

Brocade Fabric OS Access Gateway Administration Guide, 8.2.1

Supporting Fabric OS 8.2.1

Copyright © 2018 Brocade Communications Systems LLC. All Rights Reserved. Brocade and the stylized B logo are among the trademarks of Brocade Communications Systems LLC. Broadcom, the pulse logo, and Connecting everything are among the trademarks of Broadcom. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

Brocade, a Broadcom Inc. Company, reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Brocade is believed to be accurate and reliable. However, Brocade does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

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 <https://www.broadcom.com/support/fibre-channel-networking/tools/oscd>.

Contents

Introduction.....	6
About This Guide.....	6
What's New in This Document.....	6
Supported Hardware and Software.....	6
Brocade Gen 5 Platform (16Gb/s) Fixed-Port Switches.....	6
Brocade Gen 6 Platform (32Gb/s) Fixed-Port Switches.....	6
Key Access Gateway Terms.....	7
Contacting Brocade Technical Support.....	7
Document Feedback.....	8
Access Gateway Basic Concepts.....	9
Brocade Access Gateway Overview	9
Comparing Native Fabric and Access Gateway Modes.....	10
Fabric OS Features in Access Gateway Mode.....	11
Buffer Credit Recovery Support.....	14
Forward Error Correction Support.....	14
Virtual Fabrics Support.....	14
Device Authentication Support.....	14
AG Mode without All Ports on Demand Licenses.....	16
Password Distribution Support.....	17
FDMI Support.....	17
NTP Configuration Distribution to Access Gateways.....	17
Remote Fosexec Support	18
Slow-Drain Device Quarantining Support.....	19
Access Gateway Port Types.....	21
Comparison of Access Gateway Ports to Standard Switch Ports	21
Access Gateway Hardware Considerations.....	23
Configuring Ports in Access Gateway Mode.....	24
Enabling and Disabling Access Gateway Mode.....	24
Port State Description.....	25
Access Gateway Mapping.....	25
Port Mapping.....	26
F_Port Static Mapping.....	32
Device Mapping.....	33
Considerations for Access Gateway Mapping.....	39
N_Port Configurations.....	40
Displaying N_Port Configurations.....	42
Unlocking N_Ports	42
Persistent Port Online State	42
D_Port Support.....	43
Saving Port Mappings.....	44
Limitations and Considerations.....	44
Managing Policies and Features in Access Gateway Mode.....	45
Access Gateway Policies Overview.....	45
Displaying Current Policies	45
Access Gateway Policy Enforcement Matrix	45

Advanced Device Security Policy	46
How the ADS Policy Works.....	46
Enabling and Disabling the ADS Policy.....	46
Allow Lists.....	47
ADS Policy Considerations.....	48
Upgrade and Downgrade Considerations for the ADS Policy.....	49
Automatic Port Configuration Policy	49
How the APC Policy Works.....	49
Enabling and Disabling the APC Policy.....	49
APC Policy Considerations.....	50
Upgrade and Downgrade Considerations for the APC Policy.....	50
Port Grouping Policy.....	50
How Port Groups Work.....	50
Adding an N_Port to a Port Group.....	52
Deleting an N_Port from a Port Group.....	52
Removing a Port Group	52
Renaming a Port Group.....	52
Disabling the Port Grouping Policy.....	52
Port Grouping Policy Modes.....	53
Creating a Port Group and Enabling Automatic Login Balancing Mode.....	53
Enabling MFNM Mode.....	54
Disabling MFNM Mode.....	55
Displaying the Current MFNM Mode Timeout Value.....	55
Setting the Current MFNM Mode Timeout Value.....	55
Port Grouping Policy Considerations.....	55
Upgrade and Downgrade Considerations for the Port Grouping Policy.....	56
Device Load Balancing Policy.....	56
Enabling the Device Load Balancing Policy.....	56
Disabling the Device Load Balancing Policy.....	57
Device Load Balancing Policy Considerations.....	57
Persistent ALPA Policy.....	57
Enabling the Persistent ALPA Policy.....	58
Disabling the Persistent ALPA Policy.....	58
Persistent ALPA Device Data.....	58
Clearing ALPA Values.....	59
Persistent ALPA Policy Considerations.....	59
Failover Policy.....	59
Failover with Port Mapping.....	59
Failover with Device Mapping	62
Enabling and Disabling the Failover Policy on an N_Port.....	63
Enabling and Disabling the Failover Policy for a Port Group.....	64
Upgrade and Downgrade Considerations for the Failover Policy.....	64
Failback Policy.....	64
Failback Policy Configurations in Access Gateway.....	64
Enabling and Disabling the Failback Policy on an N_Port	65
Enabling and Disabling the Failback Policy for a Port Group.....	66
Forcing Failback on N_Ports.....	66
Upgrade and Downgrade Considerations for the Failback Policy.....	67
Failback Policy Disabled on Unreliable Links (N_Port Monitoring).....	67
Trunking in Access Gateway Mode.....	68

How Trunking Works.....	68
Configuring Trunking on the Edge Switch.....	68
Configuration Management for Trunk Area.....	69
Enabling Trunking	71
Disabling F_Port Trunking.....	71
Monitoring Trunking	71
AG Trunking Considerations for the Edge Switch.....	72
Trunking Considerations for Access Gateway Mode.....	74
Upgrade and Downgrade Considerations for Trunking in Access Gateway Mode.....	74
Adaptive Networking on Access Gateway.....	75
QoS: Ingress Rate Limiting	75
QoS: SID/DID Traffic Prioritization	75
Upgrade and Downgrade Considerations for Adaptive Networking in AG Mode.....	76
Adaptive Networking on Access Gateway Considerations.....	76
Per-Port NPIV Login Limit.....	77
Setting the Login Limit.....	77
Duplicate PWWN Handling during Device Login.....	78
Performance Monitoring.....	78
Flow Monitor.....	78
Flow Mirror.....	79
Considerations for the Brocade Switches.....	81
SAN Configuration with Access Gateway.....	82
Connectivity of Multiple Devices Overview.....	82
Considerations for Connecting Multiple Devices.....	82
Direct Target Attachment.....	82
Considerations for Direct Target Attachment.....	83
Target Aggregation.....	84
Access Gateway Cascading.....	85
Access Gateway Cascading Considerations.....	86
Fabric and Edge Switch Configuration.....	87
Verifying the Switch Mode.....	87
Enabling NPIV on M-EOS Switches.....	88
Connectivity to Cisco Fabrics.....	88
Enabling NPIV on a Cisco Switch.....	88
Rejoining Fabric OS Switches to a Fabric.....	89
Reverting to a Previous Configuration.....	89
Troubleshooting.....	90
Revision History.....	92
FOS-821-access-AG101; September 28, 2018.....	92
FOS-821-access-AG100; August 28, 2018.....	92

Introduction

• About This Guide.....	6
• What's New in This Document.....	6
• Supported Hardware and Software.....	6
• Key Access Gateway Terms.....	7
• Contacting Brocade Technical Support.....	7
• Document Feedback.....	8

About This Guide

This document describes the Brocade Access Gateway (AG) basic concepts, the platforms involved, configuring ports in AG mode, managing policies, and SAN configuration.

What's New in This Document

This document includes new and modified information for the Fabric OS 8.2.1 release.

Supported Hardware and Software

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems LLC for Fabric OS 8.2.1, documenting all possible configurations and scenarios is beyond the scope of this document.

Brocade Access Gateway features in Fabric OS 8.2.1 are supported on the following hardware platforms.

