISL is not supported. ¹	Yes	Connects the switch to other switches to form a fabric.

1. The switch is logically transparent to the fabric, therefore it does not participate in the SAN as a fabric switch.

Access Gateway hardware considerations

Hardware considerations for Access Gateway are as follows:

- Access Gateway is supported on the switch platforms and embedded switch platforms listed in [“Supported hardware and software”](#) on page xiii.
- Loop devices are not supported.
- Direct connections to SAN target devices are only supported if the AG-enabled module is connected to a fabric.

1 Access Gateway hardware considerations

Configuring Ports in Access Gateway mode

In this chapter

- Enabling and disabling Access Gateway mode 7
- Access Gateway mapping 10
- N_Port configurations 24

Enabling and disabling Access Gateway mode

Use the following steps to enable and disable Access Gateway mode. After you enable AG mode, some fabric information is erased, such as the zone and security databases. Enabling AG mode is disruptive because the switch is disabled and rebooted. For more information on the **ag** commands used in these steps, refer to the *Fabric OS Command Reference*.

1. Before enabling or disabling a switch to AG mode, save the current configuration file using the **configupload** command in case you might need this configuration again.
2. Ensure that no zoning or Admin Domain (AD) transaction buffers are active. If any transaction buffer is active, enabling AG mode will fail with the error, “Failed to clear Zoning/Admin Domain configuration”.
3. Verify that the switch is set to Native mode or **interopmode 0**.
 - a. Issue the **switchshow** command to verify the switch mode.
 - b. If the switch mode is anything other than 0, issue the **interopmode 0** command to set the switch to Native mode.

For more information on setting switches to Native mode, refer to the *Fabric OS Administrator's Guide*.

4. Enter the **switchdisable** command.

```
switch:admin> switchdisable
```

This command disables all user ports on a switch. All Fibre Channel ports are taken offline. If the switch was part of a fabric, the remaining switches reconfigure. You must disable the switch before making configuration changes.

5. Enter the **ag --modeenable** command.

```
switch:admin> ag --modeenable
```

The switch automatically reboots and comes back online in AG mode using a factory default port mapping. For more information on AG default port mapping, see [Table 5](#) on page 12.

6. Enter the **ag --modeshow** command to verify that AG mode is enabled.

```
switch:admin> ag --modeshow
Access Gateway mode is enabled.
```

2 Enabling and disabling Access Gateway mode

You can display the port mappings and status of the host connections to the fabric on Access Gateway.

7. Enter the **ag --mapshow** command to display all the mapped ports.

The **ag --mapshow** command shows all the N_Ports (with the **portcfgnport** value of 1) even if those N_Ports are not connected.

```
switch:admin> ag --mapshow
```

N_Port	Configured_F_Ports	Current_F_Ports	Failover	Failback	PG_ID	PG_Name
0	4;5;6	4;5;6	1	0	2	SecondFabric
1	7;8;9	7;8;9	0	1	0	pg0
2	10;11	10;11	1	0	2	SecondFabric
3	12;13	12;13	0	1	0	pg0

8. Enter the **switchshow** command to display the status of all ports. Note that the following output is an example only and may not exactly reflect output from the current Fabric OS.

```
switch:admin> switchshow
```

```
switchName:      switch
switchType:      43.2
switchState:     Online
switchMode:      Access Gateway Mode
switchWwn:       10:00:00:05:1e:03:4b:e7
```

```
switchBeacon:    OFF
```

Area	Port	Media	Speed	State	Proto
0	0	--	N4	No_Module	
1	1	cu	N4	Online	F-Port 50:06:0b:00:00:3c:b7:32 0x5a0101
2	2	cu	N4	Online	F-Port 10:00:00:00:c9:35:43:f5 0x5a0003
3	3	cu	N4	Online	F-Port 50:06:0b:00:00:3c:b6:1e 0x5a0102
4	4	cu	N4	Online	F-Port 10:00:00:00:c9:35:43:9b 0x5a0002
5	5	cu	N4	Online	F-Port 50:06:0b:00:00:3c:b4:3e 0x5a0201
6	6	cu	N4	Online	F-Port 10:00:00:00:c9:35:43:f3 0x5a0202
7	7	cu	AN	No_Sync	Disabled (Persistent)
8	8	cu	N4	Online	F-Port 10:00:00:00:c9:35:43:a1 0x5a0001
9	9	cu	AN	No_Sync	Disabled (Persistent)
10	10	cu	AN	No_Sync	Disabled (Persistent)
11	11	cu	AN	No_Sync	Disabled (Persistent)
12	12	cu	AN	No_Sync	Disabled (Persistent)
13	13	cu	AN	No_Sync	Disabled (Persistent)
14	14	cu	AN	No_Sync	Disabled (Persistent)
15	15	cu	AN	No_Sync	Disabled (Persistent)
16	16	cu	AN	No_Sync	Disabled (Persistent)
17	17	--	N4	No_Module	
18	18	--	N4	No_Module	
19	19	id	N4	No_Light	
20	20	--	N4	No_Module	
21	21	id	N4	Online	N-Port 10:00:00:05:1e:35:10:1e 0x5a0200
22	22	id	N4	Online	N-Port 10:00:00:05:1e:35:10:1e 0x5a0100
23	23	id	N4	Online	N-Port 10:00:00:05:1e:35:10:1e 0x5a0000

For a description of the port state, see [Table 3](#) on page 9.

When you disable AG mode, The switch automatically reboots and comes back online using the fabric switch configuration; the AG parameters, such as port mapping, and Failover and Failback are automatically removed. When the switch reboots, it starts in Fabric OS Native mode. To re-join the switch to the core fabric, refer to [“Rejoining Fabric OS switches to a fabric”](#) on page 67.

9. Enter the **switchDisable** command to disable the switch.

```
switch:admin> switchdisable
```

10. Enter the **ag** command with the **--modedisable** operand to disable AG mode.

```
switch:admin> ag --modedisable
```

11. Enter the **ag --modeshow** command to verify that AG mode is disabled.

```
switch:admin> ag --modeshow
Access Gateway mode is NOT enabled
```

Port state description

The following table describes the possible port states.

TABLE 3 Port state description

State	Description
No _Card	No interface card present
No _Module	No module (GBIC or other) present
Mod_Val	Module validation in process
Mod_Inv	Invalid module
No_Light	The module is not receiving light
No_Sync	Receiving light but out of sync
In_Sync	Receiving light and in sync
Laser_Flt	Module is signaling a laser fault
Port_Flt	Port marked faulty
Diag_Flt	Port failed diagnostics
Lock_Ref	Locking to the reference signal
Testing	Running diagnostics
Offline	Connection not established (only for virtual ports)
Online	The port is up and running

Access Gateway mapping

When operating in AG mode you must specify pre-provisioned routes that AG will use to direct traffic from the devices (hosts or targets) on its F_Ports to the ports connected to the fabric using its N_Ports. This is unlike Native switch mode where the switch itself determines the best path between its F_Ports. This process of pre-provisioning routes in AG mode is called “mapping.”

During mapping, device WWNs or F_Ports are assigned to N_Ports and N_Port groups on the switch running in AG mode. Mapping ensures that a device logging into the switch will always connect to the fabric through a specific N_Port or N_Port group. Two types of mapping are available:

- Port mapping

A specific F_Port is mapped to a specific N_Port. This ensures that all traffic from a specific F_Port always goes through the same N_Port. To map an F_Port to an N_Port group, simply map the port to an N_Port that belongs to that port group. All F_Ports mapped to that N_Port will be part of that port group.

- Device-based mapping (optional)

A specific device WWN is mapped to N_port groups (preferred method) or to specific N_Ports. Device mapping allows a virtual port to access its destination regardless of which F_Port on switch it resides on. Device mapping also allows multiple virtual ports on a single physical machine access multiple destinations residing in different fabrics.

Device-based mapping is optional and should be added on top of existing port maps. Port mapping must exist at all times.

Port-based mapping

An F_Port needs to be mapped to an N_Port before the F_Port can come online. When you first enable a switch to AG mode, by default, the F_Ports are mapped to a set of predefined N_Ports. For default port mapping on supported hardware platforms, refer to [Table 5](#). Refer to [Adding F_Ports to an N_Port](#) if you want to change the default mapping.

[Figure 4](#) shows a mapping with eight F_Ports evenly mapped to four N_Ports on a switch in AG mode. The N_Ports connect to the same fabric through different Edge switches.

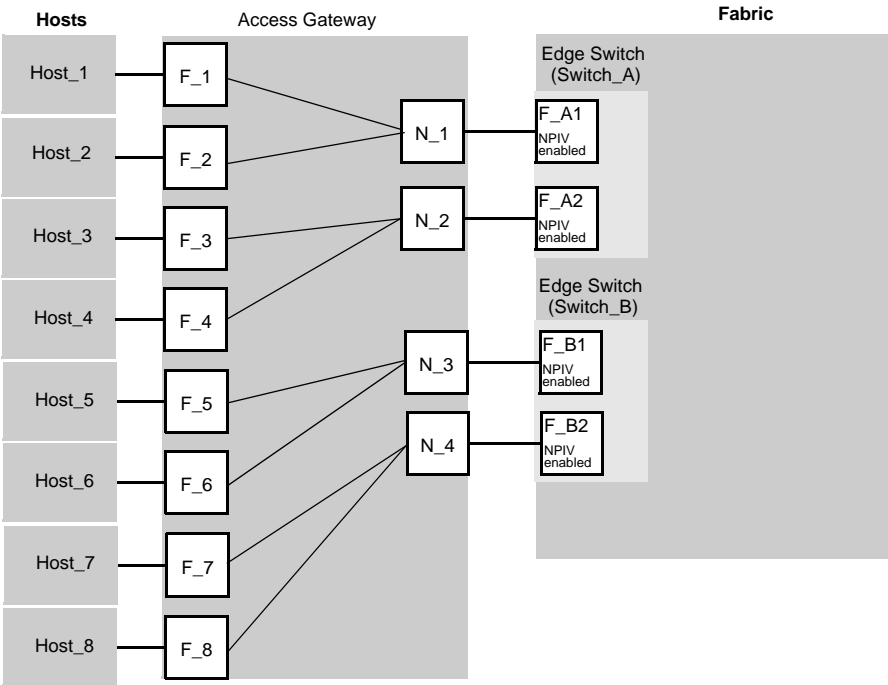


FIGURE 4 Example port-based mapping

Table 4 provides a description of the port mapping in Figure 4.

TABLE 4 Description of port mapping

Access Gateway		Fabric	
F_Port	N_Port	Edge switch	F_Port
F_1, F_2	N_1	Switch_A	F_A1
F_3, F_4	N_2	Switch_A	F_A2
F_5, F_6	N_3	Switch_B	F_B1
F_7, F_8	N_4	Switch_B	F_B2

Considerations for initiator and target ports

Following are the possible connections to FCP initiator (host) and target ports through AG:

- All F_Ports connect to all initiator ports.
- All F_Ports connected to all target ports.
- Some F_Ports connected to initiator ports and some F_Ports connected to target ports.

For the last case, communication between initiator and target ports is not supported if both are mapped to the same N_Port. Therefore, follow these recommendations for initiator and target port mapping:

- If connecting a host and target port to the same AG, you should map them to separate N_Ports and connect those N_Ports to the same fabric.
- Use separate port groups for initiator and target ports.

- When configuring secondary port mapping for failover and failback situations, make sure that initiator and target F_Ports will not fail over or fail back to the same N_Port.

Brocade 8000 mapping differences

The Brocade 8000 contains 24 internal FCoE ports and eight external Fibre Channel ports. In Access Gateway mode, the internal FCoE ports are configured logically as F_Ports, while the external Fibre Channel ports are configured as N_Ports. The FCoE ports are divided into six groups or trunks consisting of four ports each. All four ports in a group are mapped to one N_Port. Although you can change the default port mapping for these groups (refer to “[Default port mapping](#)” on page 12), consider the following when working with these FCoE ports:

- All four FCoE ports in the group are mapped to the same N_Port.
- You cannot map individual FCoE ports within the same port group to different N_Ports.
- Any Access Gateway operation that involves moving F_Ports will move all FCoE ports in the group.
- All four FCoE ports in a group will failover or failback to one N_Port.

Default port mapping

[Table 5](#) shows the default port mapping. By default, Failover and Failback policies are enabled on all N_Ports.

NOTE

All POD licenses must be present to use Access Gateway on the Brocade 5100 and 300.

TABLE 5 Access Gateway default port mapping

Brocade Model	Total Ports	F_Ports	N_Ports	Default Port Mapping
VA40-FC	40	0-31	32-39	0-3 mapped to 32 4-7 mapped to 33 8-11 mapped to 34 12-15 mapped to 35 16-19 mapped to 36 20-23 mapped to 37 24-27 mapped to 38 28-31 mapped to 39
300	24	0-15	16 -23	0, 1 mapped to 16 2, 3 mapped to 17 4, 5 mapped to 18 6, 7 mapped to 19 8, 9 mapped to 20 10, 11 mapped to 21 12, 13 mapped to 22 14, 15 mapped to 23

TABLE 5 Access Gateway default port mapping (Continued)

Brocade Model	Total Ports	F_Ports	N_Ports	Default Port Mapping
5100	40	0-31	32-39	0, 1, 2, 3 mapped to 32 4, 5, 6, 7 mapped to 33 8, 9, 10, 11 mapped to 34 12, 13, 14, 15 mapped to 35 16, 17, 18, 19 mapped to 36 20, 21, 22, 23 mapped to 37 24, 25, 26, 27 mapped to 28 28, 29, 30, 31 mapped to 39
5424	24	1-16	0, 17-23	0, 17-23 1, 2 mapped to 17 3, 4 mapped to 18 5, 6 mapped to 19 7, 8 mapped to 20 9, 10 mapped to 21 11, 12 mapped to 22 13, 14 mapped to 23 15, 16 mapped to 0
5450	26	6-25 Not all ports may be present.	0, 19-25	1, 2, 17 mapped to 19 3, 4, 18 mapped to 20 5, 6 mapped to 21 7, 8 mapped to 22 9, 10 mapped to 23 11, 12 mapped to 24 13, 14 mapped to 25 15, 16 mapped to 0
5460	26	6-25	0-5	6 and 16 mapped to 0 7 and 17 mapped to 1 8, 12, 18, and 22 mapped to 2 9, 13, 19, and 23 mapped to 3 10, 14, 20, and 24 mapped to 4 11, 15, 21, and 25 mapped to 5
5470	20	1-14	0, 15-19	1, 2 mapped to 0 3, 4 mapped to 15 5, 6, 7 mapped to 16 8, 9 mapped to 17 10, 11 mapped to 18 12, 13, 14 mapped to 19

TABLE 5 Access Gateway default port mapping (Continued)

Brocade Model	Total Ports	F_Ports	N_Ports	Default Port Mapping
5480	24	1-16	0, 17-23	1, 2 mapped to 17 9, 10 mapped to 18 3, 4 mapped to 19 11, 12 mapped to 20 15, 16 mapped to 0 5, 6 mapped to 21 13, 14 mapped to 22 7, 8 mapped to 23
8000	32	8-31 FCoE ports mapped as F_Ports.	0-7	8-11 mapped to 0 12-15 mapped to 1 16-19 mapped to 2 20-23 mapped to 3 24-27 mapped to 4 28-31 mapped to 5

Adding F_Ports to an N_Port

You can modify the default port mapping by adding F_Ports to an N_Port. Adding an F_Port to an N_Port routes that traffic to and from the fabric through the specified N_Port.

You can assign an F_Port to only one primary N_Port at a time. If the F_Port is already assigned to an N_Port, you must first remove it from the N_Port before you can add it to a different N_Port.

Use the following steps to add an F_Port to an N_Port.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag** command with the **--mapadd** *n_portnumber* "*f_port1*;*f_port2*;*...*" operand to add the list of F_Ports to the N_Port.