Brocade Gen 5 Platform (16Gb/s) Fixed-Port Switches

- Brocade 6505 Switch
- Brocade 6510 Switch
- Brocade M6505 blade server SAN I/O module
- Brocade 6542 blade server SAN I/O module
- Brocade 6543 blade server SAN I/O module
- Brocade 6545 blade server SAN I/O module
- Brocade 6546 blade server SAN I/O module
- Brocade 6547 blade server SAN I/O module
- Brocade 6548 blade server SAN I/O module
- Brocade 6558 blade server SAN I/O module

Brocade Gen 6 Platform (32Gb/s) Fixed-Port Switches

- Brocade G610 Switch
- Brocade G620 Switch

Key Access Gateway Terms

For definitions of SAN-specific terms, visit the Storage Networking Industry Association online dictionary at:

<http://www.snia.org/education/dictionary>

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

Access Gateway (AG)	A Fabric OS mode for switches that reduces storage area network (SAN) deployment complexity by leveraging N_Port ID Virtualization (NPIV).
Advanced Device Security (ADS) policy	A security policy that restricts access to the fabric at the AG level to a set of authorized devices.
device	Any host or target device with a distinct WWN. Devices may be physical or virtual.
D_Port	A port configured as a diagnostic port on an AG switch, connected fabric switch, or connected cascaded AG switch to run diagnostic tests between the ports and test the link.
edge switch	A fabric switch that connects host, storage, or other devices, such as switch running in Brocade Access Gateway mode, to the fabric.
E_Port	An inter-switch link (ISL) port. A switch port that connects switches together to form a fabric.
fabric system	A fabric system consists of interconnected nodes that look like a single logical unit when viewed collectively. A fabric system refers to a consolidated high-performance network system that consists of coupled storage devices, networking devices, and parallel processing high-bandwidth interconnects such as 8Gb/s, 10Gb/s, 16Gb/s, or 32Gb/s Fibre Channel ports.
FCoE	Fibre Channel over Ethernet (FCoE) refers to a network technology that encapsulates Fibre Channel frames over Ethernet networks. This allows Fibre Channel to use 10-Gigabit Ethernet or higher-speed networks while preserving the Fibre Channel protocol.
F_Port	A fabric port. A switch port that connects a host, host bus adapter (HBA), 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. An N_Port presents the AG-connected host or storage device to the fabric. On a Brocade Access Gateway, the N_Port connects to the edge switch.
NPIV	N_Port ID Virtualization. NPIV is a Fibre Channel facility that allows multiple F_Port IDs to share a single physical N_Port. Multiple F_Ports can be mapped to a single N_Port. This allows multiple Fibre Channel initiators to occupy a single physical port, easing hardware requirements in SAN design, especially for virtual SANs.
Port Grouping (PG) policy	The Port Grouping (PG) policy is used to partition the fabric, host, or target ports within an AG-enabled module into independently operated groups.

Contacting Brocade Technical Support

For product support information and the latest information on contacting the Technical Assistance Center, go to <https://www.broadcom.com/support/fibre-channel-networking/>. If you have purchased Brocade product support directly from Brocade, use one of the following methods to contact the Brocade Technical Assistance Center 24x7.

Online	Telephone
<p>For nonurgent issues, the preferred method is to go to MyBrocade (my.brocade.com) and then go to one of the following sites:</p> <ul style="list-style-type: none"> • My Cases • Software Downloads • Licensing tools 	<p>Required for Severity 1-Critical and Severity 2-High issues:</p> <ul style="list-style-type: none"> • North America: 1-800-752-8061 (Toll-free) • International: 1-669-234-1001 (Not toll-free) <p>Toll-free numbers are available in many countries and are listed at https://www.broadcom.com/support/fibre-channel-networking/.</p>

Online	Telephone
<ul style="list-style-type: none">• Knowledge Base	

If you purchased Brocade product support from a Brocade OEM/solution provider, contact your OEM/solution provider for all your product support needs.

- OEM/solution providers are trained and certified by Brocade to support Brocade products.
- Brocade provides backline support for issues that cannot be resolved by the OEM/solution provider.
- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information, contact Brocade or your OEM.
- For questions regarding service levels and response times, contact your OEM/solution provider.

Document Feedback

Quality is our first concern. We have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission or if you think that a topic needs further development, we want to hear from you. Send your feedback to documentation.pdl@broadcom.com. Provide the publication title, publication number, topic heading, page number, and as much detail as possible.

Access Gateway Basic Concepts

• Brocade Access Gateway Overview	9
• Fabric OS Features in Access Gateway Mode.....	11
• Access Gateway Port Types.....	21
• Access Gateway Hardware Considerations.....	23

Brocade Access Gateway Overview

Brocade Access Gateway (AG) is a Fabric OS feature that you can use to configure your fabric to handle additional devices instead of domains. You do so by configuring F_Ports to connect to the fabric as N_Ports, which increases the number of device ports that you can connect to a single fabric. When a fabric switch operates in AG mode, it is referred to as an AG device.

Multiple AG devices can connect to the following Brocade products:

- DCX 8510-4 and DCX 8510-8
- X6-4 and X6-8
- 6505
- 6510
- 6520
- 7840
- G610
- G620
- G630
- 7810

Access Gateway is compatible with Cisco-based fabrics that support standards-based N_Port ID Virtualization (NPIV). You can use the command line interface (CLI), Web Tools, or Brocade Network Advisor to enable and disable AG mode and configure the AG features on a switch. This document describes configurations using CLI commands. For more information, refer to the *Brocade Fabric OS Command Reference*, the *Brocade Fabric OS Web Tools Administration Guide*, or the *Brocade Network Advisor User Guide*.

After you enable AG mode on a switch, the F_Ports connect to the fabric as N_Ports rather than as E_Ports.

[Comparing Native Fabric and Access Gateway Modes](#) on page 10 compares the configuration that connects eight hosts to a fabric using AG mode 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. AG mode simplifies configuration and management in a large fabric by reducing the number of domain IDs and ports.

NOTE

In this document, a switch operating in Access Gateway mode is also referred to as an Access Gateway device or AG device.

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 a switch in AG mode to a Fabric OS, M-EOS, or Cisco-based fabric.

The following figures show differences between switch function in native mode and switch function in AG mode.

FIGURE 1 Switch Function in Native Mode





FIGURE 2 Switch Function in Access Gateway Mode



Fabric OS Features in Access Gateway Mode

The following table lists feature support for a switch that operates in Access Gateway mode. Feature support is indicated in the following ways:

- Yes means that the feature is supported in Access Gateway mode.
- No means that the feature is not supported in Access Gateway mode.
- N/A means that the feature is not applicable in Access Gateway mode.
- Yes* means that the feature is transparent in Access Gateway mode.
- Yes** means that the feature is available on a Brocade enterprise fabric, but possibly not available if the enterprise fabric is not a Brocade fabric.

For more information on features listed in the table, refer to the *Brocade Fabric OS Administration Guide* and the *Brocade Fabric OS Command Reference*.

TABLE 1 Fabric OS Components Supported in Access Gateway Mode

Feature	Support
Access Control	Yes (limited roles)
Adaptive Networking	Yes
Admin Domains	No Support is deprecated in Fabric OS 8.0.0 and later.
Audit	Yes
Beaconing	Yes
Bottleneck Detection	No Support is deprecated in Fabric OS 8.0.0 and later.
Buffer Credit Recovery (CR)	Yes See Buffer Credit Recovery Support on page 14.
ClearLink Diagnostic Port (D_Port)	Yes See D_Port Support on page 43.
Config Download/Upload	Yes
Device Authentication	Yes See Device Authentication Support on page 14.
DHCP	Yes
Duplicate PWWN Handling during Device Login	Yes See Duplicate PWWN Handling during Device Login on page 78.
Encryption Configuration and Management	No
Environmental Monitor	Yes
Error Event Management	Yes
Extended Fabrics	No
Fabric Assigned PWWN (FA-PWWN)	Yes
Fabric Device Management Interface (FDMI)	Yes*
Fabric Manager	Yes**
Fabric Performance Impact (FPI)	Yes On an AG switch set up with F_Port trunks, fabric performance is monitored only on those F_Ports actually present on the AG switch.
Fabric Provisioning	No
Fabric Services	No
Fabric Watch	No Support is deprecated in Fabric OS 7.4.0 and later.
Fibre Channel Routing (FCR) Services	No
Flow Vision	Yes
• Flow Monitor	
• Flow Mirror	
FICON (Includes CUP)	No
Forward Error Correction (FEC)	Yes See Forward Error Correction Support on page 14.
High Availability	Yes
Hot Code Load	Yes

TABLE 1 Fabric OS Components Supported in Access Gateway Mode (continued)

Feature	Support
IO Insight	No
License	Yes**
Lightweight Directory Access Protocol (LDAP)	Yes
Log Tracking	Yes
Management Server	NA
Manufacturing Diagnostics	Yes
Monitoring and Alerting Policy Suite (MAPS)	Yes
N_Port ID Virtualization (NPIV)	Yes
Name Server	NA
Native Interoperability Mode	NA
Network Time Protocol (NTP)	Yes
Open E_Port	NA
Performance Monitor	No Support is deprecated in Fabric OS 7.4.0 and later.
Persistent ALPA	Yes
Port Decommission	No
Port Mirroring	No Not supported in previous releases; deprecated in Fabric OS 8.0.0 and later.
QuickLoop, QuickLoop Fabric Assist	No
Remote Authentication Dial-In User Service (RADIUS)	Yes
Resource Monitor	Yes
RLS Probing for NPIV Devices	No
Security	Yes ADS/DCC Policy
Slow-Drain Device Quarantining (SDDQ)	Yes See Slow-Drain Device Quarantining Support on page 19.
SNMP	Yes
Speed Negotiation	Yes
Syslog Daemon	Yes
TACACS+	Yes
Track Changes	Yes
Trunking	Yes**
User-Defined Roles	Yes
Value Line Options (Static POD, DPOD)	Yes
Virtual Fabrics	No See Virtual Fabrics Support on page 14.
VM Insight	No
Web Tools	Yes
Zoning	NA

Buffer Credit Recovery Support

Buffer credit recovery is a Fabric OS feature supported on 8 Gbps, 16 Gbps, and 32 Gbps platforms in the following configurations:

- Between the AG device F_Port and a QLogic BR-1860 16Gbps Host Bus Adapter (HBA) port running version 3.2 or later firmware.
- Between the AG device F_Port and any device that supports credit recovery.
- Between the AG device N_Port and a Brocade fabric switch or a cascaded AG device F_Port.

NOTE

If a device that supports 16 Gbps or 32 Gbps is connected to a device that supports only 8 Gbps, buffer credit recovery is disabled, even if both devices are running 8 Gbps.

Switch platforms support buffer credit recovery in R_RDY or VC_RDY mode. In R_RDY mode, buffer credit recovery is supported without FA-PWWN and QoS. In VC_RDY mode, buffer credit recovery is supported with fabric-assigned PWWN (FA-PWWN), FEC, QoS, and trunking.

Disable this feature on the AG device before connecting to a switch running a version of Fabric OS earlier than 7.1. Use the **portcfgcreditrecovery** command to enable and disable credit recovery. Refer to the *Fabric OS Command Reference* for more information.

Forward Error Correction Support

Forward error correction (FEC) is a Fabric OS feature that is supported in the following configurations:

- Between the AG device F_Port and a QLogic BR-1860 16Gb/s HBA port running version 3.2 or later firmware.
- Between the AG device N_Port and F_Port on a Brocade 16 Gbps or 32 Gbps fabric switch or cascaded AG switch.

Consider the following limitations for FEC:

- Supported on Brocade 16Gb/s and 32Gb/s platforms only.
- Supported by Fabric OS 7.1.0 and later.
- Enabled by default.
- Disabled when Fabric OS is downgraded to a release before Fabric OS 7.1.0.
- Supported on specific switch platforms in R_RDY mode or VC_RDY mode. For information on supported platforms, refer to the "Forward error correction" section of the Brocade *Fabric OS Administration Guide*.

Virtual Fabrics Support

Virtual Fabrics is a Fabric OS feature supported on 8Gb/s, 16Gb/s, and 32Gb/s platforms with the following limitations:

- A switch cannot be enabled for both Virtual Fabrics and AG mode.

Device Authentication Support

By default, Fabric OS 6.2.0 and subsequent releases use Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP) or Fibre Channel Authentication Protocol (FCAP) for authentication. These protocols use shared secrets and digital certificates, based on switch WWN and public key infrastructure (PKI) technology, to authenticate switches.

Authentication automatically defaults to FCAP when both the devices are configured to accept FCAP authentication. To use FCAP on both devices, PKI certificates must be installed. If a PKI certificate is present on each device, FCAP has precedence over DH-CHAP.

If ports are configured for in-flight encryption, authentication defaults to DH-CHAP. Both devices must be configured to accept DH-CHAP for authentication.

If DH-CHAP or FCAP fail to authenticate, the Access Gateway port is disabled.

Authentication policy is supported in the following configurations for Access Gateway devices:

- An Access Gateway device N_Port connected to a Brocade fabric switch F_Port. The N_Port enables authentication when authentication is enabled on the connected switch. Enable the switch policy on the AG device, and enable the device policy on the fabric switch.
- An Access Gateway switch F_Port connected to an HBA. The F_Port enables authentication when the connected device sends a login request with authentication enabled. Enable the device policy on the AG device. For authentication between an AG device and an HBA, use DH-CHAP. The HBA supports DH-CHAP only.

For details on installing FCAP certificates and creating DH-CHAP secrets on the switch in AG or Native mode, refer to the Brocade *Fabric OS Administration Guide* or the Brocade *Fabric OS Command Reference*.

For general information on authentication, refer to the authentication policy for fabric elements in the Brocade *Fabric OS Administration Guide*.

Supported Policy Modes

A switch in Access Gateway mode supports the following switch and device policy modes:

- On: Strict authentication is enforced on all ports. During AG initialization, authentication is enabled on all ports. The ports on the AG device that are connected to another switch or device are disabled when the connected switch or device does not support authentication or when the policy mode is set to off.
- Off: Authentication is disabled. The AG device does not support authentication and rejects any authentication negotiation request from the connected fabric switch or HBA. Off is the default mode for both the switch and the device policy. You must configure DH-CHAP shared secrets or install FCAP certificates on the AG device and connected fabric switch before changing a policy mode from off to on.
- Passive: Incoming authentication requests are accepted. The AG device does not initiate authentication when connected to a device, but it accepts incoming authentication requests if the connecting device initiates authentication. The F_Ports on the AG device are not disabled if the connecting device does not support authentication or the policy mode is off. Passive mode is the safest mode for an AG device when the connected devices do not support authentication.

For device policy support, the AG device supports policy modes on, off, and passive.

For switch policy support, the AG device supports policy modes on and off.

The following tables describe interactions between switch policy modes on the AG device and policy modes on the connected devices for both fabric switches and HBAs.

TABLE 2 Behavior of AG Sending Device and Receiving Fabric Switch with Different Policies Configured

	Fabric Switch, Device Policy Mode On	Fabric Switch, Device Policy Mode Passive	Fabric Switch, Device Policy Mode Off
AG device, switch policy mode On	Authorization negotiation - accept DH-CHAP/FCAP: <ul style="list-style-type: none"> • Success - N_Port • Failure - disable 	Authorization negotiation - accept DH-CHAP/FCAP: <ul style="list-style-type: none"> • Success - N_Port • Failure - disable 	Authorization negotiation - reject N_Port without authentication
AG device, switch policy mode Off	No negotiation No light	No negotiation N_Port without authentication.	No negotiation N_Port without authentication

TABLE 3 Behavior of HBA Sending Device and Receiving AG Device with Different Policies Configured

	AG device, device policy mode On	AG device, device policy mode Passive	AG device, device policy mode Off
HBA, authentication enabled	Authorization negotiation - accept DH-CHAP : • Success - F_Port • Failure - disable	Authorization negotiation - accept DH-CHAP: • Success - F_Port • Failure - disable	Authorization negotiation - reject F_Port without authentication
HBA, authentication disabled	No negotiation No light	No negotiation F_Port without authentication	No negotiation F_Port without authentication

Supported Fabric OS Commands

The following Fabric OS commands for authentication policy apply to AG mode:

- **authutil --policy**
- **authutil --set**
- **authutil --show**
- **seauthsecret --set**
- **seauthsecret --show**

NOTE

Although **authutil --authinit** is not supported in AG mode, it is supported in Native mode.

For more information, refer to the Brocade *Fabric OS Command Reference*.

Limitations and Considerations

Be aware of the following limitations and considerations when configuring the authentication policy on an AG device:

- Authentication policy is not supported on cascaded AG device configurations.
- Authentication is not supported between an AG device running Fabric OS 7.1.0 (or subsequent release) and a fabric switch running a version before Fabric OS 7.1.0. If the AG device is connected to a fabric switch running a version before Fabric OS 7.1.0, the AG device N_Ports will be disabled if the authentication policy is enabled on both switches. Devices mapped to N_Ports connected to fabric switches running any version before Fabric OS 7.1.0 will also be disabled.
- If the authentication policy is disabled on the fabric switch, the AG device N_Port will come online without the authentication policy.
- Device and switch policies must be disabled on the AG device before converting it to Native mode.
- Device and switch policies must be disabled on the fabric switch in Native mode before converting it to AG mode.
- The authentication policy is disabled by default on all ports in AG mode.
- High availability (HA) reboots are supported.

AG Mode without All Ports on Demand Licenses

Consider the following points while running switches in AG mode without Ports on Demand (PoD) licenses:

- By default, configured N_Ports will come up as disabled because a PoD license is not installed for those ports.

- All F_Ports mapped to the default N_Port will come up as disabled with the following message displayed: `N-Port Offline for F-Port.`
- You must manually configure N_Ports in the AG device using the `portcfgnport` command, and you must move the cable connections accordingly.
- You can update the N_Port-to-F_Port mapping using the `ag --mapdel` and `ag --mapadd` commands.