The *f_portlist* can contain multiple F_Port numbers separated by semicolons, for example "17;18".

```
switch:admin> ag --mapadd 13 "6;7"
F-Port to N-Port mapping has been updated successfully
```

3. Enter the **ag --mapshow** command and specify the port number to display the list of mapped F_Ports. Verify that the added F_Ports appear in the list.

```
switch:admin> ag --mapshow 13

N_Port           : 13
Failover(1=enabled/0=disabled) : 1
Failback(1=enabled/0=disabled) : 1
Current F_Ports  : None
Configured F_Ports : 6;7
PG_ID           : 0
PG_Name         : pg0
```

Removing F_Ports from N_Ports

1. Connect to the switch and log in using an account assigned to the admin role.
2. Remove any preferred secondary N_Port settings for the F_Port. Refer to [“Deleting F_Ports from a preferred secondary N_Port”](#) on page 46 for instructions.
3. Enter the **ag --mapdel N_Port** command with the “[fprot;fport]” option to remove the F_Port from the N_Port.

The *f_portlist* can contain multiple F_Port numbers separated by semicolons, for example “17;18”.

```
switch:admin> ag --mapdel 17;18
F-Port to N-Port mapping has been updated successfully
```

4. Enter the **switchshow** command to verify that the F_Port is free (unassigned).

Unassigned F_Port status is Disabled (No mapping for F_Port). See port 6 in the following example.

```
switch:admin> switchshow
switchName:      fsw534_4016
switchType:      45.0
switchState:     Online
switchMode:      Access Gateway Mode
switchWwn:       10:00:00:05:1e:02:1d:b0
switchBeacon:    OFF
```

Area	Port	Media	Speed	State	Proto
0	0	cu	AN	No_Sync	
1	1	cu	AN	No_Sync	Disabled (N-Port Offline for F-Port)
2	2	cu	AN	No_Sync	Disabled (N-Port Offline for F-Port)
3	3	cu	AN	No_Sync	Disabled (N-Port Offline for F-Port)
4	4	cu	AN	No_Sync	Disabled (N-Port Offline for F-Port)
5	5	cu	AN	No_Sync	Disabled (N-Port Offline for F-Port)
6	6	cu	AN	No_Sync	Disabled (No mapping for F-Port)
7	7	cu	AN	No_Sync	
8	8	cu	AN	No_Sync	
9	9	cu	AN	No_Sync	
10	10	--	N4	No_Module	
11	11	--	N4	No_Module	
12	12	--	N4	No_Module	
13	13	id	N4	Online	N-Port 10:00:00:05:1e:35:10:1e 0x5a0a00
14	14	id	N4	Online	N-Port 10:00:00:05:1e:35:10:1e 0x5a0900
15	15	id	N4	Online	N-Port 10:00:00:05:1e:35:10:1e 0x5a0800

Device-based mapping

This feature allows you to map individual N_Port ID virtualization (NPIV) devices to N_Ports. By mapping device WWNs directly to an N_Port group (recommended) or specific N_Ports, traffic from the device will always go to the same N_Port or N_Port group, regardless of the F_Port where the device logs in. When Port Grouping Policy and WWN Load Balancing mode is enabled for a port group, WWNs mapped to that port group are automatically balanced among the online N_Ports in that group (refer to [“Port Grouping policy modes”](#) on page 36).

NOTE

Port Grouping Policy is not supported when both Automatic Login Balancing and Device Load Balancing are enabled.

Device-based mapping does not affect or replace the traditional port mapping. Device mapping is an optional mapping that will exist on top of existing port mapping. In general mapping devices to N_Port groups is recommended over mapping devices to individual N_Ports within a port group. This ensures maximum device “up-time,” especially during fail-over conditions and system power up. This is especially true when a reasonably large number of devices must connect to the same fabric through a single port group.

These aspects of device mapping are important to note:

- Logins from a device mapped to a specific N_Port or N_Port group (device mapping) always have priority over unmapped devices that log into an F_Port that has been mapped to the same N_Port or N_Port group (port mapping).
- Current device routing (dynamic mapping) may turn out different than your intended mapping (static mapping), depending on which N_Ports are online and which policies are enabled (for example, automatic port configuration, device load balancing, failover, or failback). Therefore, it is recommended to map devices to N_Port groups instead of specific N_Ports within a port group when using device mapping.

NOTE

Automatic port configuration and device load balancing cannot be enabled at the same time.

[Figure 5](#) on page 17 illustrates an example of device mapping to port groups. In the example, WWNs 1, 2, and 3 can connect to any N_Port in Port Group 1 (PG1), while WWNs 4 and 5 can connect with any N_Port in Port Group 2 (PG2).




FIGURE 5 Example of device mapping to N_Port groups

Figure 6 shows an example of device mapping to specific N_Ports. Note that you can map one or multiple WWNs to one N_Port to allow multiple devices to log in through one N_Port.




FIGURE 6 Example device mapping to an N_Port

Static versus dynamic mapping

Device mapping can be classified as either “static” or “dynamic” as follows:

- Device mapping to an N_Port and to an N_Port Group are considered static. Static mappings persist across reboots and can be saved and restored with Fabric OS **configUpload** and **configDownload** commands.
- Automatic WWN load balancing, if enabled, is considered dynamic. These mappings exist only while a device is logged in. Dynamic mappings cannot be saved or edited by the administrator and do not persist across reboots. Dynamic mapping shows the current mapping for devices as opposed to original static mapping, if one had been specified. If a device is mapped to N_port group, then all mapping is dynamic.

NOTE

These mappings only apply to NPIV devices and cannot redirect devices that are directly attached to Access Gateway, since physically-attached devices use the port maps to connect to the fabric.

Device mapping to port groups (recommended)

Mapping NPIV devices to a port group is an ideal choice when a reasonably sized set of devices must connect to the same group of N_Ports, and you want the flexibility of moving the devices to any available F_Port. This type of mapping is recommended because the device will automatically connect to the least-loaded N_Port in the group if the N_Port to which the device is currently connected goes offline or is not yet online. For more information on port groups, refer to “[Port Grouping policy](#)” on page 33.

Use the following steps to map one or more devices to an N_Port group or remove device mapping from an N_Port group.

1. Connect to the switch and log in using an account assigned to the admin role.
2. To add one or multiple device WWNs to an N_Port group, enter the **ag --addwnpgmapping Port_Group** command with the “[WWN];[WWN]” option.

All the listed device WWNs will use the least loaded N_Port in the port group when they log in, unless a specific device mapping can be used instead. This command can only map devices currently connecting through NPIV.

The following example adds two devices to port group 3.

```
ag --addwnpgmapping 3 "10:00:00:06:2b:0f:71:0c;10:00:00:05:1e:5e:2c:11"
```

To change all currently existing device mappings to a different port group use the **--all** option instead of listing all the WWNs.

The following example changes all the currently mapped devices to use port group 3 instead of the current port group mappings.

```
ag --addwnpgmapping 3 --all
```

3. To remove one or multiple devices to an N_Port group, enter the **ag --delwnpgmapping Port_Group** command with the “[WWN];[WWN]” option.

All the listed devices will stop using the least-loaded N_Port in the group when they log in,

The following example removes mapping for two devices to port group 3.

```
ag --delwnpgmapping 3 "10:00:00:06:2b:0f:71:0c;10:00:00:05:1e:5e:2c:11"
```

To remove all devices mapped to an N_Port group, enter the command with the **--all** option instead of listing all WWNs. All of the devices will cease automatic use of the least loaded port in the port group when they log in. The **--all** option is a shortcut for specifying all of the devices that are already mapped with the **addwnpgmapping** command.

The following example removes all devices mapped to port group 3.

```
ag --delwnpgmapping 3 --all
```

4. Enter the **ag --wnmapshow** command to display the list of WWNs mapped to port groups and verify that the correct devices have been mapped to the desired port group.

Device mapping to N_Ports

Use the following steps to add one or more devices to an N_Port to route all device traffic to and from the device through the specified N_Port. Also use these steps to remove device mapping to an N_Port.

1. Connect to the switch and log in using an account assigned to the admin role.
2. To add one or multiple devices to an N_Port, enter the **ag --addwwnmapping N_Port** command with the “[WWN];[WWN]” option. All the listed device WWNs will use the N_Port if it is available.

The following example adds two devices to N_Port 17.

```
ag --addwwnmapping 17 "10:00:00:06:2b:0f:71:0c;10:00:00:05:1e:5e:2c:11"
```

The --all options edit all the currently existing mappings. none of the --all options have any way to detect what devices are using the switch. This option just edits the mappings that are in the list.

To change all current device mappings to a different N_Port, enter the **ag --addwwnmapping N_Port** command with the --all option.

The following command changes all the existing device mappings to use port 17.

```
ag --addwwnmapping 17 --all
```

3. To remove mapping for one or multiple devices to an N_Port, enter the **ag --delwwnmapping N_Port** command with the “[WWN];[WWN]” option. All the listed device WWNs will no longer try to use the N_Port unless a device logs in through an F_Port that is mapped to the N_Port.

The following example removes two devices from N_Port 17.

```
ag --delwwnmapping 17 "10:00:00:06:2b:0f:71:0c;10:00:00:05:1e:5e:2c:11"
```

To remove all devices currently mapped to an N_Port, enter the **ag --delwwnmapping N_Port** command with the --all option. All the listed devices will no longer try to use the N_Port unless a device logs in through an F_Port that is mapped to the N_Port. The -all option is a shortcut for specifying all of the devices that are already mapped with the **addwwnpgmapping** command.

The following command removes all devices currently mapped to port 17.

```
ag --delwwnmapping 17 --all
```

4. Enter the **ag --wwnmapshow** command to display the list of N_Ports mapped to WWNs and verify that the correct WWNs have been mapped or removed from the desired N_Port(s).

Disabling device mapping

Use the following procedures to disable device mapping for all or only specific devices. These procedures are useful when you want to temporarily disable device mapping, then enable this at a later time without reconfiguring your original mapping. To enable disabled mapping, refer to [“Enabling device mapping”](#) on page 21.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --wwnmappingdisable** with the “[WWN]; [WWN]” option to disable mapping for specific WWNs. The device mappings will be ignored for all the listed device WWNs without removing the entry from the WWN mapping database.

The following example disables device mapping for two WWNs.

```
switch:admin> ag --wwnmappingdisable "10:00:00:06:2b:0f:71:0c;
10:00:00:05:1e:5e:2c:11"
```

Enter the **ag** command with the **ag--wwnmappingdisable** with the **--all** option to disable mapping for all available WWNs. The **-all** option will not affect mappings made in the future, Disabled mappings can be modified without automatically enabling them.

The following example removes device mapping for all available WWNs.

```
switch:admin> ag --wwnmappingdisable --all
```

Enabling device mapping

Use the following steps to enable device mapping for all or specific devices that were previously disabled.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --wwnmappingenable** command with the "[WWN]; [WWN]" option to enable mapping for specific WWNs.

The following example enables two device WWNs.

```
switch:admin> ag --wwnmappingenable "10:00:00:06:2b:0f:71:0c;
10:00:00:05:1e:5e:2c:11"
```

Enter the **ag --wwnmappingenable** with the **--all** option to enable mapping for all currently available WWNs. The **-all** option will not affect mappings made in the future, any mapping added for a new device, that is one who's mapping is not disabled, will be enabled by default. Disabled mappings can be modified with out automatically enabling them.

The following command enables all previously disabled device mappings.

```
switch:admin> ag --wwnmappingenable --all
```

Display device mapping information

Use the **--wwnmapshow** command to display static and dynamic mapping information about all device WWNs that have been mapped to N_Ports or N_Port groups. For each WWN, this command displays the following:

- WWN - Device WWNs that are mapped to N_Ports
- 1st N_Port - First or primary mapped N_Port (optional)
- 2nd N_Port - Secondary or failover N_Port (optional)
- PG_ID - Port Group ID where the device is mapped (mapped)
- Current - The N_Port that the device is using (none displays if device not logged in)
- Enabled - Indicates whether device mapping is enabled or disabled

Note that new device mappings will only be enabled and display the next time the device logs into the switch.

Pre-provisioning

You can use Fabric OS commands, Web Tools, and Fabric Manager to map devices that do not yet exist. This allows applicable management programs to push configuration changes without worrying about the order in which they are received. For example, if system administrators need to push a set of port group changes and a set of device mapping changes, they could push them in either order without error. This also applies to using Fabric OS commands for device mapping. You could map several devices to a new port group then create the group without error. Removing a device twice can also be accomplished without error.

VMware configuration

To use the device mapping feature for connecting VMware systems, refer to the Technical Brief: *How to Configure NPIV on VMware ESX Server 3.5* at following link:

http://www.brocade.com/downloads/documents/brocade_vmware_technical_briefs/Brocade_NPIV_ESX3.5_WP.pdf

The following is a *summary* of the steps involved.

1. Make sure that virtual port names (VWWPN) of virtual machines (VM) are mapped to the correct port group (or N_Port). Map all VWWPNs to N_Ports to avoid confusion.
2. Make sure all VWWPNs are mapped for LUN access for array-based targets.
3. Make sure to include all VWWPNs in the zone configuration.
4. Zone the server's physical port to the storage device.
5. Finally check the traffic that originates from virtual node PID (VN PID). if configuration is correct, traffic will flow from VN PID.

Failover and Failback considerations

When using device mapping with VMware, the base device initiates PLOGI and PRLI to the target, and then discovers the LUN. The virtual device also initiates a PLOGI and PRLI to the target, but LUN discovery does not occur. Therefore, when the device-mapped port is toggled and failover or failback takes place, traffic will resume from the base device. We recommend one of the following when using device mapping with VMware:

- Targets should also be reachable by the base device so that I/Os can resume if the mapped device fails over and I/Os will move over to the base PID.
- Reboot the server so that it initializes and uses device mapping

Considerations for Access Gateway mapping

This section outlines considerations and limitations for Access Gateway mapping types.

Mapping priority

To avoid potential problems when both port-based and device-based mapping are implemented, AG uses the following priority system when verifying policies to select the N_Port where a FLOGI is routed. Access Gateway considers all available mappings in the following order until one can be used.

1. Static device mapping to N_Port (if defined)
2. Device mapping to N_Port group (if defined)
For more information, refer to [“Port Grouping policy”](#) on page 33.
3. Automatic WWN load balancing within a port group (if enabled)
For more information, refer to [“Port Grouping policy”](#) on page 33.

NOTE

Only NPIV devices can use device mapping and the automatic WWN Load Balancing policy.

NOTE

In Fabric OS v6.4.0, the device load balancing policy is enabled per module rather than per port group.

4. Port mapping to an N_Port
5. Port mapping to an N_Port in a port group (if defined)
For more information, refer to [“Port Grouping policy”](#) on page 33.

Device mapping considerations

Consider the following points when using device mapping:

- If the N_Port is disabled, all devices that are mapped to it will be disabled. Depending on the effective failover policy, the devices will be enabled on other N_Ports.
- Similar to Port-based mappings, device-based mappings are affected by changes to underlying F_Ports. In other words, if an F_Port needs to be taken offline, both the physical device and all virtual nodes behind it will momentarily go offline.
- Once devices are mapped to an N_Port rather than an N_Port group, they cannot be automatically rebalanced to another N_Port if an additional N_Port comes online.
- There can be cases where two NPIV devices logging through the same F_Port are mapped to two different N_Ports that are connected to two different fabrics. In this case, both NPIV devices may be allocated the same PID by their respective fabric. Once Access Gateway detects this condition, it will disable that F_Port, and the event will be logged.

NOTE

Access Gateway algorithms reduce the chances of PID collisions, but they cannot be totally eliminated. In some cases, you may be able to configure your virtual or physical fabrics to further reduce this condition.