Password Distribution Support

In Fabric OS 7.3.0 and subsequent releases, you can distribute the password database to all switches that are connected to the same fabric whether they are in AG mode or Native mode. Use the `distribute` command on any switch that is in Native mode to distribute the password database to all AG devices and switches connected to the same fabric. You can selectively distribute the password database by specifying the AG device name, or you can use a wildcard-matching character (*) to distribute it to all switches and AG devices.

Consider the following points when configuring password distribution:

- On the AG device, the `fddcfg` command is used to accept or reject the password database from any of the switches in the same fabric.
- To accept the password database from any switch, the AG device must be running Fabric OS 8.0.0 or later, and at least one of the switches in Native mode within the same fabric must be running Fabric OS 7.3.0 or later.
- Other databases supported by `fddcfg` are not supported on AG devices.
- Virtual Fabrics (VF) mode distribution does not apply to an AG device.
- The `distribute` command is not supported in AG mode. Therefore, an AG device cannot distribute its password database to switches that are in Native mode.

FDMI Support

In Fabric OS 7.3.0 and subsequent releases, an AG device can register its N_Port with FDMI devices. Use the `fdmishow` command to display the device details in AG mode. The `fdmishow` command displays only the local devices. Remote device details cannot be displayed.

NTP Configuration Distribution to Access Gateways

Any switch running Fabric OS 7.4.0 or later can distribute the NTP server configuration to core Access Gateway (AG) devices connected to the fabric. The core AG devices must be running Fabric OS 7.4.0 or later to distribute the NTP configuration to other cascaded or edge AG devices. However, AG devices are not capable of distributing the NTP configuration to other switches in the fabric.

A switch running Fabric OS 7.4.0 or later can distribute the NTP server configuration to AG devices connected to the fabric. However, AG devices cannot distribute the NTP configuration to other switches in the fabric.

In switch mode, the principal or primary FCS switch synchronizes its time with the external NTP server every 64 seconds and sends time updates to other switches in the fabric. The time updates are not sent in-band to AG devices. An AG device need not sync with the external NTP server since it can receive the NTP server configuration from the connected Fabric OS switch. If the AG device is connected to more than one fabric, the latest clock server request received is used.

Consider the following points when distributing the NTP server configuration:

- The `tsClockServer` command distributes the NTP server configuration to all switches within the fabric and to AG devices connected to the same fabric.
- Already distributed NTP server configurations will persist on the AG device after downgrading the firmware from Fabric OS 8.0.1 or later to an earlier version on supported platforms.

- Already distributed NTP server configurations will persist on the AG device after downgrading the firmware from Fabric OS 8.0.1 or later to an earlier version on supported platforms. However, the AG device cannot distribute any configuration to the edge AG devices.

Remote Fosexec Support

For Fabric OS 8.1.0 and later, you can enable the Remote Fosexec feature and use the **fosExec** command to issue Fabric OS commands on AG devices in remote domains across the fabric. Before this release, you needed to log in to individual remote AG devices to issue commands.

NOTE

You must use the **configure** command to enable the Remote Fosexec feature on both the sending and receiving AG device. By default, Remote Fosexec is disabled. Refer to the *Brocade Fabric OS Command Reference* for procedures.

When Remote Fosexec is enabled, you can use **fosExec** to enable a command on all AG devices in a fabric or on a specific remote AG device while logged into a local device. The following is an example of using the **fosexec** with **switchshow** to display switch data.

- To issue **switchShow** on a remote AG device named "Core_AG_2", enter the following.

```
switch:admin> fosexec --ag Core_AG_2 -cmd "switchshow"
```

- To issue **switchShow** on all remote AG devices in the fabric, enter the following.

```
switch:admin> fosexec --ag all -cmd "switchshow"
```

You can use **fosExec** to issue commands that display statistics and monitoring (show commands). You cannot issue commands that modify a switch state or display security-relevant states.

The following commands are supported in Fabric OS v8.1.0 and later in AG mode:

- **ag --show**
- **ag --modeshow**
- **ag --mapshow**
- **ag --policyshow**
- **ag --failbackshow**
- **ag --pgshow**
- **ag --failovershow**
- **ag --prefshow**
- **ag --adsshow**
- **ag --wwnmapshow**
- **ag --reliabilitycountershown**
- **ag --reliabilityshow**
- **ag --backupmappingshow**

NOTE

The **fosExec** command cannot be issued on a device configured in AG mode. Therefore, you cannot issue commands on a remote AG device from a device configured in AG mode.

Commands issued using **fosExec** have the following limitations:

- They can fetch 64 KB of data from a remote switch or domain only.

- They are not supported in FIPS mode.
- They cannot be interactive.
- They will not work if they take longer than 15 seconds to complete.


For details on the Remote Fosexec feature and using the `fosExec` command for issuing commands on remote AG devices, refer to the *Brocade Fabric OS Command Reference*.

Slow-Drain Device Quarantining Support

In a fabric, many flows share the same link or virtual channel (VC). However, the credits used to send traffic across the link are common to all flows that use the same link. Therefore, a slow-draining device slows down the return of credits and impacts the healthy flows using the link.

The following figure illustrates how this functions. In the figure, the inability of the slow-draining device (1) to clear frames quickly enough that causes backpressure not just to the edge switch or AG device (2), but also to the core switch (3) and the target (4) that is trying to send data to the slow-draining device.

FIGURE 3 Slow-Drain Dataflow



To isolate the traffic destined to slow-drain the device or port, the Slow-Drain Device Quarantine (SDDQ) feature of Fabric Performance Impact (FPI) quarantines the port. SDDQ moves all the traffic destined for the slow-drain devices to a low priority VC and adds the port to the predefined ALL_QUARANTINED_PORTS MAPS logical group. The RASLOG messages are displayed to notify that the traffic has been successfully moved to a lower priority VC.

The following example displays the RASLOG messages for SDDQ:

```
[NS-1017], 6, SLOT 6 | FID 21, INFO, DCX_120_21, Domain 34, Port index 171, All devices zoned with the slow
drain device 0x220d00 have been quarantined.
[MAPS-1022], 7, SLOT 6 | FID 21, WARNING, DCX_120_21, Port 3/27 (Port index 171) has been marked as Slow
Drain Device.
[AG-1086], 513, FID 128, INFO, G610, Slow drain device 0x011205 connected to F-port (8) have been
quarantined.
[AG-1087], 528, FID 128, INFO, G610, Slow drain device 0x011205 connected to F-port (8) have been
unquarantined.
```

For more information on RASLOG messages, refer to the *Brocade FOS Message Reference Guide*.

Once MAPS quarantine a port, it continues to monitor the slow-drain port and if the condition gets cleared then MAPS automatically performs the un-quarantine action (if configured). The user can also manually un-quarantine the port. For automatic or manual clearing of quarantined ports, refer to the *Brocade Monitoring and Alerting Policy Suite Configuration Guide*.

Access Gateway (AG) has the capability to detect slow-drain behavior for the devices that log on through the AG switch.

AG cannot detect the slow-drain behavior of devices that do not log on through the AG switch. For example, if host A login through AG1 and talks to a target of another switch with the slow-drain behavior, then the AG cannot detect the target as the slow-drain device. The slow-drain behavior is detected only by the switch used to login into the device.

The following table shows the SDDQ behavior with different firmware loaded in Edge AG, Core AG, and Edge Fabric Switch.

TABLE 4 Port configurations

Edge AG	Result/Behavior
	o r e A G
FOS 8.2.1	No device will be quarantined O S 8 .2 .0 x
FOS 8.2.0x	No device will be quarantined O S 8 .2 .1
FOS 8.2.0x	Devices in Core AG will be quarantined O S 8 .2 .1
FOS 8.2.1	Devices in both Core AG and Edge AG will be quarantined O S 8 .2 .1
FOS 8.2.1	No device will be quarantined O S 8 .1

TABLE 4 Port configurations (continued)

Edge AG	Result/Behavior o r e A G
	2 .1
FOS 8.2.1	Devices in Edge AG will be quarantined O S 8 .2 .0 x

Access Gateway Port Types

An AG device connects to the fabric by means of node ports (N_Ports). Typically, switches connect to the fabric using inter-switch link (ISL) ports, such as E_Ports.

AG devices use the following Fibre Channel (FC) ports:

- F_Port: A fabric port that connects a host, HBA, or storage device to an AG device.
- N_Port: A node port that connects an AG device to the F_Port of the fabric switch.
- D_Port: A diagnostic port that is configured in diagnostic mode to run tests between it and a connected D_Port on another switch or HBA.

NOTE

Use the `portcfgpersistenable` command on all external (outward-facing) ports to ensure that these ports come back online after a switch reboot or power failure. To ensure that this command persists on an embedded switch, enter the `portcfgpersistenable` command through the chassis management console and not the CLI. Refer to [Persistent Port Online State](#) on page 42 for more information.


Comparison of Access Gateway Ports to Standard Switch Ports

The Access Gateway device multiplexes host connections to the fabric. It presents an F_Port to the host and an N_Port to the edge fabric switch. Using N_Port ID Virtualization (NPIV), the AG device allows multiple FC hosts to access the SAN on the same physical port. This feature 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 it participates in fabric management and zoning distribution. A fabric switch requires more physical ports than an AG device to connect the same number of hosts.

The following figure compares the ports that are used between AG mode and native mode.

FIGURE 4 Port Comparison



You can test the link between ports on an AG device and another AG device, fabric switch, or HBA by configuring each connection in the link as a D_Port. When you configure the ports at each end of the link as D_Ports, diagnostic tests automatically initiate on the link when the D_Ports come online. You can view results by using Fabric OS commands, such as **portdporttest**, during or after testing. When configured as a D_Port, the port does not participate in fabric operations, log in to a remote device, or transmit data traffic. The following figure shows D_Port links between multiple devices.

FIGURE 5 Diagnostic Port Configurations

The following table compares port configurations between AG devices and a standard fabric switch.

TABLE 5 Port Configurations

Port Type	Available on the Access Gateway?		Available on the Fabric Switch?	
F_Port	Yes	Connects hosts and targets to the Access Gateway.	Yes	Connects devices, such as hosts, HBAs, and storage, to the fabric.
N_Port	Yes	Connects Access Gateway to a fabric switch.	N/A	N_Ports are not supported.
E_Port	N/A	ISL is not supported. The switch is logically transparent to the fabric; therefore it does not participate in the SAN as a fabric switch.	Yes	Connects the switch to other switches to form a fabric.
D_Port	Yes	Allows diagnostic testing across link to connected AG device, fabric switch, or HBA.	Yes	Allows diagnostic testing across the link to the connected AG device.

Access Gateway Hardware Considerations

Hardware considerations for the Access Gateway devices are as follows:

- The Access Gateway feature is supported on the switch platforms and embedded switch platforms listed in [Supported hardware and software](#).
- Loop devices are not supported.
- Direct connections to SAN target devices are supported only when the AG device is connected to a fabric. Cascaded AG devices do not support SAN target devices.

Configuring Ports in Access Gateway Mode

• Enabling and Disabling Access Gateway Mode.....	24
• Access Gateway Mapping.....	25
• N_Port Configurations.....	40
• D_Port Support.....	43

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. After enabling, you can display the default port mappings and the status of the host connections to the fabric. For more information on the `ag` commands used in these steps, refer to the *Brocade Fabric OS Command Reference*.

When you disable AG mode, the switch automatically reboot in Fabric OS Native mode and comes back online using the fabric switch configuration. The AG mode parameters, such as port mapping, and failover and fallback, are automatically removed. To rejoin the switch to the core fabric, refer to the [Rejoining Fabric OS Switches to a Fabric](#) on page 89.

1. Connect to the switch and log in using an account with admin permissions.
2. Before enabling or disabling AG mode on a switch, use the `configUpload` command to save the current configuration file.
3. Ensure that no zoning transaction buffers are active. If any transaction buffer is active, enabling AG mode will fail with the error `Failed to clear Zoning configuration.`
4. Verify that the switch is set to Native mode.
 - a) Use the `switchShow` command to verify the switch mode.
 - b) If the switch mode is not Native, use the `ag --modedisable` command to set the switch to Native mode.

For more information on setting switches to Native mode, refer to the *Brocade Fabric OS Administration Guide*.

5. 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 is part of a fabric, the remaining switches reconfigure. You must disable the switch before making configuration changes.

6. Enter the `ag --modeenable` command.

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

The switch automatically reboots and comes back online in AG mode using the factory default port mapping. For more information on AG mode default port mapping, refer to [Table 8](#) on page 28.

7. Enter the `ag --modeshow` command to verify that AG mode is enabled.

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

8. Use the `ag --mapshow` command to display all mapped ports.

The `ag --mapshow` command shows all enabled N_Ports, even if those N_Ports are not connected.

9. Use the **switchshow** command to display the status and port state of all ports.

Refer to the Brocade *Fabric OS Command Reference* for examples of output. For a description of the port state, refer to [Table 6](#) on page 25.

10. Enter the **switchdisable** command to disable the switch.

```
switch:admin> switchdisable
```

11. Enter the **ag --modedisable** command to disable AG mode.

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

12. 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 lists and describes the port states.

TABLE 6 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	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.
Lock_Ref	Locking to the reference signal.
Testing	Running diagnostics.
Offline	Connection not established (only for virtual ports).
Online	Port is up and running.

Access Gateway Mapping

When a switch operates in AG mode, you must specify routes that the AG device uses to direct traffic from the devices on its F_Ports to the fabric Ports connected to its N_Ports. The routes must be preprovisioned. The process of preprovisioning routes in AG mode is called mapping. (By comparison, a switch operating in Native mode determines the best routing path between its F_Ports.)

You can create two types of maps; Port maps and device maps. Port maps are required. Device maps are optional and assign device WWNs to N_Ports and N_Port groups. Port mapping and device mapping operate as follows:

- Port mapping

Port mapping ensures that all traffic from a specific F_Port always goes through the same N_Port. A single F_Port is mapped to a single N_Port, or to an N_Port group. To map an F_Port to an N_Port group, map the F_Port to an N_Port that belongs to that port group. All F_Ports mapped to an N_Port group are part of that N_Port group.

- Device mapping

Device mapping is optional. Port maps must exist before you can create device maps. Device mapping allows a virtual Port to access its destination device regardless of the F_Port where the device resides. Device mapping also allows multiple virtual Ports on a single physical machine to access multiple destinations residing in different fabrics.


The preferred method is to map a device WWN to an N_Port group. When a device WWN is mapped to an N_Port, and a failover N_Port is also specified, the device can reach the fabric through the primary or secondary N_Port only. However, when a device WWN is mapped to a port group, it can log in to the fabric until the last N_Port in the particular port group remains online.

You can map a device to multiple groups. Alternatively, you can map a device to a specific N_Port.

Port Mapping

F_Ports must be mapped to N_Ports before the F_Ports can come online. The following figure shows an example in which eight F_Ports are mapped evenly to four N_Ports on a switch in AG mode. The N_Ports connect to the same fabric through different edge switches.

FIGURE 6 Port Mapping Example



The following table lists and describes the port mapping.

TABLE 7 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

Default Port Mapping

When you enable AG mode on a switch, a default mapping is used for the F_Ports and N_Ports. The following table describes the default port mapping for all supported hardware platforms.

NOTE

By default, Failover and Failback policies are enabled on all N_Ports.

To change the default mapping, refer to [Adding F_Ports to an N_Port](#) on page 31. Note that all F_Ports must be mapped to an N_Port before the F_Port can come online.

NOTE

For Fabric OS 7.3.0 and later, all PoD licenses are not required to run in AG mode.

TABLE 8 Access Gateway Default Port Mapping

Brocade Model	Total Ports	F_Ports	N_Ports	Default port mapping
6505	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
M6505	24	1-16	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
6510	48	0-39	40-47	0-4 mapped to 40 5-9 mapped to 41 10-14 mapped to 42 15-19 mapped to 43 20-24 mapped to 44 25-29 mapped to 45 30-34 mapped to 46 35-39 mapped to 47
6543	24	0, 9-23	1-8	9,17 mapped to 1 10,18 mapped to 2 11,19 mapped to 3 12,20 mapped to 4 13,21 mapped to 5 14,22 mapped to 6 15,23 mapped to 7 0,16 mapped to 8
6545	26	1-18	0, 19-25	1,2,17 mapped to 19 3,4,18 mapped to 20

TABLE 8 Access Gateway Default Port Mapping (continued)

Brocade Model	Total Ports	F_Ports	N_Ports	Default port mapping
				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
6546	24	0-15	16-23	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
6547	48	1-28	0,29-47	1, 21 mapped to 0 2, 22 mapped to 29 3, 23 mapped to 30 4, 24 mapped to 31 5, 25 mapped to 32 6, 26 mapped to 33 7, 27 mapped to 34 8, 28 mapped to 35 9 mapped to 36 10 mapped to 37 11 mapped to 38 12 mapped to 39 13 mapped to 40 14 mapped to 41 15 mapped to 42 16 mapped to 43 17 mapped to 44 18 mapped to 45 19 mapped to 46 20 mapped to 47
6548	48	1-16	0, 17-27	1, 13 mapped to 0 2, 14 mapped to 17 3, 15 mapped to 18

TABLE 8 Access Gateway Default Port Mapping (continued)