- Device mapping is not supported when firmware is downgraded to Fabric OS 6.3.x or lower. You must delete device mappings before downgrading or disable Device Load Balancing.
- Static and dynamic device mapping are only supported on the edge module in a cascaded Access Gateway configuration.
- When mapping devices to a port group, make sure that all ports in the group have the same NPIV login limit. If some ports have a lower login limit than the other ports, and there are many logins to the group, some devices will repeatedly attempt to connect to the device with the lower limit (because it has the fewest logins) and fail to connect.

N_Port configurations

By default, on embedded switches, only the internal ports of Access Gateway are configured as F_Ports. All external ports are configured (locked) as N_Ports. On standalone switches with AG support, a preset number of ports are locked as N_Ports and the rest of the ports operate as standard F_Ports. Although some ports are locked as N_Ports, these ports can be converted to F_Ports. For example, [Figure 7](#) shows a host connected to external ports of an Embedded Switch with the switch in AG mode. To convert a N_Port to an F_Port first remove all the F_Ports that are mapped to that N_Port, then unlock the port from N_Port state. Finally, define a map for the port. It is highly recommended that all F_Ports mapped to the N_Port first be remapped to other N_Ports before that port is converted into F_Port. Also note that if APC policy is enabled, the port conversion is done automatically and no user intervention is necessary. For more information on which ports are locked as N_Ports by default, see [Table 5](#) on page 12.




FIGURE 7 Example of adding an external F_Port (F9) on an embedded switch

NOTE

A switch in Access Gateway mode must have at least one port configured as an N_Port. Therefore, the maximum number of F_Ports that can be mapped to an N_Port is the number of ports on the switch minus one.

Displaying N_Port configurations

1. Connect to the switch and log in using an account assigned to the admin role.

Enter the **portcfgnport** command.

```
switch:admin> portcfgnport

Ports          0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
-----+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---
Locked N_Port  .. .. .. .. .. .. .. .. .. .. ON ON ON ON ON ON
```

Unlocking N_Ports

By default, on embedded switches all external ports are configured in N_Port lock mode when you enable Access Gateway. Access Gateway connects only FCP initiators and targets to the fabric. It does not support other types of ports, such as ISL (inter switch link) ports.

By default, on fabric switches the port types are not locked. Fabric OS Native mode dynamically assigns the port type based on the connected device: F_Ports and FL_Ports for hosts, HBAs, and storage devices; and E_Ports, EX_Ports, and VE_Ports for connections to other switches.

Unlocking the N_Port configuration automatically changes the port to an F_Port. When you unlock an N_Port, the F_Ports are automatically unmapped and disabled.

Following are procedures for unlocking N_Ports that are in locked mode.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **portcfgnport** command.

NOTE

The **portcfgnport** command only works when the Port Grouping policy is enabled.

```
switch:admin> portcfgnport

Ports          0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
-----+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---
Locked N_Port  .. .. .. .. .. .. .. .. .. .. ON ON ON ON ON ON
```

3. Enter the **portcfgnport** command and specify the port number and 0 (zero) to unlock N_Port mode.

```
switch:admin> portcfgnport 10 0
```

Alternatively, to lock a port in N_Port mode, enter the **portcfgnport** and specify the port number and 1.

```
switch:admin> portcfgnport 10 1
```

2 N_Port configurations

Managing Policies and Features in Access Gateway Mode

In this chapter

• Access Gateway policies overview	27
• Advanced Device Security policy	28
• Automatic Port Configuration policy	31
• Port Grouping policy	33
• Device Load Balancing Policy	40
• Persistent ALPA Policy	41
• Failover	44
• Failback	48
• Trunking in Access Gateway mode	50
• Adaptive Networking on Access Gateway	58
• Per Port NPIV login limit	60
• Considerations for the Brocade 8000	60

Access Gateway policies overview

This chapter provides detailed information on all Access Gateway policies. These policies can be used to control various advanced features, such as Failover, Failback, and Trunking when used in Access Gateway mode.

Displaying current policies

You can run the following command to display policies that are currently enabled or disabled on a switch.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --policyshow** command.

The following example shows that Port Grouping, Automatic Port Configuration, and Advanced Device Security policies are enabled.

```
switch:admin> ag --policyshow
Policy_Description      Policy_Name    State
-----
Port Grouping           pg             Enabled
Auto Port Configuration auto           Disabled
Advanced Device Security ads            Enabled
WWN Based Load Balancing wwnloadbalance Disabled
```

Access Gateway policy enforcement matrix

The following table shows which combinations of policies can co-exist with each other.

TABLE 6 Policy enforcement matrix

Policies	Auto Port Configuration	Port Grouping	N_Port Trunking	ADS Policy
Auto Port Configuration	N/A	Cannot co-exist	Can co-exist	Can co-exist
N_Port Grouping	Mutually exclusive	N/A	Can co-exist	Can co-exist
N_Port Trunking	Can co-exist	Can co-exist	N/A	Can co-exist
ADS Policy ¹	Can co-exist	Can co-exist	Can co-exist	N/A
Device Load Balancing ²	Cannot co-exist	Can co-exist	Can co-exist	Can co-exist

1. The ADS policy is not supported when using Device mapping.

2. Device Load Balancing and Automatic Login Balancing cannot be enabled for the same port group.

Advanced Device Security policy

ADS is a security policy that restricts access to the fabric at the AG level to a set of authorized devices. Unauthorized access is rejected and the system logs a RASLOG message. You can configure the list of allowed devices for each F_Port by specifying their Port WWN (PWWN). The ADS policy secures virtual and physical connections to the SAN.

How the ADS policy works

When you enable this policy, it applies to all F_Ports on the AG-enabled module. By default, all devices have access to the fabric on all ports. You can restrict the fabric connectivity to a particular set of devices where AG maintains a per-port allow list for the set of devices whose PWWN you define to log in through an F_Port. You can view the devices with active connections to an F_Port using the **ag --show** command.

NOTE

The **ag --show** command only displays the Core AGs, such as the AGs that are directly connected to fabric. The **agshow --name name** command displays the F_Ports of both the Core and Edge AGs.

Alternatively, the security policy can be established in the Enterprise fabric using the DCC policy. For information on configuring the DCC policy, see [“Enabling the DCC policy on trunk”](#) on page 53. The DCC policy in the Enterprise fabric takes precedence over the ADS policy. It is generally recommended to implement the security policy in the AG module rather than in the main fabric, especially if Failover and Failback policies are enabled.

Enabling and disabling the Advanced Device Security policy

By default, the ADS policy is disabled. When you manually disable the ADS policy, all of the allow lists (global and per-port) are cleared. Before disabling the ADS policy, you should save the configuration using the **configupload** command in case you need this configuration again.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --policyenable ads** command to enable the ADS policy.

```
switch:admin> ag --policyenable ads
The policy ADS is enabled
```

3. Enter the **ag --policydisable ads** command to disable the ADS policy.

```
switch:admin> ag --policydisable ads
The policy ADS is disabled
```

NOTE

Use the **ag --policyshow** command to determine the current status of the ADS policy.

Setting the list of devices allowed to log in

You can determine which devices are allowed to log in on a per F_Port basis by specifying the device's port WWN (PWWN). Lists must be enclosed in double quotation marks. List members must be separated by semicolons. The maximum number of entries in the allowed device list is twice the per port maximum log in count. Replace the WWN list with an asterisk (*) to indicate all access on the specified F_Port list. Replace the F_Port list with an asterisk (*) to add the specified WWNs to all the F_Ports' allow lists. A blank WWN list ("") indicates no access. The ADS policy must be enabled for this command to succeed.

NOTE

Use an asterisk enclosed in quotation marks, "*", to set the Allow list to "All Access" to all F_Ports; use a pair of double quotation marks ("") to set the Allow list to "No Access".

Note the following characteristics of the Allow List:

- The maximum device entries allowed in the Allow List is twice the per port max login count.
 - Each port can be configured to "not allow any device" or "to allow all the devices" to log in.
 - If the ADS policy is enabled, by default, every port is configured to allow all devices to log in.
 - The same Allow List can be specified for more than one F_Port.
1. Connect to the switch and log in using an account assigned to the admin role.
 2. Enter the **ag --adsset** command with the appropriate operands to set the list of devices allowed to log into specific ports. In the following example, ports 1, 10, and 13 are set to "all access."

```
switch:admin> ag --adsset "1;10;13" "*"
WWN list set successfully as the Allow Lists of the F_Port[s]
```

Setting the list of devices not allowed to log in

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --adsset** command with the appropriate operands to set the list of devices not allowed to log into specific ports. In the following example, ports 11 and 12 are set to “no access.”

```
switch:admin > ag --adsset "11;12" ""
WWN list set successfully as the Allow Lists of the F_Port[s]
```

Removing devices from the list of allowed devices

Use the **ag --adsdel** command to delete the specified WWNs from the list of devices allowed to log in to the specified F_Ports. Lists must be enclosed in double quotation marks. List members must be separated by semicolons. Replace the F_Port list with an asterisk (*) to remove the specified WWNs from all the F_Ports' allow lists. The ADS policy must be enabled for this command to succeed.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --adsdel** command to remove one or more devices from the list of allowed devices.

Use the following syntax:

```
ag--adsdel "F_Port [;F_Port2;...]" "WWN [;WWN2;...]"
```

In the following example, two devices are removed from the list of allowed devices (for ports 3 and 9).

```
switch:admin> ag --adsdel "3;9"
"22:03:08:00:88:35:a0:12;22:00:00:e0:8b:88:01:8b"
WWNs removed successfully from Allow Lists of the F_Port[s]Viewing F_Ports
allowed to login
```

Adding new devices to the list of allowed devices

You can add the specified WWNs to the list of devices allowed to log in to the specified F_Ports. Lists must be enclosed in double quotation marks. List members must be separated by semicolons. Replace the F_Port list with an asterisk (*) to add the specified WWNs to all the F_Ports' allow lists. The ADS policy must be enabled for this command to succeed.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --adsadd** command with appropriate operands to add one or more new devices to the list of allowed devices.

Use the following syntax:

```
ag--adsadd "F_Port [;F_Port2;...]" "WWN [;WWN2;...]"
```

In the following example, two devices are added to the list of allowed devices (for ports 3 and 9).

```
switch:admin> ag --adsadd "3;9"
"20:03:08:00:88:35:a0:12;21:00:00:e0:8b:88:01:8b"
WWNs added successfully to Allow Lists of the F_Port[s]
```

Displaying the list of allowed devices on the switch

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --adsshow** command.

```
switch:admin> ag --adsshow
  F_Port      WWNs Allowed
-----
  1           ALL ACCESS
  3           20:03:08:00:88:35:a0:12
              21:00:00:e0:8b:88:01:8b
  9           20:03:08:00:88:35:a0:12
              21:00:00:e0:8b:88:01:8b
 10           ALL ACCESS
 11           NO ACCESS
 12           NO ACCESS
 13           ALL ACCESS
-----
```

ADS policy considerations

The following are considerations for setting the ADS policy:

- In cascading configurations, you should set the ADS policy on the AG module that directly connects to the servers.
- ADS policy can be enabled or disabled independent of status of other AG policies.
- The ADS policy is not currently supported with device-based mapping.

Upgrade and downgrade considerations for the ADS policy

Downgrading to Fabric OS v6.3.0 or earlier is supported.

Upgrading from v6.3.0 to v6.4.0 or downgrading from v6.4.0 to v6.3.0 will not change the APC policy settings.

Automatic Port Configuration policy

APC provides the ability to automatically discover port types (host, target, or fabric) and dynamically update the port maps when a change in port-type connection is detected. This policy is intended for a fully hands-off operation of Access Gateway. APC dynamically maps F_Ports across available N_Ports so they are evenly distributed.

How the APC policy works

When the APC policy is enabled and a port on AG is connected to a Fabric switch, AG configures the port as an N_Port. If a host is connected to a port on AG, then AG configures the port as an F_Port and automatically maps it to an existing N_Port with the least number of F_Ports mapped to it. When the APC policy is enabled, it applies to all ports on the switch.

Enabling and disabling the APC policy

Use the following steps to enable and disable Automatic Port Configuration policy. This policy is disabled by default in Access Gateway.

Enabling APC policy

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **switchdisable** command to ensure that the switch is disabled.
3. Enter the **configupload** command to save the switch's current configuration.
4. Enter the **ag --policydisable pg** command to disable the port grouping policy.
5. Enter the **ag --policyenable auto** command to enable the APC policy.

```
switch:admin> ag --policyenable auto
All Port related Access Gateway configurations will be lost.
Please save the current configuration using configupload.
Do you want to continue? (yes, y, no, n): [no] y
```

6. At the command prompt, type **Y** to enable the policy.
The switch is ready; a reboot is not required.

Disabling APC policy

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **switchdisable** command to ensure that the switch is disabled.
3. Enter the **configupload** command to save the switch's current configuration.
4. Enter the command **ag --policydisable auto** to disable the APC policy.
5. At the command prompt, type **Y** to disable the policy.

```
switch:admin> ag --policydisable auto
Default factory settings will be restored.
Default mappings will come into effect.
Please save the current configuration using configupload.
Do you want to continue? (yes, y, no, n): [no] y
Access Gateway configuration has been restored to factory default
```

6. Enter the **switchenable** command to enable the switch.

Automatic Port Configuration policy considerations

Following are the considerations for the Automatic Port Configuration policy:

- The APC and the PG policies cannot be enabled at the same time. You can still benefit from the automatic port mapping feature of the APC policy when the port grouping policy is enabled by enabling the auto distribution feature for each port group.
- You cannot manually configure port mapping when this policy is enabled.

- The APC policy applies to all ports on the switch. Enabling the APC policy is disruptive and erases all existing port-based mappings. Therefore, before enabling the APC policy, you should disable the AG module. When you disable the APC policy, the N_Port configuration and the port-based mapping revert back to the default factory configurations for that platform. It is recommended that before you either disable or enable APC policy to save the current configuration file using the **configupload** command in case you might need this configuration again.

Upgrade and downgrade considerations for the APC policy

The following are supported:

- Downgrading to a Fabric OS level that supports the APC policy.
- Upgrading from Fabric OS v6.3.0 to Fabric OS v6.4.0 will maintain the policy that was enabled in Fabric OS 6.3.0.

Port Grouping policy

Use the PG policy to partition the fabric, host, or target ports within an AG-enabled module into independently operated groups. Use the PG policy in the following situations:

- When connecting the AG module to multiple physical or virtual fabrics.
- When you want to isolate specific hosts to specific fabric ports for performance, security, or other reasons.

How port groups work

Create port groups using the **ag --pgcreate** command. This command groups N_Ports together as “port groups.” By default, any F_Ports mapped to the N_Ports belonging to a port group will become members of that port group. Port grouping fundamentally restricts failover of F_Ports to the N_Ports that belong to that group. For this reason an N_Port cannot be member of two port groups. The default PGO group contains all N_Ports that do not belong to any other port groups.

[Figure 8](#) on page 34 shows that if you have created port groups and then an N_Port goes offline, the F_Ports being routed through that port will fail over to any of the N_Ports that are part of that port group and are currently online. For example, if N_Port 4 goes offline then F_Ports 7 and 8 are routed through to N_Port 3 as long as N_Port 3 is online because both N_Ports 3 and 4 belong to the same port group, PG2. If no active N_Ports are available, the F_Ports are disabled. The F_Ports belonging to a port group do not fail over to N_Ports belonging to another port group.

3 Port Grouping policy




FIGURE 8 Port grouping behavior

When a dual redundant fabric configuration is used, F_Ports connected to a switch in AG mode can access the same target devices from both of the fabrics. In this case, you must group the N_Ports connected to the redundant fabric into a single port group. It is recommended to have paths fail over to the redundant fabric when the primary fabric goes down. Refer to [Figure 9](#).




FIGURE 9 Port group 1 (pg1) setup

Adding an N_Port to a port group

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --pgadd** command with the appropriate operands to add an N_Port to a specific port group. In the following example N_Port 14 is added to port group 3.

Note that if you add more than one N_Ports, you must separate them with a semicolon.

```
switch:admin> ag --pgadd 3 14
N_Port[s] are added to the port group 3
```

Deleting an N_Port from a port group

Before deleting an N_Port, all F_Ports mapped to that N_Port should be remapped before that N_Port is deleted from a port group.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --pgdel** command with the appropriate operands to delete an N_Port from a specific port group. In the following example, N_Port 13 is removed from port group 3.

```
switch:admin> ag --pgdel 3 13
N_Port[s] are deleted from port group 3
```

3. Enter the command **ag --pgshow** to verify the N_Port was deleted from the specified port group.

```
switch:admin> ag --pgshow
PG_ID PG_Name      PG_Mode  N_Ports  F_Ports
-----
0      pg0             lb,mfnm  1;3      10;11
2      SecondFabric    -        0;2      4;5;6
-----
```

Removing a port group

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --pgremove** command with appropriate operands to remove a port group. In the following example, port group 3 is removed.

```
switch:admin> ag --pgremove 3
Port Group 3 has been removed successfully
```

Renaming a port group

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --pgrename** command with appropriate operands to rename port group. In the following example, port group pgid 2 is renamed to MyEvenFabric.

```
switch:admin> ag --pgrename 2 MyEvenFabric
Port Group 2 has been renamed as MyEvenFabric successfully
```

Disabling the Port Grouping policy

The Port Grouping (PG) policy is enabled by default for Access Gateway. To disable this policy, use the following steps.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --policydisable** command.

```
switch:admin> ag --policydisable pg.
```

Port Grouping policy modes

You can enable and disable the following Port Grouping policy modes when you create port groups using the **pgcreate** command. Alternately, you can enable these policies using the **ag-pgsetmodes** command.

Automatic Login Balancing

If Automatic Login Balancing mode is enabled for a port group and an F_Port goes offline, logins in the port group are redistributed among the remaining F_Ports. Similarly, if an N_Port comes online, port logins in the PG are redistributed to maintain a balanced N_Port-to-F_Port ratio.

Considerations for Automatic Login Balancing

Please consider the following facts about this feature:

- Automatic Login Balancing is disruptive. However, you can minimize disruption by disabling or enabling rebalancing of F_Ports on F_Port offline or N_Port online events. Refer to [“Rebalancing F_Ports”](#) on page 37.
- You must explicitly enable Automatic Login Balancing on a port group.
- If an N_Port is deleted from a port group enabled for Automatic Login Balancing, the F_Ports mapped to that N_Port stay with the port group as long as there are other N_Ports in the group. Only the N_Port is removed from the port group. This is because the F_Ports are logically associated with the port groups that are enabled for Login Balancing. This is not the case for port groups not enabled for Automatic Login Balancing. When you delete an N_Port from one of these port groups, the F_Ports that are mapped to the N_Port move to PGO along with the N_Port. This is because the F_Ports are logically associated with the N_Ports in port groups not enabled for Login Balancing.

Managed Fabric Name Monitoring (MFNM)

Fabric Name Monitoring mode automatically detects whether all the N_Ports within a port group are physically connected to the same physical or virtual fabric. Once a misconnection is detected there are two methods to handle it, depending on the operating mode. For “default” mode a message is logged into RASLOG. For “managed” mode (MFNM), automatic failover disables on all N_Ports within the N_Port group, and a message displays in the RAS log about multiple fabrics.

In both default and managed mode, the system queries the fabric name once every 120 seconds to detect inconsistencies such as a port group being connected to multiple fabrics. You can configure the monitoring timeout value to something other than 120 seconds using the **ag-pgfnmtov** command. Refer to [“Setting the current fabric name monitoring timeout value”](#) on page 39.

The **ag-pgfnmtov** command is blocked on a Brocade 8000 during and after creation of a port group.

Creating a port group and enabling Automatic Login Balancing mode

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --pgcreate** command with appropriate operands to create a port group. In the following example, a port group named “FirstFabric” is created that includes N_Ports 1 and 3 and has automatic login balancing (lb) enabled.

```
switch:admin> ag --pgcreate 3 "1;3" -n FirstFabric1 -m "lb"
```



```
Port Group 3 created successfully
```

3. Enter the **ag --pgshow** command to verify the port group was created.

```
switch:admin> ag --pgshow
PG_ID PG_Name      PG_Mode  N_Ports  F_Ports
-----
0      pg0           lb,mfnn  none     none
2      SecondFabric  -        0;2      4;5;6
3      FirstFabric   lb       1;3      10;11
```

Rebalancing F_Ports

To minimize disruption that could occur once F_Ports go offline or when additional N_Ports are brought online you can modify the default behavior of the automatic login balancing feature by disabling or enabling rebalancing of F_Ports when F_Port offline or N_Port online events occur.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **agautomapbalance --enable** command with appropriate operands to enable automatic login redistribution of F_Ports. In the following example, rebalancing of F_Ports in port group 1 in Access Gateway is enabled when an F_Port online event occurs.

```
switch:admin> agautomapbalance --enable -fport -pg 1
```

3. Enter the **agautomapbalance --disable -all** command with appropriate operands to disable automatic login distribution of N_Ports for all PGs in the Access Gateway when an N_Port online event occurs.

```
switch:admin> agautomapbalance --disable -nport -all
```

4. Enter the **agautomapbalance --disable -all** command with appropriate operands to disable automatic login distribution of F_Ports for all port groups in the Access Gateway when an F_Port online event occurs.

```
switch:admin> agautomapbalance --disable -fport -all
```

5. Enter the **agautomapbalance --show** command to display the automatic login redistribution settings for port groups. In the following example, there are two port groups, 0 and 1.

```
switch:admin> agautomapbalance --show

AG Policy: pg
-----
PG_ID LB mode nport fport
-----
0 Enabled    Enabled    Disabled
1 Disabled   -         -
```

This command also displays the automatic login redistribution settings for N_Ports and F_Ports as shown in the following example.

```
switch:admin> agautomapbalance --show

-----
AG Policy: Auto
```

```
-----
automapbalance on N_Port Online Event: Disabled
automapbalance on F_Port Offline Event: Enabled
-----
```

Considerations when modifying automatic login balancing

Consider the following when disabling automatic login balancing:

- Be aware that modifying the APC policy default setting using the **agautomapbalance** command may yield to uneven distribution of F_Ports to N_Ports. In such cases you may want to consider a manual login distribution that forces a rebalancing of F_Ports to N_Ports.
- To control automatic rebalancing to avoid disruptions when the Port Grouping policy is enabled, refer to [“Rebalancing F_Ports”](#) on page 37.

Enabling Managed Fabric Name Monitoring mode

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --pgsetmodes** command with appropriate operands to enable MFNM mode. In the following example, MFNM mode is enabled for port group 3.

```
switch:admin> ag --pgsetmodes 3 "mfnm"
Managed Fabric Name Monitoring mode has been enabled for Port Group 3
```

Disabling Managed Fabric Name Monitoring mode

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --pgdelmodes** command with appropriate operands to disable MFNM mode. In the following example, MFNM mode is disabled for port group 3.

```
switch:admin> ag --pgdelmodes 3 "mfnm"
Managed Fabric Name Monitoring mode has been disabled for Port Group 3
switch:admin> ag --pgshow
PG_ID PG_Name PG_Mode N_Ports F_Ports
-----
0 pg0 lb,mfnm 0;2 4;5;6
3 FirstFabric lb 1;3 10;11
-----
```

Displaying the current fabric name monitoring timeout value

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --pgfnmtov** command.

```
switch:admin> ag --pgfnmtov

Fabric Name Monitoring TOV: 120 seconds
```

Setting the current fabric name monitoring timeout value

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --pgfnmtov** command, followed by a value.

```
switch:admin> ag --pgfnmtov 100
```

This sets the timeout value to 100 seconds.

NOTE

The **pgfnmtov** command is blocked on a Brocade 8000.

Port Grouping policy considerations

Following are the considerations for the Port Grouping policy:

- A port cannot be a member of more than one port group.
- The PG policy is enabled by default in Fabric OS 6.0 and higher. A default port group “0” (PG0) is created, which contains all ports on the AG.
- APC policy and PG policy are mutually exclusive. You cannot enable these policies at the same time.
- If an N_Port is added to a port group or deleted from a port group and login balancing is enabled or disabled for the port group, the N_Port maintains its original failover or failback setting. If an N_Port is deleted from a port group, it automatically gets added to port group 0.
- When specifying a preferred secondary N_Port for a port group, the N_Port must be from the same group. If you specify an N_Port as a preferred secondary N_Port and it already belongs to another port group, the operation fails. Therefore, it is recommended to form groups before defining the preferred secondary path.
- If the PG policy is disabled while a switch in AG mode is online, all the defined port groups are deleted, but the port mapping remains unchanged. Before disabling the PG policy, you should save the configuration using the **configupload** command in case you might need this configuration again.
- If N_Ports connected to unrelated fabrics are grouped together, N_Port failover within a port group can cause the F_Ports to connect to a different fabric and the F_Ports may lose connectivity to the targets they were connected to before the failover, thus causing I/O disruption as shown in [Figure 9](#) on page 34. Ensure that the port group mode is set to [Managed Fabric Name Monitoring \(MFNM\)](#) mode. This monitors the port group to detect connection to multiple fabrics and disables failover of the N-ports in the port group. For more information on MFNM, refer to [“Enabling Managed Fabric Name Monitoring mode”](#) on page 38.

Upgrade and downgrade considerations for the Port Grouping policy

Downgrading to Fabric OS v6.3.0 or earlier is supported. Note the following considerations when upgrading and downgrading from Fabric OS v6.4.0 to Fabric OS v6.3.0 and earlier:

- When upgrading to Fabric OS v6.4.0, the PG policy that was enforced in Fabric OS v6.3.0 continues to be enforced in Fabric OS v6.4.0 and the port groups are retained. You should save the configuration file using the **configupload** command in case you might need this configuration again.
- If you upgrade from Fabric OS 5.3.0 to 6.0 or higher, you will not see any change in device behavior where the Port Grouping policy is enabled by default.

Device Load Balancing Policy

When Device Load Balancing is enabled, devices mapped to a port group always log into the least-loaded N_Port in that port group. This helps to distribute the login load on each of the N_Ports. This policy is intended for use in conjunction with device-based mapping. It provides an automatic approach to mapping devices to the least loaded N_Port within an N_Port group. To effectively use this policy, we recommend that you map devices to desired N_Port groups before enabling this policy. The Port Grouping policy must be enabled before you can enable Device Load Balancing.

Manually created mappings from devices an N_Port take precedence over automatically created mappings. Refer to [“Mapping priority”](#) on page 22 for details on connection priority for AG port mapping. For more information on device mapping, refer to [“Device-based mapping”](#) on page 15.

Enabling WWN Load Balancing

Use the following steps to enable Device Load Balancing.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **configupload** command to save the switch's current configuration.
3. The Port Grouping policy must be enabled to enable Device Load Balancing. Enter the **ag --policyshow** command to determine if the Port Grouping policy is enabled. If it is not enabled, enter **ag --policyenable pg** to enable this policy.
4. Enter the **ag --policyenable wwnloadbalance** command to enable the Device Load Balancing policy. Note that since in Fibre Channel devices are identified by their WWNs, CLI commands use device WWNs.

Disabling Device Load Balancing

Before disabling this policy, you should save the configuration using the **configupload** command in case you need this configuration again.

1. Connect to the switch and log in using an account assigned to the admin role.

2. Enter the **ag --policydisable wwnloadbalance** command to enable the Device Load Balancing policy.

```
switch:admin> ag --policydisable wwnloadbalance
The policy WWN load balancing is disabled
```

NOTE

Use the **ag --policyshow** command to determine the current status of the WWN Load Balancing policy.

Device Load Balancing considerations

- This policy should be enabled on the edge AG of a cascaded AG configuration.
- This policy is not applicable on a port group when the APC policy or Automatic Login Balancing are enabled.
- This policy is not supported on the Brocade 8000 switch for Fabric OS v6.4.0. This is because MFNM is enabled on the default port group and any created port groups on the Brocade 8000. As a result, the **pgsetmodes**, **pgdelmodes**, and **pgcreate** commands are blocked for the **-m** option, and Automatic Login Balancing cannot be enabled.
- If a device is mapped to a port that is currently part of a trunk, then the device will use that trunk. When trunking is used with Device Load Balancing Policy, then the load on each trunk will be proportional to the number of ports in that trunk. Use the **ag -show** command to determine the devices using a particular trunk.
- When using this policy make sure that all ports in the port group have the same NPIV login limit. If some ports have a lower login limit than the other ports, and there are many logins to the group, some devices will repeatedly attempt to connect to the device with the lower limit (because it has the fewest logins) and fail to connect.

Persistent ALPA Policy

This policy is meant for host systems with operating systems that cannot handle different PID addresses across login sessions when booting over SAN. The persistent ALPA policy for switches in Access Gateway mode lets you configure the AG module so that the host is more likely to get the same PID when it logs out of and into the same F_Port. Since the ALPA field makes up a portion of PID, the PID may possibly change across switch module or the server power cycles. This policy, if enabled, will help reduce the chances of a different PID issued for the same host.

The benefit of this feature is that it will ensure a host has the same ALPA on the F_Ports though the host power cycle. You may also achieve the same behavior and benefit by setting the same policy in the main (core) fabric. When this feature is enabled, AG will request the same ALPA from the core fabric. However, depending on the Fabric, this request may be denied. When this occurs, the host is assigned a different ALPA. One of the following settings deal with this situation:

- In “Flexible” mode the AG only log s an event that it did not receive the same ALPA from the core fabric and continues bringing up the device with the new ALPA.
- In the “Stringent” mode, if the requested ALPA is not available, the server login will be rejected and the server port will not be able to log in into the fabric.

Enabling Persistent ALPA

By default, Persistent ALPA is disabled. You can enable Persistent ALPA using the **ag --persistentalpaenable** command with the following syntax and with one of the following value types:

```
ag -persistentalpaenable 1/0[On/Off] -s/-f[Stringent/Flexible]
```

- **Flexible ALPA** assigns an unassigned ALPA value when the ALPA assigned to the device is taken by another host.
- **Stringent ALPA** causes the host login request to be rejected by AG if assignment of the same ALPA is not possible.

To enable Persistent ALPA, use the following steps.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --persistentalpaenable** command to enable persistent ALPA in flexible or stringent mode.

```
switch:admin> ag --persistentalpaenable 1 -s/-f
```

To ensure consistency among the different devices, after Persistent ALPA is enabled, all the ALPAs become persistent whether they were logged in before the Persistent ALPA feature was enabled or not.

Disabling Persistent ALPA

When you disable this feature, do not specify the value type, for example flexible ALPA or stringent ALPA. Use the following steps.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --persistentalpadisable** command.

```
switch:admin> ag --persistentalpaenable 0
```

Persistent ALPA device data

Access Gateway uses a table to maintain a list of available and used ALPAs. When the number of entries in this table is exhausted, the host receives an error message. You can remove some of the entries to make space using instructions under [“Removing device data from the database”](#) next.

Removing device data from the database

Use the following steps to remove device data from the database.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --deletepwwnfromdb** command.

```
switch:admin> ag --deletepwwnfromdb PWWN
```

Displaying device data

You can view the ALPA of the host related to any ports you delete from the database.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --printalpamap** command with the appropriate operand to display a database entry for a specific F_Port. The following example will display an entry for F_Port 2.

```
switch:admin> ag --printalpamap 2
```

Clearing ALPA values

You can clear the ALPA values for a specific port.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --clearalpamap** command with the appropriate operand to remove the PWW-to-ALPA mapping for a specific port. In the following example, the mapping for port 2 is cleared from the database.

```
switch:admin> ag --clearalpamap 2
```

NOTE

All the data must be persistent in case of a reboot. During a reboot, the tables will be dumped to the persistent_NPIV_config file.