Brocade Model	Total Ports	F_Ports	N_Ports	Default port mapping
				4, 16 mapped to 19 5 mapped to 20 6 mapped to 21 7 mapped to 22 8 mapped to 23 9 mapped to 24 10 mapped to 25 11 mapped to 26 12 mapped to 27
6558	48	1-24	0, 25-47	24 mapped to 32 23 mapped to 33 22 mapped to 34 21 mapped to 35 20 mapped to 36 19 mapped to 37 18 mapped to 38 17 mapped to 39 16 mapped to 25 15 mapped to 26 14 mapped to 27 13 mapped to 28 12 mapped to 29 11 mapped to 30 10 mapped to 31 9 mapped to 0 8 mapped to 40 7 mapped to 41 6 mapped to 42 5 mapped to 43 4 mapped to 44 3 mapped to 45 2 mapped to 46 1 mapped to 47
G610	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

TABLE 8 Access Gateway Default Port Mapping (continued)

Brocade Model	Total Ports	F_Ports	N_Ports	Default port mapping
				8,9 mapped to 20 10,11 mapped to 21 12,13 mapped to 22 14,15 mapped to 23
G620	64	0-39	40-47	0-4 mapped to 40 5-9 mapped to 41 10-14 mapped to 42 15-19 mapped to 43 20-24 mapped to 44 25-29 mapped to 45 30-34 mapped to 46 35-39 mapped to 47

Considerations for Initiator and Target Ports

The following connections are possible for the Fibre Channel Protocol (FCP) initiator (host) and target ports through an AG device:

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

NOTE

Communication between host and target ports is not supported if both are mapped to the same N_Port.

Use the following recommendations for mapping between host and target ports:

- Use separate port groups for the host and target ports.
- If connecting a host and target port to the same AG device, map the host and target to separate N_Ports and connect those N_Ports to the same fabric.
- When configuring secondary port mapping for failover and fallback situations, make sure that host and target F_Ports do not fail over or fail back to the same N_Port.

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 with admin permissions.

2. Enter the **ag --mapadd n_portnumber f_port1;f_port2;...** command to add the list of F_Ports to the N_Port.

The F_Port list can contain multiple F_Port numbers separated by semicolons. In the following example, F_Ports 6 and 7 are mapped to N_Port 13.

```
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.

Removing F_Ports from an 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 remove an F_Port from an N_Port.

1. Connect to the switch and log in using an account assigned with admin permissions.
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 62 for instructions.

3. Enter the **ag --mapdel n_portnumber f_port1;f_port2;...** command to remove F_Ports from an N_Port.

The F_Port list can contain multiple F_Port numbers separated by semicolons. In the following example, F_Ports 17 and 18 are removed from the N_Port where they were mapped.

```
switch:admin> ag --mapdel 40 "17;18"
F-Port to N-Port mapping has been updated successfully
```

4. Enter the **switchShow** command to verify that the F_Port mapping is deleted and not assigned.

The output for this command displays the following information in the Proto column for each unassigned F_Port: **Disabled** (No mapping for F_Port)

F_Port Static Mapping

The F_Port Static Mapping feature adds the **staticadd** and **staticdel** keywords to the **ag** Fabric OS command. These keywords simplify how you change N_Port to F_Port mapping. Using the **ag staticadd** command, any existing mappings are removed and replaced with the new mapping information. The actions of the **ag --mapdel** and **ag --mapadd** commands are combined.

Use the following steps to change F_Port to N_Port mapping.

1. Connect to the switch and log in using an account with admin permissions.
2. Use the **ag** command with the **staticadd** keyword to add or change static port mapping. In the following example, F_Port 1 is added to N_Port 17 and F_Port 1 is removed from any existing N_Port mapping.

```
switch:admin> ag --staticadd 17 1
```

When the F_Port static mapping is changed or added, the F_Port and all attached devices log out of the previously mapped N_Port and log in to the new N_Port.

3. Use the **ag** command with the **staticdel** keyword to remove static port mapping.

In the following example, F_Ports 3, 4, and 5 are removed from N_Port 17.

```
ag --staticdel 17 "3;4;5"
```

Considerations for using F_Port Static Mapping with Other AG Features and Policies

Consider the following when using F_Port Static Mapping with Access Gateway features and policies:

- F_Port Static Mapping works with cascaded Access Gateway configurations.
- Failover, fallback, and preferred secondary N_Port settings are disabled for F_Ports that are statically mapped.
- Statically mapped ports are blocked from using the Automatic Port Configuration (APC) and Advanced Device Security (ADS) policies. You cannot enable the APC policy until all static mappings are deleted using the `ag --staticdel` command.
- F_Port Static Mapping works with the Port Grouping (PG) policy with some modifications to policy behavior. If static mapping is applied to an F_Port already mapped to an N_Port, the F_Port loses its mapping to the N_Port applied through the Port Grouping policy. Therefore, the F_Port does not have the failover, fallback, or preferred N_Port settings that other F_Ports have when mapped to an N_Port in that port group. To remap to an N_Port with PG policy attributes, use the `ag --staticdel` command to remove the static mapping, and then remap to another N_Port using the `ag --mapadd` command.
- F_Port Static Mapping does not work with Device Load Balancing. Because F_Port Static Mapping forces the F_Port to map to a specific N_Port, NPIV devices that log in to the F_Port cannot redistribute themselves among N_Ports in the port group.
- F_Port Static Mapping does not work with port trunking. If an F_Port is statically mapped to an N_Port and trunking is enabled, the F_Port goes offline. If port trunking is enabled for an F_Port already, you cannot configure static mapping for the F_Port.

Upgrade and Downgrade Considerations

- All static mappings are maintained when upgrading to the latest Fabric OS version.
- When downgrading, you must remove all static mappings or downgrade will not be allowed.

Device Mapping

Device mapping allows you to map individual N_Port ID Virtualization (NPIV) devices to N_Ports. By mapping device WWNs directly to an N_Port group, traffic from the device will always go to the same N_Port group, independently of the F_Port where the device logs in. When the Port Grouping and Device Load Balancing policies are 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 53).

NOTE

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

Device mapping does not affect or replace the traditional port mapping. Device mapping is optional and is in addition to the existing port mapping. In general, it is recommended that you map devices to N_Port groups rather than map devices to individual N_Ports within a port group. Group mapping ensures maximum device up-time during failover conditions and system power up. Connections occur more quickly when a large number of devices must connect to the same fabric through a single port group.

The following aspects of device mapping are important to note:


- Device mapping has priority over port mapping. That is, logins from a device mapped to a specific N_Port group or N_Port always have priority over unmapped devices that log in to an F_Port that has been mapped to the same N_Port group or N_Port.
- 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.

The following figure 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 7 Example of device mapping to N_Port groups



The following figure 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 8 Example device mapping to an N_Port

Static Versus Dynamic Mapping

Device mapping is of two types, static or dynamic. Static device mapping has the following considerations:

- Static device mapping is when a device is mapped to an N_Port group or N_port.
- Static device mappings persist across reboots.
- Static device mappings can be saved and restored with the **configUpload** and **configDownload** commands.

Dynamic device mapping has the following considerations:

- Dynamic device mapping is when Automatic Device Load Balancing is enabled.
- Dynamic device mapping exists only while a device is logged in. When the device logs out, the mapping for that device no longer exists.
- Dynamic device mapping cannot be saved or edited, and does not persist when the switch reboots.
- Dynamic mapping shows the current mapping for devices as opposed to the original static mapping. If a device is mapped to an N_Port group, then all mapping is dynamic.

NOTE

Static and dynamic mapping only applies to NPIV devices and cannot redirect devices that are physically attached to Access Gateway devices because physically-attached devices use port maps to connect to the fabric.

Device Mapping to Port Groups (Recommended)

Device mapping is recommended when a number of devices must connect to the same group of N_Ports. This approach provides the flexibility to move the devices to any available F_Port. Device mapping provides load balancing and connects the device to the least-loaded N_Port in the group when the device comes online. For more information on port groups, refer to [Port Grouping Policy](#) on page 50.

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

1. Connect to the switch and log in using an account assigned to the admin role.
After you have logged in, you can apply any of the following commands to configure how devices are mapped to port groups.
2. To add one or multiple device WWNs to an N_Port group , use the **ag --addwwnpgmapping** command.

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

3. To change all currently existing device mappings to a different port group, use the **--all** keyword 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 --addwwnpgmapping 3 --all
```

4. To remove one or multiple devices to an N_Port group , use the **ag --delwwnpgmapping** command.

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

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

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

5. To remove all devices mapped to an N_Port group, use the command with the **--all** keyword 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** keyword is a shortcut for specifying all of the devices that are already mapped with the **ag --addwwnpgmapping** command.