Persistent ALPA policy considerations

The Persistent ALPA feature is not supported in the following:

- When AG N_Ports are connected to the shared ports of 48-port Director blades
- CISCO fabrics. Enable Persistent FCID mode on the connecting Cisco switch to achieve the same functionality.
- Persistent ALPA configuration will not change to the default when the **configdefault** command is used, but will retain the previous configuration.

Upgrade and downgrade considerations for Persistent ALPA

Downgrading to Fabric OS v6.2.X or earlier is not supported. When downgrading to Fabric OS v6.2.X or earlier, if the Persistent ALPA feature is enabled, clear all the data from the database, and then disable this feature before downgrading. For information on how to clear data from the database, see [“Removing device data from the database”](#) on page 42.

Failover

Access Gateway Failover ensures maximum uptime for the servers. When a port is configured as an N_Port, failover is enabled by default and is enforced during power-up. Failover allows hosts and targets to automatically remap to another online N_Port if the primary N_Port goes offline.

NOTE

For port-based mapping, the Failover policy must be enabled on an N_Port for failover to occur. For device-based mapping, if a device is mapped to an N_Port in a port group, the device will always reconnect to the least-loaded online N_Port in the group (or secondary N_Port in the group if configured) if the primary N_Port goes offline. This occurs regardless of whether the Failover policy is enabled or disabled for the primary N_Port.

Failover with port-based mapping

The Failover allows F_Ports to automatically remap to an online N_Port if the primary N_Port goes offline. If multiple N_Ports are available for failover, the failover policy evenly distributes the F_Ports to available N_Ports belonging to the same N_Port group. If no other N_Port is available, failover does not occur and the F_Ports mapped to the primary N_Port go offline as well.

AG provides an option to specify a secondary failover N_Port for an F_Port.

Failover configurations in Access Gateway

The following sequence describes how a failover event occurs:

- An N_Port goes offline.
- All F_Ports mapped to that N_Port are temporarily disabled.
- If the Failover policy is enabled on an offline N_Port, the F_Ports mapped to it will be distributed among available online N_Ports. If a secondary N_Port is defined for any of these F_Ports, these F_Ports will be mapped to those N_Ports. If port group policy is enabled, then the F_Ports only fail over to N_Ports that belong to the same port group as the originally offline N_Port.

Example : Failover

This example shows the failover behavior in a scenario where two fabric ports go offline, one after the other. Note that this example assumes that no preferred secondary N_Port is set for any of the F_Ports.

- First the Edge switch F_A1 port goes offline, as shown in [Figure 10](#) on page 45 Example 1 (left), causing the corresponding Access Gateway N_1 port to be disabled.
The ports mapped to N_1 fail over; F_1 fails over to N_2 and F_2 fails over to N_3.
- Next the F_A2 port goes offline, as shown in [Figure 10](#) on page 45 Example 2 (right), causing the corresponding Access Gateway N_2 port to be disabled.
The ports mapped to N_2 (F_1, F_3, and F_4) fail over to N_3 and N_4. Note that the F_Ports are evenly distributed to the remaining online N_Ports and that the F_2 port did not participate in the failover event.

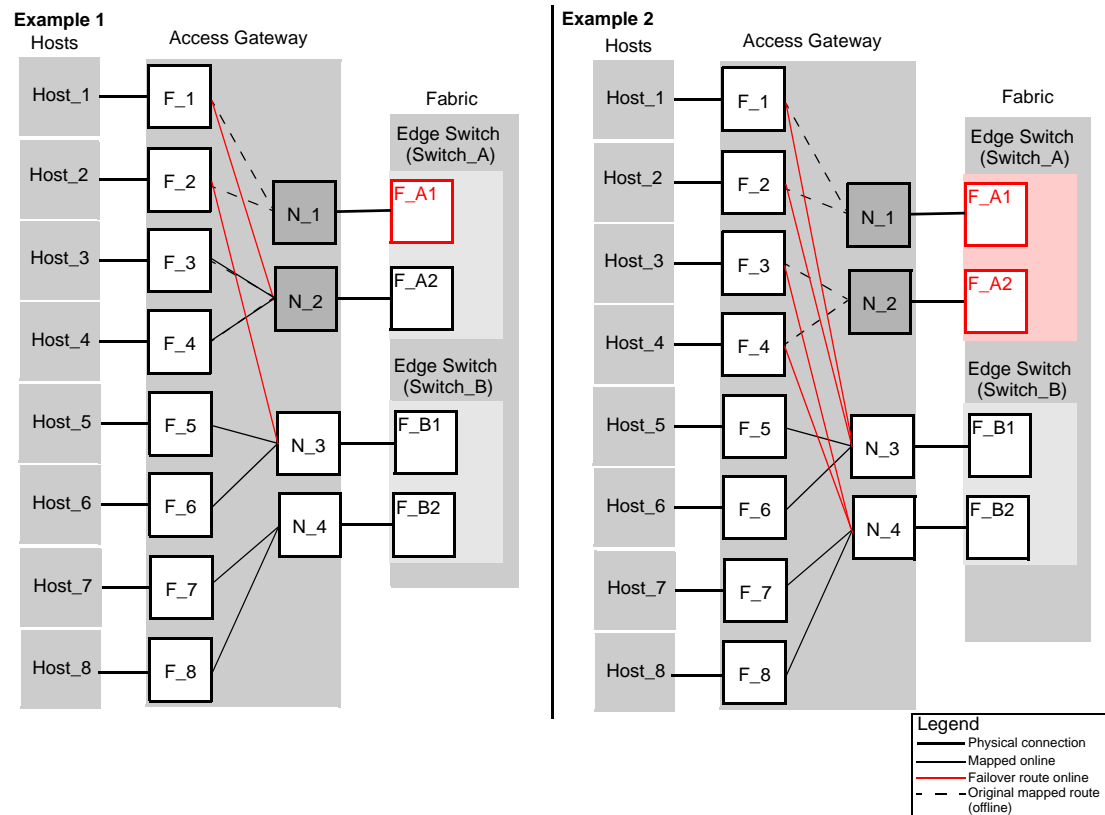


FIGURE 10 Example 1 and 2 Failover behavior

Adding a preferred secondary N_Port (optional)

F_Ports automatically fail over to any available N_Port. Alternatively, you can specify a preferred secondary N_Port in case the primary N_Port fails. If the primary N_Port goes offline, the F_Ports fail over to the preferred secondary N_Port (if it is online), then re-enable. If the secondary N_Port is offline, the F_Ports will disable. Define the preferred secondary N_Ports per F_Port. For example, if two F_Ports are mapped to a primary N_Port 1, you can define a secondary N_Port for one of those F_Ports and not define a secondary N_Port for the other F_Port. F_Ports must have a primary N_Port mapped before a secondary N_Port can be configured.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --prefset** command with the "*F_Port1;F_Port2; ...*" N_Port operands to add the preferred secondary F_Ports to the specified N_Port.

The F_Ports must be enclosed in quotation marks and the port numbers must be separated by a semicolon, for example:

```
switch:admin> ag --prefset "3,9" 4
Preferred N_Port is set successfully for the F_Port[s]
```

NOTE

Preferred mapping is not allowed when automatic login balancing mode is enabled for a port group. All N_Ports are the same when automatic login balancing is enabled.

Deleting F_Ports from a preferred secondary N_Port

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --prefdel** command with the "F_Port1;F_Port2;..." N_Port operands to delete F_Ports from an N_Port.

The list of F_Ports must be enclosed in quotation marks. Port numbers must be separated by a semicolon. In the following example, F_Ports 3 and 9 are deleted from preferred secondary N_Port 4.

```
switch:admin> ag --prefdel "3;9" 4
Preferred N_Port is deleted successfully for the F_Port[s]
```

Failover with device-based mapping

Failover is handled similarly for port-based and device-based mapping, if devices are mapped to N_Port groups. If a device is mapped to an N_Port in a group, and an N_Port goes offline, the devices mapped to that N_Port will reconnect on the least loaded online N_Ports in the group.

Enabling or disabling Failover or Failback policies for N_Ports have no effect on device-based mapping. A device will always fail over to an online N_Port in the port group, regardless of whether Failback is enabled for an N_Port or not. Whereas, with port-based mapping, if you disable the Failover or Failback policy on an N_Port, the F_Port will not failover or failback to other N_Ports.

Failover behavior is different if a device is mapped to an specific N_Port instead of to a N_Port group. If mapping a device to a specific N_Port, you can define a secondary N_Port that will be used if the primary N_Port is offline. To maximize the device uptime it is recommended to map the device to a port group rather than to specific N_Ports.

Adding a preferred secondary N-Port for device mapping (optional)

Use the following steps to configure a secondary N_Port where devices will connect if their first or primary N_Port, if defined, is unavailable.

1. Connect to the switch and log in using an account assigned to the admin role.
2. To configure an N_Port as a failover port for one or multiple devices mapped to a specific N_Port, enter the **ag --addwwnfailovermapping N_Port** command with the "[WWN];[WWN]" option. All of the listed device WWNs will use the listed N_Port if it is available and the first mapped N_Port is unavailable.

The following example configures N_Port 32 as the failover port for two devices already mapped to a primary N_Port.

```
ag --addwwnfailovermapping 32
"10:00:00:06:2b:0f:71:0c;10:00:00:05:1e:5e:2c:11"
```

To configure N_Port 32 as a failover port for all WWNs mapped to the N_Port, enter the **ag --delwwnfailovermapping N_Port** command with the **--all** operand.

```
ag --delwwnfailovermapping 32--all
```

Deleting a preferred secondary N_Port for device mapping (optional)

Use the following steps to remove a secondary N_Port where devices will connect if their first or primary N_Port, if defined, is unavailable.

1. Connect to the switch and log in using an account assigned to the admin role.
2. To delete an N_Port configured as a failover port for one or multiple devices mapped to a specific N_Port, enter the **ag --delwwnfailovermapping N_Port** command with the “[WWN];[WWN]” option. All of the listed devices will stop using the N_Port if the first N_Port mapped to the devices is unavailable unless they log in through F_Ports that are mapped to the N_Port.

The following example removes N_Port 32 as the secondary N_Port for two devices already mapped to a primary N_Port.

```
ag --delwwnfailovermapping 32
"10:00:00:06:2b:0f:71:0c;10:00:00:05:1e:5e:2c:11"
```

To remove an N_Port as a failover port for all devices mapped to the N_Port, enter the **ag --delwwnfailovermapping N_Port** command with the **--all** option.

The following command removes N_Port 32 as the secondary N_Port for all available devices.

```
ag --delwwnfailovermapping 32--all
```

Enabling and disabling Failover on a N_Port

Use the following steps to enable or disable failover policy on a specific N_Port.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --failovershow N_Port** command to display the failover setting.

```
switch:admin> ag --failovershow 13
Failover on N_Port 13 is not supported
```

3. Enter the **ag --failoverenable N_Port** command to enable failover.

```
switch:admin> ag --failoverenable 13
Failover policy is enabled for port 13
```

4. Enter the **ag --failoverdisable N_Port** command to disable failover.

```
switch:admin> ag --failoverdisable 13
Failover policy is disabled for port 13
```

Enabling and disabling Failover for a port group

Failover policy can be enabled on a port group. To enable or disable use the following steps to enable or disable failover on all the N_Ports belonging to the same port group.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --failoverenable -pg pgid** command to enable failover.

```
switch:admin> ag --failoverenable -pg 3
Failover policy is enabled for port group 3
```

3. Enter the **ag --failoverdisable -pg pgid** command to disable failover.

```
switch:admin> ag --failoverdisable -pg 3
Failover policy is disabled for port group 3
```

Upgrade and downgrade considerations for Failover

Consider the following when upgrading or downgrading Fabric OS versions.

- Downgrading to Fabric OS v6.3.0 or earlier is supported.
- Upgrading from v6.3.0 to v6.4.0 or downgrading from v6.4.0 to v6.3.0 will not change failover settings.

Failback

Failback policy provides a means for hosts that have failed over to automatically reroute back to their intended mapped N_Ports when these N_Ports come back online. Failback policy is an attribute of an N_Port and is enabled by default when a port is locked to the N_Port.

Only the originally mapped F_Ports fail back. In the case of multiple N_Port failures, only F_Ports that were mapped to a recovered N_Port experience failback. The remaining F_Ports are not redistributed.

NOTE

For port-based mapping, the Failback policy must be enabled on an N_Port for failback to occur. For device-based mapping, the Failback policy has no effect. If a device is mapped to a port group, it will always fail over to an online N_Port in the port group (or secondary N_Port if configured) and will remain connected to this failover N_Port when the original N_Port comes back online.

Failback configurations in Access Gateway

The following sequence describes how a failback event occurs:

- When an N_Port comes back online, with Failback enabled, the F_Ports that were originally mapped to it are temporarily disabled.
- The F_Port is rerouted to the primary mapped N_Port, and then re-enabled.
- The host establishes a new connection with the fabric.

NOTE

The failback period is quite fast and rarely causes an I/O error at the application level.

Example : Failback

In Example 3, described in [Figure 11](#) on page 49, the Access Gateway N_1 remains disabled because the corresponding F_A1 port is offline. However, N_2 comes back online. See [Figure 10](#) on page 45 for the original fail over scenario.

The ports F_1 and F_2 are mapped to N_1 and continue routing to N_3. Ports F_3 and F_4, which were originally mapped to N_2, are disabled and rerouted to N_2, and then enabled.

Example 3

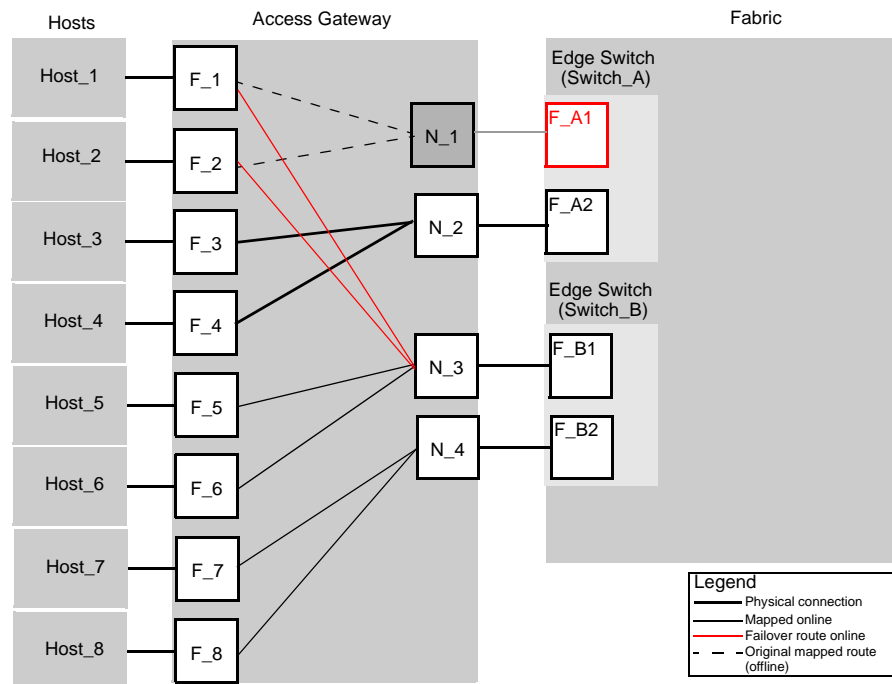


FIGURE 11 Failback behavior

Enabling and disabling Failback on an N_Port

Use the following steps to enable or disable Failback on N_Ports.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **ag --failbackshow n_portnumber** command to display the failover setting.

```
switch:admin> ag --failbackshow 13
Failback on N_Port 13 is not supported
```

3. Use the following commands to enable or disable Failback:

- Enter the **ag --failbackenable n_portnumber** command to enable failback.

```
switch:admin> ag --failbackenable 13
Failback policy is enabled for port 13
```

- Enter the **ag --failbackdisable n_portnumber** command to disable failback.

```
switch:admin> ag --failbackdisable 13
Failback policy is disabled for port 13
```

Enabling and disabling Failback for a port group

Use the following steps to enable or disable Failback policy on all the N_Ports belonging to the same port group.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the following commands to enable or disable Failback for a port group:

- Enter the **ag --failbackenable pg pgid** command to enable failback on a port group.

```
switch:admin> ag --failbackenable -pg 3  
Failback policy is enabled for port group 3
```

- Enter the **ag --failbackdisable pg pgid** command to disable failback on a port group.

```
switch:admin> ag --failbackdisable -pg 3  
Failback policy is disabled for port group 3
```

Upgrade and downgrade considerations for Failback

- Downgrading to Fabric OS v6.3.0 or earlier is supported.
- Upgrading from Fabric OS v6.3.0 is supported.

Trunking in Access Gateway mode

Brocade's hardware-based Port Trunking feature enhances management, performance, and reliability of Access Gateway N_Ports when they are connected to Brocade fabrics. Port Trunking combines multiple links between the switch and AG module to form a single, logical port. This enables fewer individual links, thereby simplifying management. This also improves system reliability by maintaining in-order delivery of data and avoiding I/O retries if one link within the trunk fails. Equally important is that framed-based trunking provides maximum utilization of links between the AG module and the core fabric.

Trunking allows transparent failover and failback within the trunk group. Trunked links are more efficient because of the trunking algorithm implemented in the switching ASICs that distributes the I/O more evenly across all the links in the trunk group.

Trunking in Access Gateway is mostly configured on the Edge switch. To enable this feature, you must install the Brocade ISL license on both the Edge switch and the module running in AG mode and ensure that both modules are running the same Fabric OS version. If a module already has an ISL Trunking license, no new license is required. After the trunking license is installed on a switch in AG mode and you change the switch to standard mode, you can keep the same license.

How Trunking works

Trunking in Access Gateway mode provides a trunk group between N_Ports on the AG module and F_Ports on the Edge switch module. With trunking, any link within a trunk group can go offline or become disabled, but the trunk remains fully functional and no re-configuration is required. Trunking prevents reassignments of the Port ID when N_Ports go offline.

Configuring Trunking on the Edge switch

Since AG Trunking configuration is mostly on the Edge switch, information in this section is applicable to the Edge switch module and not the AG module. On the AG module you only need to ensure that the trunking license is applied and enabled. On the Edge switch, you must first configure an F_Port Trunk group and statically assign an Area_ID to the trunk group. Assigning a Trunk Area (TA) to a port or trunk group enables F_Port masterless trunking on that port or trunk group. On switches running in Access Gateway mode, the masterless trunking feature trunks N_Ports because these are the only ports that connect to the Enterprise fabric. When a TA is assigned to a port or trunk group, the ports will immediately acquire the TA as the area of its process IDs (PID). When a TA is removed from a port or trunk group, the port reverts to the default area as its PID.

NOTE

By default, Trunking is enabled on all N_Ports of the AG; ensure that this feature is enabled on N_Ports that are part of port trunk group.

Trunk group creation

Port trunking is enabled between two separate Fabric OS switches that support trunking and where all the ports on each switch reside in the same quad and are running the same speed. Trunk groups form when you connect two or more cables on one Fabric OS switch to another Fabric OS switch with ports in the same port group or quad. A port group or a quad is a set of sequential ports, for example ports 0-3. The Brocade 300 switch supports a trunk group with up to eight ports. The trunking groups are based on the user port number, with contiguous eight ports as one group, such as 0-7, 8-15, 16-23 and up to the number of ports on the switch.

Setting up trunking

Trunking is enabled between two separate Fabric OS switches that support trunking and where all the ports on each switch reside in the same quad and are running the same speed. Trunk groups form when you connect two or more cables on one Fabric OS switch to another Fabric OS switch with ports in the same port group or quad. A port group or a quad is a set of sequential ports, for example ports 0-3 in the figure shown below. For example, the Brocade 300 platform supports a trunk group with up to eight ports. The trunking groups are based on the user port number, with contiguous eight ports as one group, such as 0-7, 8-15, 16-23 and up to the number of ports on the switch.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Ensure that both modules (Edge switch and the switch running in AG mode) have the trunking licenses enabled.
3. Ensure that the ports have trunking enabled by issuing the **portcfgshow** command. If Trunking is not enabled, issue the **portcfgtrunkport port 1** command.
4. Ensure that ports within a trunk have the same speed.
5. Ensure that ports within an ASIC trunk group are used to group the ports as part of a trunk on the Edge switch or on an AG.
6. Ensure that both modules are running the same Fabric OS versions.

Configuration management for trunk areas

The **porttrunkarea** command does not allow ports from different admin domains (ADs) and ports from different logical switches to join the same trunk area (TA) group.

When you assign a TA, the ports within the TA group will have the same Index. The Index that was assigned to the ports is no longer part of the switch. Any Domain,Index (D,I) AD that was assumed to be part of the domain may no longer exist for that domain because it was removed from the switch.

Example : How Trunk Area assignment affects the port Domain,Index

If you have AD1: 3,7; 3,8; 4,13; 4,14 and AD2: 3,9; 3,10, and then create a TA with Index 8 with ports that have index 7, 8, 9, and 10. Then index 7, 9, and 10 are no longer with domain 3. This means that AD2 does not have access to any ports because index 9 and 10 no longer exist on domain 3. This also means that AD1 no longer has 3,7 in effect because Index 7 no longer exists for domain 3. AD1's 3,8, which is the TA group, can still be seen by AD1 along with 4,13 and 4,14.

A port within a TA can be removed, but this adds the Index back to the switch. For example, the same AD1 and AD2 with TA 8 holds true. If you remove port 7 from the TA, it adds Index 7 back to the switch. That means AD1's 3,7 can be seen by AD1 along with 3,8; 4,13 and 4,14.

Assigning a Trunk Area

You must enable trunking on all ports to be included in a Trunk Area before you can create a Trunk Area. Use the **portCfgTrunkPort** or **switchCfgTrunk** command to enable trunking on a port or on all ports of a switch.

Issue the **porttrunkarea** command to assign a static TA on a port or port trunk group, to remove a TA from a port or group of ports in a trunk, and to display masterless trunking information.

You can remove specified ports from a TA using the **porttrunkarea --disable** command; however this command does not unassign a TA if its previously assigned Area_ID is the same address identifier (Area_ID) of the TA unless all the ports in the trunk group are specified to be unassigned. For more information on the **porttrunkarea** command, enter **help porttrunkarea** or see the *Fabric OS Command Reference*. F_Port trunking will not support shared area ports 16-47 on the Brocade FC8-48 blades.

The following table shows an example of the Address Identifier.

TABLE 7 Address identifier

23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Domain ID								Area_ID								Port ID							
Address Identifier																							

1. Connect to the switch and log in using an account assigned to the admin role.
2. Disable the ports to be included in the TA.
3. Enable TA for the appropriate ports. In the following example, TA is enabled for ports 13 and 14 on slot 10 with port index of 125.

```
switch:admin> porttrunkarea --enable 10/13-14 -index 125
Trunk index 125 enabled for ports 10/13 and 10/14
```

4. Show the TA port configuration (ports still disabled).

```
switch:admin> porttrunkarea --show enabled
```


Slot	Port	Type	State	Master	TI	DI
10	13	--	--	--	125	125
10	14	--	--	--	125	126

5. Enable ports specified in step 3. Continuing with the example shown in step 3, this would mean enabling ports 13 and 14.

```
switch:admin> portenable 10/13
switch:admin> portenable 10/14
```

6. Show the TA port configuration after enabling the ports:

```
switch:admin> porttrunkarea --show enabled
Slot  Port  Type   State  Master  TI  DI
-----
10    13    F-port Master  10/13   125 125
10    14    F-port Slave   10/13   125 126
```

Enabling the DCC policy on trunk

1. After you assign a Trunk Area, the **porttrunkarea** command checks whether there are any active DCC policies on the port with the index TA, and then issues a warning to add all the device WWNs to the existing DCC policy with index as TA.

All DCC policies that refer to an Index that no longer exist will not be in effect.

2. Add the WWN of all the devices to the DCC policy against the TA.
3. Enter the **secpolicyactivate** command to activate the DCC policy.

You must enable the TA before issuing the **secpolicyactivate** command in order for security to enforce the DCC policy on the trunk ports.

4. Turn on the trunk ports.

Trunk ports should be turned on after issuing the **secpolicyactivate** command to prevent the ports from becoming disabled in the case where there is a DCC security policy violation.

Enabling trunking

1. Connect to the switch and log in using an account assigned to the admin role.
2. Disable the desired ports by entering the **portdisable port** command for each port to be included in the TA.
3. Enter the **porttrunkarea --enable 3** command with appropriate operands to form a trunk group for the desired ports. For example, if ports 36-39 were disabled in step 2, then the example command shown below forms a trunk group for ports 36-39 with index 37. These will be connected to N_Ports on an AG module.

```
switch:admin> porttrunkarea --enable 36-39 -index 37
Trunk area 37 enabled for ports 36, 37, 38 and 39.
```

4. Enter the **portenable port** command for each port in the TA to re-enable the desired ports, such as ports 36-39.
5. Enter the **switchshow** command to display the switch or port information, including created trunks.

Disabling F_Port trunking

Use the following steps to disable F_Port Trunking.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **porttrunkarea --disable 36-39** command.

```
switch:admin> porttrunkarea --disable 36-39
ERROR: port 36 has to be disabled
```

Disable each port prior to removing ports from the TA. Then reissue the command:

```
switch:admin> porttrunkarea --disable 36-39
Trunk area 37 disabled for ports 36, 37, 38 and 39.
```

Trunking monitoring

For F_Port masterless trunking, you must install Filter, EE or TT monitors on the F_Port trunk port. Whenever the master port changes, it is required to move the monitor to the new master port. For example, if a master port goes down, a new master is selected from the remaining slave ports. APM must delete the monitor from the old master and install the monitor on new master port. If you attempt to add a monitor to a slave port, it is automatically added to the master port.

Trunking considerations for the Edge switch

Table 8 describes the Access Gateway trunking considerations for the Edge switch.

TABLE 8 Access Gateway trunking considerations for the Edge switch

Category	Description
Area assignment	You statically assign the area within the trunk group on the Edge switch. That group is the F_Port masterless trunk.
	The static trunk area you assign must fall within the F_Port trunk group starting from port 0 on a Edge switch or blade.
	The static trunk area you assign must be one of the port's default areas of the trunk group.
Authentication	Authentication occurs only on the F_Port trunk master port and only once per the entire trunk. This behavior is same as E_Port trunk master authentication. Because only one port in the trunk does FLOGI to the switch, and authentication follows FLOGI on that port, only that port displays the authentication details when you issue the portshow command. Note: Switches in Access Gateway mode do not perform authentication.
Management Server	Registered Node ID (RNID), Link Incident Record Registration (LIRR), and (QSA) Query Security Attributes ELs are not supported on F_Port trunks.

TABLE 8 Access Gateway trunking considerations for the Edge switch (Continued)

Category	Description
Trunk area	<p>The port must be disabled before assigning a Trunk Area on the Edge switch to the port or removing a Trunk Area from a trunk group.</p> <p>You cannot assign a Trunk Area to ports if the standby CP is running a firmware version earlier than Fabric OS V6.2.0.</p>
PWWN	The entire Trunk Area trunk group share the same Port WWN within the trunk group. The PWWN is the same across the F_Port trunk that will have 0x2f or 0x25 as the first byte of the PWWN. The TA is part of the PWWN in the format listed in Table 9 on page 57.
Downgrade	<p>You can have trunking on, but you must disable the trunk ports before performing a firmware downgrade.</p> <p>Note: Removing a Trunk Area on ports running traffic is disruptive. Use caution before assigning a Trunk Area if you need to downgrade to a firmware earlier than Fabric OS v6.1.0.</p>
Upgrade	No limitations on upgrade to Fabric OS v6.4.0 if the F_Port is present on the switch. Upgrading is not disruptive.
HA Sync	If you plug in a standby-CP with a firmware version earlier than Fabric OS v6.1.0 and a Trunk Area is present on the switch, the CP blades will become out of sync.
Port Types	Only F_Port trunk ports are allowed on a Trunk Area port. All other port types that include F/FL/E/EX are persistently disabled.
Default Area	Port X is a port that has its Default Area the same as its Trunk Area. The only time you can remove port X from the trunk group is if the entire trunk group has the Trunk Area disabled.
<code>portCfgTrunkPort port, 0</code>	<code>portCfgTrunkPort port, 0</code> will fail if a Trunk Area is enabled on a port. The port must be Trunk Area-disabled first.
<code>switchCfgTrunk 0</code>	<code>switchCfgTrunk 0</code> will fail if a port has TA enabled. All ports on a switch must be TA disabled first.
Port Swap	When you assign a Trunk Area to a trunk group, the Trunk Area cannot be port swapped; if a port is swapped, then you cannot assign a Trunk Area to that port.
Trunk Master	No more than one trunk master in a trunk group. The second trunk master will be persistently disabled with reason "Area has been acquired".
Fast Write	When you assign a Trunk Area to a trunk group, the trunk group cannot have fast write enabled on those ports; if a port is fast write enabled, the port cannot be assigned a Trunk Area.
FICON	FICON is not supported on F_Port trunk ports. However, FICON can still run on ports that are not F_Port trunked within the same switch.
FC8-48 blades	F_Port Trunking does not support shared area ports on the Brocade FC8-48 blades in a 48000. F_Port Trunking is supported on all ports on the Brocade FC8-48 in the DCX and DCX-4S.

TABLE 8 Access Gateway trunking considerations for the Edge switch (Continued)

Category	Description
FC4-32 blade	If an FC4-32 blade has the Trunk Area enabled on ports 16 - 31 and the blade is swapped with a FC8-48 blade, the Trunk Area ports will be persistently disabled. You can run the porttrunkarea command to assign a Trunk Area on those ports.
Trunking	You must first enable Trunking on the port before the port can have a Trunk Area assigned to it.
PID format	F_Port masterless trunking is only supported in CORE PID format.
Long Distance	Long distance is not allowed when AG is enabled on a switch. This means you cannot enable long distance on ports that have a Trunk Area assigned to them.
Port mirroring	Port mirroring is not supported on Trunk Area ports or on the PID of an F_Port trunk port.
Port speed	Ports within a trunk must have the same port speed for a trunk to successfully be created.
configdownload and configupload	<p>If you issue the configdownload command for a port configuration that is not compatible with F_Port trunking, and the port is Trunk Area enabled, then the port will be persistently disabled.</p> <p>Note: Configurations that are not compatible with F_Port trunking are long distance, port mirroring, non-CORE_PID, and Fastwrite.</p> <p>If you issue the configupload command, consider the following:</p> <ul style="list-style-type: none"> • A configuration file uploaded when AG mode is disabled cannot be downloaded when AG mode is enabled. • A configuration file uploaded when AG mode is enabled cannot be downloaded when AG mode is disabled. • A configuration file uploaded when the PG policy is enabled cannot be downloaded when the APC policy is enabled. • A configuration file uploaded when the APC policy is enabled cannot be downloaded when the PG policy is enabled.
ICL port	F_Port trunks are not allowed on ICL ports. The porttrunkarea command does not allow it.
AD	You cannot create a Trunk Area on ports with different Admin Domains. You cannot create a Trunk Area in AD255.
DCC Policy	DCC policy enforcement for the F_Port trunk is based on the Trunk Area; the FDISC requests to a trunk port is accepted only if the WWN of the attached device is part of the DCC policy against the TA. The PWWN of the FLOGI sent from the AG will be dynamic for the F_Port trunk master. Because you do not know ahead of time what PWWN AG will use, the PWWN of the FLOGI will not go through DCC policy check on an F_Port trunk master. However, the PWWN of the FDISC will continue to go through DCC policy check.