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

```
ag --delwwnpgmapping 3 --all
```

6. Optionally, you can use the **ag --wwnmapshow** 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 with admin permissions.
After you have logged in, you can apply any of the following commands to configure how devices are mapped to port groups.

2. To add one or multiple devices to an N_Port, use the **ag --addwwnmapping** command. All listed device WWNs 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** option edits all the currently existing mappings. None of the **--all** options can detect what devices are using the switch. This option affects the mappings that are in the list.

3. To change all current device mappings to a different N_Port, use the **ag --addwwnmapping** command with the **--all** keyword.

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

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

4. To remove mapping for one or multiple devices from an N_Port, use the **ag --delwwnmapping** command. All listed device WWNs no longer attempt 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"
```

5. To remove all devices currently mapped from an N_Port, use the **ag --delwwnmapping** command with the **--all** keyword. All listed devices no longer attempt to use the N_Port unless a device logs in through an F_Port that is mapped to the N_Port. The **--all** keyword is a shortcut for specifying all of the devices that are already mapped with the **ag --addwwnmapping** command.

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

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

6. Optionally, use the **ag --wwnmaphow** 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_Ports.

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, and then enable this at a later time without reconfiguring your original mapping.

1. Connect to the switch and log in using an account assigned to the admin role.
After you have logged in, you can one of the following commands to disable how devices are mapped to port groups.
2. Use the **ag --wwnmappingdisable** command to disable mapping for specific WWNs. The device mappings are 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"
```

3. Use the **ag --wwnmappingdisable** command with the **--all** to disable mapping for all available WWNs. The **--all** keyword 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 devices or specific devices that were previously disabled.

1. Connect to the switch and log in using an account with admin permissions.
After you have logged in, you can one of the following commands to enable how devices are mapped to port groups.
2. Use the **ag --wwnmappingenable** command 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"
```

3. Use the **ag --wwnmappingenable** command with the **--all** keyword to enable mapping for all currently available WWNs. The **--all** keyword does not affect mappings made in the future. Any mapping added for a new device for which mapping is not specifically disabled is enabled by default. Disabled mappings can be modified without automatically enabling them.

The following example enables all previously disabled device mappings.

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

Displaying Device Mapping Information

The **ag --wwnmapshow** command displays 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
- Current - The N_Port that the device is using (none displays if the device is not logged in)
- Enabled - Indicates whether device mapping is enabled or disabled

NOTE

New device mappings are only enabled and displayed the next time the device logs in to the switch.

To display device mapping information, use the **ag --wwnmapshow** command.

Pre-Provisioning

Pre-provisioning is when you use Fabric OS commands, Web Tools, and Fabric Manager to map devices not yet connected to the fabric. Pre-provisioning allows management programs to push configuration changes with no concern for the order in which the changes are received. For example, if system administrators must push port group changes and device mapping changes, those changes can be pushed in either order without error. Pre-provisioning also applies to using Fabric OS commands for device mapping.

VMware Configuration Considerations

When device mapping is enabled for individual virtual machines (VMs) running on a VMware ESX server connected to an F_Port, network traffic for those VMs is redirected, but only when the ESX server is configured to use Raw Device Mapped storage. All traffic originates from a VM WWN and follows any mapping configured for the WWN.

If anything interrupts the virtual port connection for the VM, such as a failover port being used, traffic will originate from the ESX server base device port ID and not the VM port ID.

If there are any additional disruptions, the server does not switch back to the virtual port, and the VM traffic does not follow the configured device mapping. Note that this behavior can also occur when a VM first boots, prior to any failover.

When this behavior occurs, the VM WWN will be logged in to the fabric. The WWN appears in the output of **ag --show** and **ag --wwnmapshow**, as well as on the switch. The **portperfshow** command displays all traffic on the port to which the ESX server port is mapped (base PID).

Configuring Device Mapping

To configure WWN mapping on VMware ESX systems, use the following steps.

1. Make sure that virtual world wide port names (VWWPNs) of virtual machines (VMs) are mapped to the correct port group (or N_Port). Brocade recommends that you 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. Reboot the VM.
5. Zone the physical port on the server to the storage device.
6. Check the traffic that originates from the virtual node PID (VN PID). If the configuration is correct, traffic will flow from the VN PID.

Failover and Failback Considerations

When using device mapping with VMware, the base device initiates port login (PLOGI) and process login (PRLI) to the target, and then discovers the logical unit number (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 resumes from the base device. One of the following actions is recommended when using device mapping with VMware:

- Make sure targets can be reached by the base device so that network traffic can resume if the mapped device fails over and traffic moves over to the base PID.
- Reboot the server so that it initializes and uses configured device mapping.

Considerations for Access Gateway Mapping

This section presents considerations and limitations for Access Gateway mode mapping types.

Mapping Priority

To avoid problems when both port and device mapping are implemented, AG mode uses a priority system to select the N_Port where a fabric login (FLOGI) is routed. Access Gateway mode considers mappings in the following order until one can be used.

NOTE

Only NPIV devices can use device mapping and the automatic Device Load Balancing policy. Device Load Balancing policy is enabled per module rather than per port group.

For more information, refer to [Port Grouping Policy](#) on page 50.

1. Static device mapping to N_Port (if defined)
2. Device mapping to N_Port group (if defined)
3. Automatic Device Load Balancing within a port group (if enabled)
4. Port mapping to an N_Port

5. Port mapping to an N_Port in a port group (if defined)

Device Mapping Considerations

Consider the following points when using device mapping:

- When the N_Port is disabled, all devices that are mapped to it are disabled. Depending on the effective failover policy, the devices are enabled on other N_Ports.
- Similar to port mappings, device mappings are affected by changes to underlying F_Ports. In other words, if an F_Port is taken offline, both the physical device and all virtual nodes behind it momentarily go offline.
- When 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.
- In case of switch toggle or reboot, if the device comes up before its primary N_Port, it will go online through its failover N_Port, any available N_Port in the port group (device mapping) or through its physically mapped N_Port (port mapping) based on the port's availability.
- In some instances, two NPIV devices logging in to the same F_Port are mapped to two different N_Ports, which in turn are connected to two different fabrics. When this occurs, both NPIV devices can be allocated the same PID by their respective fabrics. 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 PID collisions.

- Static and dynamic device mapping are only supported on the edge module in a cascaded Access Gateway device configuration.
- When mapping devices to a port group, make sure all ports in that 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 can repeatedly attempt to connect to the device with the lower limit, because it has the fewest logins, and fail to connect.

N_Port Configurations

For embedded switches, the internal ports on an Access Gateway device are configured as F_Ports by default. All external ports are configured, or locked, as N_Ports. On standalone switches that support AG mode, a preset number of ports are locked as N_Ports and the remaining ports operate as standard F_Ports.

The following figure shows a host connected to external ports of an embedded switch that is in AG mode.

FIGURE 9 Example of adding an external F_Port (F9) on an embedded switch

Although some ports are locked as N_Ports, you can convert N_Ports to F_Ports, as follows.

1. Before converting an N_Port to an F_Port, remap all F_Ports on that N_port to another N_Port.
2. Remove all the F_Ports that are mapped to the selected N_Port.
3. Unlock the selected port from N_Port state.
4. Define a map for the port.

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.

NOTE

If the Automatic Port Configuration (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, refer to [Table 8](#) on page 28.

Displaying N_Port Configurations

Use the following steps to determine which ports on a switch are locked as N_Ports.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **portcfgnport** command. The command output displays ON for locked N_Ports.

Unlocking N_Ports

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

On fabric switches in Native mode, 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.

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

The following steps show how to unlock and lock N_Ports.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the **portcfgnport** command to display which ports on the switch are locked as N_Ports. Command output will display ON for locked N_Ports.

NOTE

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

3. To unlock N_Port mode, enter the **portcfgnport** command and specify the N_Port port number and 0 (zero).

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

4. To lock a port in N_Port mode, enter the **portcfgnport** command and specify the port number and 1.

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

Persistent Port Online State

Use the **portCfgPersistentEnable** command on all external or outward facing ports to ensure that these ports come back online after a switch reboot or power failure. For an embedded switch, execute this command through the chassis management console and not the Fabric OS CLI or the command may not persist.

NOTE

If the port is connected to another switch when this command is entered, the fabric may reconfigure.

After the port is persistently enabled, devices connected to the port can again communicate with the fabric when the port comes back online. You can identify a single port to be configured by its port number or by its port index number. Port ranges are supported with index numbers or by specifying a slot or a slot range. Use the **switchShow** command for a list of valid ports, slots, and port index numbers.

As an example, to persistently enable a single port, enter the following command.