TABLE 8 Access Gateway trunking considerations for the Edge switch (Continued)

Category	Description
D.I. Zoning (D,I) AD (D, I) DCC and (PWWN, I) DCC	<p>Creating a Trunk Area may remove the Index ("I") from the switch to be grouped to the Trunk Area. All ports in a Trunk Area share the same "I". This means that Domain,Index (D,I), which refer to an "I", that might have been removed, will no longer be part of the switch.</p> <p>Note: Ensure to include AD, zoning and DCC when creating a Trunk Area.</p> <p>You can remove the port from the Trunk Area to have the "I" back into effect. D,I will behave as normal, but you may see the effects of grouping ports into a single "I".</p> <p>Also, D,I continues to work for Trunk Area groups. The "I" can be used in D,I if the "I" was the "I" for the Trunk Area group.</p> <p>Note: "I" refers to Index and D,I refers to Domain,Index.</p>
Two masters	Two masters is not supported in the same F_Port trunk group.
QoS	Supported.

The following table describes the PWWN format for F_Port and N_Port trunk ports.

TABLE 9 PWWN format for F_Port and N_Port trunk ports

NAA = 2	2f:xx:nn:nn:nn:nn:nn	Port WWNs for:	The valid range of xx is [0 - FF],
	(1)	switch's FX_Ports.	for maximum of 256.
NAA = 2	25:xx:nn:nn:nn:nn:nn	Port WWNs for:	The valid range of xx is [0 - FF],
	(1)	switch's FX_Ports	for maximum of 256.

Trunking considerations for Access Gateway module

Consider the following for Trunking in Access Gateway mode:

- Access Gateway trunking is not supported on M-EOS or third-party switches.
- Trunk groups cannot span across multiple N_Port groups within an AG module in AG mode. Multiple trunk groups are allowed within the same N_Port group. All ports within a trunk group must be part of the same port group; ports outside of a port group cannot form a trunk group
- The **ag -wwnmapshow** command will not display trunking for device-mapped ports. If a device is mapped to a port with device mapping and that port is currently part of a trunk, then the device will use that trunk. When trunking is used with Device Load Balancing Policy, then the load on each trunk will be proportional to the number of ports in that trunk. Use the **ag -show** command to determine the devices using a particular trunk.

Upgrade and downgrade considerations for Trunking in Access Gateway mode

Upgrading and downgrading from Fabric OS v6.4.0 to Fabric OS v6.3.0 and earlier is supported.

Adaptive Networking on Access Gateway

Adaptive Networking (AN) ensures bandwidth for critical servers, virtual servers, or applications in addition to reducing latency and minimizing congestion. Adaptive Networking in Access Gateway works in conjunction with the Quality of Service (QoS) feature on Brocade fabrics. Fabric OS provides a mechanism to assign traffic priority, (high, medium, or low) for a given source and destination traffic flow. By default, all flows are marked as medium.

The following must be appropriately installed:

- The Adaptive Networking (AN) license must be installed on all switches operating in Access Gateway mode to take advantage of the QoS and Ingress Rate Limiting features.
- The Server Application Optimization (SAO) license must be installed to extend QoS features to supported HBAs.

To determine if these licenses are installed on the connected switch, issue the Fabric OS **licenseshow** command. Refer to the *Fabric OS Administrator's Guide* for detailed information about QoS.

You can configure the ingress rate limiting and SID/DID traffic prioritization levels of QoS for the following configurations:

- Supported HBA to AG to switch
- Unsupported HBA to AG to switch
- HBA (all) to Edge AG to Core AG to switch

For additional information on the Brocades adapters, refer to your *HBA Administrator's Guide*.

QoS: Ingress Rate Limiting on AG

Ingress rate limiting restricts the speed of traffic from a particular device to the switch port. On switches in AG mode, you must configure ingress rate limiting on F_Ports.

For more information and procedures for configuring this feature, refer to “QoS: Ingress Limiting” in the *Fabric OS Administrator's Guide*.

QoS: SID/DID traffic prioritization

SID/DID traffic prioritization allows you to categorize the traffic flow between a given host and target as having a high or low priority; the default is medium. For example, you can assign online transaction processing (OLTP) to a high priority and the backup traffic to a low priority.

For detailed information on this feature, refer to “QoS: SID/DID traffic prioritization” in the *Fabric OS Administrator's Guide*.

[Figure 12](#) on page 59 shows the starting point for QoS in various Brocade and Non-Brocade configurations.




FIGURE 12 Starting point for QoS

Upgrade and downgrade considerations with Adaptive Networking in AG mode enabled

Downgrading to Fabric OS v6.3.0 is supported. Note the following considerations when upgrading and downgrading from Fabric OS v6.4.0 to Fabric OS v6.2.X and earlier:

- If any of the AG QoS enabled ports are active and you attempt a firmware downgrade, the downgrade is prevented. You must disable the QoS-enabled ports before performing a firmware downgrade.
- Upgrades from earlier versions to Fabric OS v6.4.0 are allowed, but AG QoS-enabled ports do not become effective until the ports are disabled or enabled so that QoS mode can be negotiated on the ISL links.

Adaptive Networking on Access Gateway considerations

- QoS is configured in the fabric, as normal, and not on the AG module. To extend QoS benefits to AG and devices behind it you only need to ensure that the AN and/or SAO licenses are applied and enabled on the AG module.
- QoS on Access Gateway is only supported on Fabric OS 6.3 and later.
- You should disable HBA QoS if connected to a 6.2 version AG.
- Disable QoS on an AG port if it connects with a switch running Fabric OS 6.2. Otherwise, the port will automatically disable with an error. To recover, disable QoS on the port, then enable the port.
- Disabling QoS on online N_Ports in the same trunk can cause the slave N_Port ID virtualization (NPIV) F_Port on the edge switch to become persistently disabled with "Area has been acquired." This is expected behavior because after QOS is disabled, the slave NPIV F_Port on the edge switch also tries to come up as a master. To avoid this issue, simply persistently enable the slave F_Port on the switch.
- QoS takes precedence over ingress rate limiting
- Ingress rate limiting is not enforced on trunked ports.

Per Port NPIV login limit

This feature allows you to set a specific maximum NPIV login limit on individual ports. This feature works in both Native Fabric Switch and Access Gateway mode. Using this feature, you can use additional tools to design and implement a virtual infrastructure. In Access Gateway mode, this feature allows smaller login limits for F_Ports and larger limits for N_Ports. Note that N_Ports are restricted by the NPIV login limit of the connecting port on the Edge switch.

Note the following aspects of this feature:

- Upgrading from Fabric OS v6.3.0 to v6.4.0 will retain the NPIV login limit set in v6.3.0
- Downgrading from Fabric OS v6.4.0 to v6.3.0 will reset the NPIV login limit back to 255.
- The value that you set is persistent across reboots and firmware upgrades.
- This feature supports virtual switches, so each port can have a specific NPIV login limit value in each logical switch.
- The login limit default is 126. This value will be set for a port when the **portCfgDefault** command is used to reset port default values.
- Before changing the login limits, you must disable the port.
- This feature only applies to ports enabled for NPIV operation. To enable NPIV functionality for a port, you can use the **portCfgNPIVPort –enable** command when the switch is in Fabric OS Native mode. For details, refer to the *Fabric OS Command Reference Manual*.

Setting the login limit

Use the following procedure to set the NPIV login limit for a port.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Disable the port by entering the **portdisable port** command.
3. Enter the **portcfgnpiv –setloginlimit [Slot/]Port loginlimit** command to set the login limit. For example, the following sets the login limit on port 12 to 200.

```
portcfgnpivport --setloginlimit 12 200
```

Considerations for the Brocade 8000

This section provides information on differences in operation, Fabric OS command function, and features on the Brocade 8000 when operating in Access Gateway mode.

Port mapping

The Brocade 8000 contains FCoE and Fibre Channel ports. In Access Gateway mode, the FCoE ports are configured logically as F_Ports, while the Fibre Channel ports are configured as N_Ports. For details on how this affects port-based mapping, refer to [“Brocade 8000 mapping differences”](#) on page 12.

Policy and feature support

The following AG policies and features are not supported on the Brocade 8000.

- Access Gateway Cascading

NOTE

This is not supported on the Brocade 8000 Core AG (the Brocade 8000 is only supported on an Edge AG).

- Automatic Load Balancing
- Auto Port Configuration Policy
- Persistent ALPA
- Device Load Balancing

Fabric OS commands

This section describes differences in using Fabric OS commands on the Brocade 8000 in AG mode.

- The following commands are not supported on the Brocade 8000 in AG mode:
 - `ag -pgmapadd`
 - `ag -pgmapdel`
 - `ag -pgsetmodes`
 - `ag -pgdelmodes`
 - `ag -pgfnmtov`
 - `ag -persistentalpaenable`
 - `ag -printalpamap`
 - `ag -deletepwwnfromdb`
 - `ag -clearalpamap`
 - `ag -wwnmapshow`
 - `ag -addwwnmapping`
 - `ag -delwwnmapping`
 - `ag -addwwnpgmapping`
 - `ag -delwwnpgmapping`
 - `ag -wwnmappingenable`
 - `ag -wwnmappingdisable`
 - `ag -delwwnfailovermapping`
 - `agautomapbalance`
 - `portcfgnport`

3 Considerations for the Brocade 8000

- The following commands have restricted usage, mostly because the Brocade 8000 contains only eight Fibre Channel ports and does not support the Automatic Port Configuration policy:
 - **ag –pgcreate**
 - **ag –policyenable**
 - **ag –policydisable**
 - **ag –portcfgdefault**
- To enable or disable FCoE (F) ports, use **fcoe –enable** and **fcoe –disable** instead of **portdisable** and **portenable**.
- The **portcfgdefault** command resets the degraded state and NPIV PerPort and clears the BufferLimitedMode on a port. For other AG platforms, this command restores the port configuration to factory default values.

Port Trunking and QoS features

Because the Brocade 8000 has limited available buffers and Port trunking and QoS require more buffers than normal, consider the following points:

- Do not enable QoS by itself on more than six Fibre Channel ports at a time. If you attempt to enable on more than six ports, the Brocade 8000 may enter buffer-limited mode.
- To enable both Trunking and QoS on the Brocade 8000, we recommend that you enable QoS first. If you enable Trunking first, both features will compete for buffers and you will not be able to enable QoS on more than two ports. If you enable QoS first, adequate buffers will be available for Trunking due to the function of QoS.

Automatic Login Balancing

MFNM is enabled by default on all port groups and cannot be disabled on the Brocade 8000. Because of this, the **pgsetmodes**, **pgdelmodes**, and **pgcreate** commands are blocked for the **-m** option, and Automatic Login Balancing cannot be enabled.

SAN Configuration with Access Gateway

In this chapter

- Connectivity of multiple devices overview 63
- Direct target attachment. 63
- Target aggregation. 64
- Access Gateway cascading. 64
- Fabric and Edge switch configuration 65
- Connectivity to Cisco Fabrics 67
- Rejoining Fabric OS switches to a fabric 67

Connectivity of multiple devices overview

This chapter describes how to connect multiple devices to a switch in Access Gateway (AG) mode, and discusses Edge switch compatibility, target aggregation, direct target attachment, port requirements, NPIV HBA, and interoperability. AG does not support daisy chaining when two AG devices are connected to each other in a loop configuration. Switches in AG mode can connect to third-party fabrics with the following firmware versions:

- M-EOSc v9.6.2 or later and M-EOSn v9.6 or later.
- Cisco MDS Switches with SAN OS v3.0(1).
- Loop devices and FICON channels/control unit connectivity are not supported.
- When a switch is in AG mode, it can be connected to NPIV-enabled HBAs, or F_Ports that are NPIV-aware. Access Gateway supports NPIV industry standards per FC-LS-2 v1.4.

Direct target attachment

FCP targets can directly connect to an AG module instead of through a fabric connection.

Even though target devices can directly be connected to AG ports we recommend that target devices be connected to the core Fabric. Follow the “[Considerations](#)” below when connecting target devices directly to an AG.

Considerations

- Direct Target attachment to AG is only supported if the AG module is also connected to a core fabric. A switch module running in AG mode does not provide Name Services on its own, and routing to the target devices will need to be established by the core fabric.
- Hosts and targets can not be mapped to the same N-port.

- Redundant configurations should be maintained so that when hosts and targets fail over or fail back, they should not get mapped to a single N_Port.
- Hosts and targets should be in separate port groups.
- Configuration is not enforced.

Target aggregation

Access Gateway mode is normally used as host aggregation. In other words, a switch module in AG mode aggregates traffic from a number of host systems onto a single uplink N_Port. Similarly, many targets can be aggregated onto to a single uplink N_port. This feature has many applications. As one example, you can consolidate targets with various lower Fibre Channel speeds (such as 1, 2 or 4 Gbps) onto a single high-speed uplink port to the core fabric. This reduces the number of core fabric ports used by target devices and allows higher scalability.

Access Gateway cascading

Cascading is an advanced configuration supported in Access Gateway mode. You can use cascading to further increase the ratio of hosts to fabric ports beyond what a single switch model in AG mode can support.

Access Gateway cascading lets you connect two Access Gateway (AG) switches linking them back to back. The AG switch that is directly connected to the fabric is referred to as the Core AG. In this document, the AG switch connected to the device is referred to as the Edge AG. [Figure 13](#) on page 64 illustrates Access Gateway cascading.




FIGURE 13 Access Gateway cascading

AG cascading provides higher over-subscription because it allows you to consolidate the number of ports going to the main fabric. There is no license requirement to use this feature.

Note the following configuration considerations when cascading Access Gateways:

- Only one level of cascading is supported. Note that several Edge AGs can connect into a single Core AG to support even a higher consolidation ratio.

- AG trunking between the Edge and Core AG switches is not supported. Trunking between the Core AG switch and the fabric is supported.
- It is recommended that you enable Advanced Security Policy (ADS) on all AG F_Ports that are directly connected to devices.
- APC policy is not supported when cascading.
- Loopbacks (Core AG N_Port to Edge AG F_Port) are not allowed.
- The **agshow** command issued on the fabric will discover only the Core AG switches. If issued as **agshow --name AG name**, then the F_Ports of both the Core and Edge AG switches will be shown for the Core AG switch.
- Due to high subscription ratios that could occur when cascading AGs, ensure there is enough bandwidth for all servers when creating such configurations. The subscription ratio becomes more acute in a virtual environment.

Fabric and Edge switch configuration

To connect devices to the fabric using Access Gateway, configure the fabric and Edge switches within the fabric that will connect to the AG module using the following parameters. These parameters apply to Fabric OS, M-EOS, and Cisco-based fabrics:

- Install and configure the switch as described in the switch's Hardware Reference manual before performing these procedures.
- Verify that the interop mode parameter is set to Brocade Native mode.
- Configure the F_Ports on the Edge switch to which Access Gateway is connected as follows:
 - Enable NPIV.
 - Disable long distance mode.
 - Allow multiple logins for M-EOS switches. The recommended fabric login setting is the maximum allowed per port and per switch.
- Use only WWN zoning for devices behind AG.
- If DCC security is being used on Edge switches that directly connect to AG, make sure to include the Access Gateway WWN or the port WWN of the N_Ports. Also include the HBA WWNs that will be connected to AG F_Ports to the ACL list in the ACL policy. It is recommended to use AG ADS policy instead of the DCC policy on the Edge switch.
- Allow inband queries for forwarded fabric management requests from the hosts. Add the Access Gateway switch WWN to the access list if inband queries are restricted.

Before connecting Access Gateway to classic Brocade switches, disable the Fabric OS Management Server Platform Service to get accurate statistical and configuration fabric data,

Verifying the switch mode

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **switchShow** command to display the current switch configuration.

The following example shows partial output for this command for a switch in the Fabric OS Native mode where **switchMode** displays as Native.

```
switch:admin> switchshow
switchName:      switch
```

4 Fabric and Edge switch configuration

```
switchType:      76.6
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:     13
switchId:        fffc01
switchWwn:       10:00:00:05:1e:03:4b:e7
zoning:          OFF
switchBeacon:    OFF
-----=
```

See [Table 3](#) on page 9 for a description of the port state.

If the switch is in Native mode, you can enable AG mode; otherwise, set the switch to Native mode, and then reboot the switch.

Enabling NPIV on M-EOS switches

1. Connect to the switch and log in as admin on the M-EOS switch.
2. Enable the MS services by entering the following commands:

For the Mi10K switch, enter the following command.

```
fc osmsState <vfid> <state>
```

where

vfid Virtual fabric identification number.

state Can be *enable* for the enabled state or *disable* for the disabled state.

For other McDATA switches, enter the following command.

```
config OpenSysMs setState <osmsState>
```

where

osmsState Can be *enable* or *1* for the enabled state or *disable* or *0* for the disabled state.

3. Enable NPIV functionality on the Edge fabric ports so that multiple logins are allowed for each port. Enter the following command on the M-EOS switch to enable NPIV on the specified ports.

```
config NPIV
```

Your M-EOS switch is now ready to connect.

NOTE

You can run the **agshow** command to display Access Gateway information registered with the fabric. When an Access Gateway is exclusively connected to Non-Fabric-OS-based switches, it will not show up in the **agshow** output on other Brocade Switches in the fabric.

Connectivity to Cisco Fabrics

When connecting a switch in Access Gateway mode to a Cisco fabric Fabrics you only need to make sure NPIV is enabled on the connecting switch and that Fabric OS version 3.1 or higher is used.

Enabling NPIV on a Cisco switch

1. Log in as admin on the Cisco MDS switch.
2. Enter the **show version** command to determine that you are using the correct SAN-OS version and to see if NPIV is enabled on the switch.
3. Enter the following commands to enable NPIV:

```
conf t
enable npiv
```

4. Press **Ctrl-Z** to exit.
5. Enter the following commands to save the MDS switch connection:

```
copy run start
```

Your Cisco switch is now ready to connect to a switch in Access Gateway mode.

Rejoining Fabric OS switches to a fabric

When a switch reboots after AG mode is disabled, the Default zone is set to no access. Therefore, the switch does not immediately join the fabric to which it is connected. Use one of the following methods to re-join a switch to the fabric:

- If you saved a Fabric OS configuration before enabling AG mode, download the configuration using the **configDownload** command.
 - If you want to re-join the switch to the fabric using the fabric configuration, use the following procedure.
1. Connect to the switch and log in using an account assigned to the admin role.
 2. Enter the **switchDisable** command to disable the switch.
 3. Enter the **defZone --allAccess** command to allow the switch to merge with the fabric.
 4. Enter the **cfgSave** command to commit the defzone changes.
 5. Enter the **switchEnable** command to enable the switch and allow it to merge with the fabric.

The switch automatically re-joins the fabric.

Reverting to a previous configuration

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **switchDisable** command to disable the switch.
3. Enter the **configDownload** command to revert to the previous configuration.
4. Enter the **switchEnable** command to bring the switch back online.

4 Rejoining Fabric OS switches to a fabric

The switch automatically joins the fabric.

Troubleshooting

This appendix provides troubleshooting instructions.

TABLE 10 Troubleshooting

Problem	Cause	Solution
Switch is not in Access Gateway mode	Switch is in Native switch mode	<p>Disable switch using the switchDisable command.</p> <p>Enable Access Gateway mode using the ag --modeenable command.</p> <p>Answer yes when prompted; the switch reboots.</p> <p>Log in to the switch.</p> <p>Display the switch settings using the switchShow command. Verify that the field switchMode displays Access Gateway mode.</p>
NPIV disabled on Edge switch ports	Inadvertently turned off	<p>On the Edge switch, enter the portCfgShow command.</p> <p>Verify that NPIV status for the port to which Access Gateway is connected is ON.</p> <p>If the status displays as “–” NPIV is disabled. Enter the portCfgNpivPort port_number command with the enable operand to enable NPIV.</p> <p>Repeat this step for each port as required.</p>
Need to reconfigure N_Port and F_Ports	Default port setting not adequate for customer environment	<p>Enter the portCfgShow command.</p> <p>For each port that is to be activated as an N_Port, enter the portCfgNport port_number command with the 1 operand.</p> <p>All other ports remain as F_Ports.</p> <p>To reset the port to an F_Port, enter the portCfgNpivPort port_number command with the disable operand.</p>
LUNs are not visible	<p>Zoning on fabric switch is incorrect.</p> <p>Port mapping on Access Gateway mode switch is incorrect.</p> <p>Cabling not properly connected.</p>	<p>Verify zoning on the Edge switch.</p> <p>Verify that F_Ports are mapped to an online N_Port. See “Access Gateway default port mapping” on page 12.</p> <p>Perform a visual inspection of the cabling, check for issues such as wrong ports, twisted cable, or bent cable. Replace the cable and try again. Ensure the F_Port on AG module is enabled and active.</p>

TABLE 10 Troubleshooting (Continued)

Problem	Cause	Solution
Failover is not working	Failover disabled on N_Port.	<p>Verify that the failover and failback policies are enabled, as follows:</p> <p>Enter the ag --failoverShow command with the <i>port_number</i> operand.</p> <p>Enter the ag --failbackShow command with the <i>port_number</i> operand.</p> <p>Command returns “Failback (or Failover) on N_Port <i>port_number</i> is supported.”</p> <p>If it returns, “Failback (or Failover) on N_Port <i>port_number</i> is not supported.” See “F_Ports automatically fail over to any available N_Port. Alternatively, you can specify a preferred secondary N_Port in case the primary N_Port fails. If the primary N_Port goes offline, the F_Ports fail over to the preferred secondary N_Port (if it is online), then re-enable. If the secondary N_Port is offline, the F_Ports will disable. Define the preferred secondary N_Ports per F_Port. For example, if two F_Ports are mapped to a primary N_Port 1, you can define a secondary N_Port for one of those F_Ports and not define a secondary N_Port for the other F_Port. F_Ports must have a primary N_Port mapped before a secondary N_Port can be configured.” on page 45.</p>
Access Gateway is mode not wanted	Access Gateway must be disabled.	<p>Disable switch using the switchDisable command.</p> <p>Disable Access Gateway mode using the ag --modeDisable command.</p> <p>Answer yes when prompted; the switch reboots.</p> <p>Log in to the switch.</p> <p>Display the switch settings using the switchShow command. Verify that the field switchMode displays Fabric OS Native mode.</p>
“Login Rejected by FC stack” messages on console may be seen during F_Port and N_Port disruptions on Brocade 8000 in AG Mode.	The CNA host is retrying a log in before the switch has finished precessing a previous fabric logout (LOGO) attempt.	Working as designed. After the switch has completed LOGO processing, it will accept another login.

NOTE

If a Fabric OS switch is in AG mode and is also set to McDATA Fabric mode, when that switch is connected to an M-EOS switch, the Fabric OS switch does not display in the output when you run the **agshow** command.

Index

A

Access Gateway

- cascading, 64
- comparison to standard switches, 4
- compatible fabrics, 1
- connecting devices, 63
- connecting two AGs, 64
- description, 1
- displaying information, 66
- features, 3
- limitations, 5
- mapping description, 11
- port types, 4

Access Gateway mode

- comparison, 2
- disabling, 9
- port types, 4
- supported firmware versions, 63
- terms, *xvi*
- verifying, 7

ACL policies, settings, 65

adding devices to fabric, 30

address Identifier, 52

admin domain, 56

ADS Policy

- adding devices, 30
- displaying devices, 30, 31
- enabling, 29
- removing devices, 30

APC Policy

- disabling, 33
- rebalancing F_Ports, 37
- support for port groups, 36

area assignment, 54

authentication, limitations, 54

B

behavior, failover policy, 49

Brocade 8000

- AG considerations, 60
- default mapping, 12
- mapping differences, 12

C

Cisco fabric

- connectivity, 67
- enabling NPIV on Cisco switch, 67

code, *xv*

commands

- ag --addwwnfailovermapping, 46
- ag --addwwnpgmapping, 19
- ag --delwwnfailovermapping, 47
- ag --delwwnpgmapping, 19
- ag --failbackEnable, 49, 50
- ag --failbackShow, 49, 70
- ag --failoverDisable, 47
- ag --failoverEnable, 47, 48
- ag --failoverShow, 47, 70
- ag --mapAdd, 14
- ag --mapDel, 15
- ag --mapShow, 8, 14
- ag --modeDisable, 9, 70
- ag --modeEnable, 7, 69
- ag --modeShow, 7
- ag --policydisable wwnloadbalance, 41
- ag --policyenable wwnloadbalance, 40
- ag --wwnmapping, 19, 20, 46, 47
- ag --wwnmappingdisable, 20
- ag --wwnmappingenable, 21
- ag --wwnmapshow, 19, 20
- cfgSave, 67
- configDownload, 67
- configUpload, 18
- defZone --allAccess, 67
- portCfgNpivPort, 69
- portCfgNport, 25, 69
- portCfgShow, 69
- switchDisable, 9, 67, 69, 70
- switchEnable, 67
- switchMode, 69, 70
- switchShow, 8, 15, 65, 69, 70

compatibility, fabric, 65

configurations

- enabling switch, 67
- limitations with configdownload command, 56
- merging switch with fabric, 67
- re-joining switch to fabric, 67
- saving, 67
- using configdownload command, 67

D

daisy chaining, 63

DCC policy

- adding WWN, 53
- enabling, 53
- limitation creating TA, 56

default area, removing ports, 55

device load balancing, 36

device load balancing policy, 40

- APC policy, 41
- Brocade 8000, 41
- considerations, 41
- disabling, 40
- enabling, 40
- trunking, 41, 57

device mapping, 10

- adding a secondary N_Port, 46
- adding devices to N_Ports, 20
- considerations, 23
- disabling, 20
- display mapping information, 21
- enabling, 21
- failover, 46
- feature overview, 15
- pre-provisioning, 22
- removing secondary N_Port, 47
- static vs. dynamic mapping, 18
- to port group, 18
- to ports, 20
- VMware configuration, 22
- VMware considerations, 22

devices

- attaching multiple devices, 63

disabling switch

- switchDisable, 67

domain,Index, 52

downgrading, 55

downgrading considerations, 31, 33

dynamic vs. static mapping, 18

E

Edge switch

- FLOGI, 65
- long distance mode setting, 65
- NPIV, 65
- settings, 65

F

F_Port

- adding external port on embedded switch, 24
- description, 4
- mapping, example, 11
- maximum number mapped to N_Port, 24
- settings, Edge switch, 65
- shared area ports, 52
- trunking setup, 51

fabric

- compatibility, 65
- inband queries, 65
- join, 67
- logins, 65
- management server platform, 65
- zoning scheme, 65

Fabric OS management server platform service settings, 65

failback policy

- upgrade and downgrade considerations, 50

failback policy example, 44, 48

failover

- device mapping, 46

failover policy

- behavior, 45
- configurations for port mapping, 44
- enabling, 47
- example, 45, 49
- port mapping, 44

fast write limitation, 55

FICON, F_Port trunk ports, 55

H

HA sync, TA present, 55

I

ICL ports, limitations, 56

inband queries, 65

initiator and target port considerations, 11

J

join fabric, 67

L

limitations

- device load balancing, 41
- direct connections to target devices, 5
- loop devices not supported, 5
- login balancing considerations, 38
- long distance mode, Edge switch, 65

M

managed fabric name monitoring

- disabling, 38
- displaying current timeout value, 38
- enabling, 38
- setting timeout values, 39

management server, 54

mapping

- Brocade 8000 differences, 12
- considerations, 22
- default mapping for Brocade 8000, 12
- device, 10
- device to port groups, 18
- devices to ports, 20
- example, 11
- port, 10
- ports, 10

mapping priority, 10

masterless trunking, 56

M-EOS switch, enabling NPV, 66

N

N_Port

- configurations, 24
- description, 4
- displaying configurations, 25
- failover in a PG, 39
- mapping example, 11
- masterless trunking, 51
- maximum number supported, 24
- multiple trunk groups, 57
- trunk groups, 57
- unlock, 25
- unlocking, 25

N_Port configurations

- displaying, 25

N_Ports

- unlocking, 25
- native switchMode, 66
- non disruptive, 55
- NPIV
 - Edge switch, 65
 - enabling on Cisco switch, 67
 - enabling on M-EOS switch, 66
 - login limit, 60
 - support, 63

O

- optional features, *xviii*

P

- per port NPIV login limit, 60
- Persistent ALPA
 - support, 41
- persistent ALPA
 - clearing ALPA values, 43
 - considerations, 43
 - deleting hash table data, 42
 - disabling, 42
 - enabling, 42
 - flexible ALPA value, 42
 - reboot, 43
 - stringent ALPA value, 42
 - tables, 42
 - upgrade and downgrade considerations, 43
 - value types, 42
- policies
 - advance device security, 28
 - enabling DCC policy, 53
 - enforcement matrix, 28
 - port grouping, 33
 - showing current policies, 27
 - using policyshow command, 27
- port
 - comparison, 4
 - mapping, 10
 - requirements, 63
 - types, 4

port group

- add N_Port, 34, 36
- create, 36
- delete N_Port, 35
- disabling, 35
- enabling logging balancing mode, 36
- login balancing mode, 36
- managed fabric name monitoring mode, 36
- remove port group, 35
- rename, 35
- Port Grouping policy
 - using portcfgnport command, 25
- port grouping policy
 - considerations, 39
 - downgrading considerations, 40
- port mapping, 10
 - adding F_Ports to N_Ports, 14
 - adding ports, 14
 - adding secondary N_Port, 45
 - considerations for initiator and target ports, 11
 - default F_Port-to-N_Port, 12
 - deleting secondary N_Port, 46
 - maximum number of F_Ports, 24
 - removing F_Ports from N_Ports, 15
 - removing F_Ports fromn N_Ports, 15
- Port mirroring, not supported, 56
- port state, description, 9
- port swap, not swapping TA, 55
- port types, limitations, 55
- preferred secondary N_Port
 - login balancing mode, 45
 - online, 44
- PWWN
 - format, 57
 - sharing TA trunk group, 55

Q

QoS

- firmware downgrade, 59
- ingress rate limiting, 58
- SID/DID traffic prioritization, 58

R

- removing devices from switch, 30
- removing trunk ports, 55
- requirements, ports, 63

S

settings

- ACL policies, 65
 - FLOGI, 65
 - inband queries, 65
 - management server platform, 65
 - zone, no access, 67
- static vs. dynamic mapping, 18
- supported hardware and software, *xiii*
- switch mode, verify, 65

T

terms, *xvi*

trunk area

- assign, 52
- configuration management, 52
- disabling, 55
- remove ports, 52
- standby CP, 55
- using the porttrunkarea command, 56

trunk groups, create, 51

trunk master, limitation, 55

trunking, 50

- configuring on edge switch, 51
- considerations in AG module, 57
- considerations on edge switch, 54
- disabling, 54
- enabling, 53, 56
- license, 50
- monitoring, 54

U

unlock N_Port, 25

upgrading, 55

V

VMware configuration for device mapping, 22

Z

zoning

- schemes, 65
- setting, 67