```
switch:admin> portcfgpersistentenable 4
```

Refer to the *Brocade Fabric OS Command Reference* for more information.

D_Port Support

The D_Port (diagnostic port) feature is supported on 16 Gbps and 32 Gbps ports in the configurations shown in the following table:

TABLE 9 D_Port Support for Access Gateway

D_Port mode	Switch to AG	AG to AG	AG to HBA
Static D_port	Supported	Supported	Supported
Dynamic D_port	Not supported	Not supported	Supported
On-Demand D_port	Not supported	Not supported	Not supported

D_Port modes function as follows:

- Static D_Port: You can convert a Fibre Channel port to a D_Port on an AG device, an HBA, or a fabric switch or another AG device in a cascaded configuration to test the link between the ports. When you configure the ports on each end of the link as a D_Port, diagnostic tests automatically initiate on the link when the D_Ports come online. After the port is in D_Port mode, the port does not participate in fabric operations, log in to a remote device, or run data traffic. Refer to [Comparison of Access Gateway Ports to Standard Switch Ports](#) on page 21 for information on the supported D_Port configurations.

NOTE

N_Port-to-F_Port mappings must be removed from an AG device F_Port before configuring it as a D_Port. Refer to [Saving Port Mappings](#) on page 44 for more information.

- Dynamic D_Port: Supported on AG ports, starting with Fabric OS 7.3.0, to support on-demand D_Port mode on attached HBAs. If the attached HBA port supports the D_Port FLOGI profile, the AG device port automatically enters dynamic D_Port mode when requested for diagnostic testing. The AG port automatically reverts to F_Port mode when the HBA or device port reverts back to a normal port.

To verify dynamic D_Port support on the switch, enter the **configure** command:

```
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
F-Port login parameters (yes, y, no, n): [no]
D-Port parameters (yes, y, no, n): [no] y
Dynamic D-Port (on, off): [on]
```

You can view the results from D_Port testing by using Fabric OS commands such as **portDPortTest --show** during or after testing. The adapter, fabric switch, or AG device is the initiator and the other device is the responder. You can view detailed results from the initiator.

The following tests automatically run on the link between configured D_Ports:

- Electrical loopback in SFP
- Optical loopback in SFP
- Link traffic test to test frame flow to ASIC or CPU
- Link latency and distance measurement

For details on configuring D_Ports, using D_Ports, and D_Port limitations and considerations, refer to the *Brocade Fabric OS Troubleshooting and Diagnostics Guide*. For details on D_Port-related commands, refer to the *Fabric OS Command Reference*.

Saving Port Mappings

Before configuring static D_Ports, you must remove all mappings between the ports and devices because mappings are not retained on D_Ports. This includes port (N_Port-to-F_Port), device (WWN), static, and dynamic mapping. You can use Fabric OS commands to save N_Port mappings and to delete saved N_Port mappings. When you disable the D_Ports, use the saved mapping information to assist manual reconfiguration of the deleted port mappings. Refer to the *Brocade Fabric OS Command Reference* for more details on the backup mapping commands.

To save port mappings use the **ag --backupmappingsave** command with the *N_Port* variable. The following example saves mappings on N_port 44.

```
switch:admin>ag --backupmappingsave 44
Configured static and preferred mappings have been saved for the N_port successfully
N_Port          44
Backed-up Configured F_Ports          20:21:22
Backed-up Static F_Ports             23:24
Backed-up Preferred F_Ports          26:27:28:29
```

Displaying port mappings

To display saved port mappings, use the **ag --backupmappingshow** command with the *N_Port* variable to display saved N_Port mappings. You can use the displayed information to reconfigure the ports. The following example displays mappings on N_port 44.

```
switch:admin>ag --backupmappingshow 44
N_Port          44
Backed-up Configured F_Ports          20:21:22
Backed-up Static F_Ports             23:24
Backed-up Preferred F_Ports          26:27:28:29
```

Deleting port mappings

To delete port mappings, use the **ag --backupmappingdel** command with the *N_Port* variable to delete configured N_Port mappings. The following example deletes mappings on N_port 44.

```
switch:admin>ag --backupmappingdel 44
```

Limitations and Considerations

Consider the following limits and conditions when using D_Ports in AG device configurations:

- Any D_Port must be configured on the AG device, fabric switch, cascaded AG device, or HBA before enabling D_Ports on both sides of the link. Otherwise, the port will be persistently disabled.
- After configuring a D_Port for an AG device port, mapping is not retained. Static D_Port configuration cannot be made unless mappings are removed from the port. This includes F_Port-to-N_Port, static, preferred, and device (WWN) mapping. Therefore, all mappings must be manually removed on the Access Gateway port before configuring the port as a D_Port.

For a complete list of D_Port limitations and considerations, refer to the *Fabric OS Administration Guide*.

Managing Policies and Features in Access Gateway Mode

• Access Gateway Policies Overview.....	45
• Advanced Device Security Policy	46
• Automatic Port Configuration Policy	49
• Port Grouping Policy.....	50
• Device Load Balancing Policy.....	56
• Persistent ALPA Policy.....	57
• Failover Policy.....	59
• Failback Policy.....	64
• Trunking in Access Gateway Mode.....	68
• Adaptive Networking on Access Gateway.....	75
• Per-Port NPIV Login Limit.....	77
• Duplicate PWWN Handling during Device Login.....	78
• Performance Monitoring.....	78
• Considerations for the Brocade Switches.....	81

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 use the **ag --policyshow** 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.

Access Gateway Policy Enforcement Matrix

The following table shows which policies can be enabled at the same time. For example, in the Auto Port Configuration policy row, only N_Port Trunking and Advanced Device Security can be enabled with this policy.

TABLE 10 Policy enforcement matrix

Policies	Auto Port Configuration	N_Port Grouping	N_Port Trunking	Advanced Device Security
Auto Port Configuration	N/A	No	Yes	Yes
N_Port Grouping	Mutually exclusive	N/A	Yes	Yes
N_Port Trunking	Yes	Yes	N/A	Yes
Advanced Device Security ²	Yes	Yes	Yes	N/A
Device Load Balancing ³	No	Yes	Yes	Yes

Advanced Device Security Policy

Advanced Device Security (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 the ADS policy, it applies to all F_Ports on the AG-enabled device. 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 the AG device 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 **agshow** command as follows. Note that the **agshow** command is not supported in AG mode.

- Issuing the **agshow** command without operands displays a summary of F_Ports on all AG Core and Edge switches attached to the fabric, such as the AG devices that are directly connected to fabric.
- Issuing the **agshow --name** command displays details, such as Access Gateway and attached F_Port information, for a specific AG Core or Edge switch attached to the fabric.
- Issuing the **agshow --all** command displays details, such as Access Gateway and attached F_Port information, for all AG Core and Edge switches connected to the fabric.

Alternatively, the security policy can be established in the Enterprise fabric using the Device Connection Control (DCC) policy. For information on configuring the DCC policy, refer to [Enabling the DCC Policy on a Trunk](#) on page 70. 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 the Failover and Failback policies are enabled.

Enabling and Disabling the ADS 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, 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.

² The ADS policy is not supported when using device mapping.

³ Device Load Balancing and Automatic Login Balancing cannot be enabled for the same port group.

Allow Lists

You can determine which devices are allowed to log in on a per-F_Port basis by specifying lists of F_Ports and device WWNs in the **ag --adsset** command. The ADS policy must be enabled for this command to succeed.

```
ag --adsset "F_Port [ ;F_Port2;... ]" "WWWN [ ;WWN2;... ]"
```

Lists must be enclosed in 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 login count.

Use an asterisk (*) instead of port numbers in the F_Port list to add the specified WWNs to all the F_Ports allow lists. Use an asterisk (*) instead of WWNs to indicate access to all devices from the specified F_Port list. A blank WWN list ("") indicates no access.

Use an asterisk enclosed in quotation marks ("*") to set the allow list to "all access." Use a pair of double quotation marks ("") to set the allow list to "no access."

Consider the following characteristics of the allow list:

- The maximum number of device entries allowed in the allow list is twice the per-port maximum 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.

Setting the List of Devices Allowed to Login

The following steps show how to set the allowed devices.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the **ag --adsset** command with the appropriate options to set the list of devices allowed to log in to 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 Login

The following steps show how to set the list of devices that are 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 options to set the list of devices not allowed to log in to 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

Remove specified WWNs from the list of devices allowed to log in to the specified F_Ports using the **ag --adsdel** command. Lists must be enclosed in 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.

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

For more details on this command and its operands, refer to the *Fabric OS Command Reference*.

The following steps show how to remove devices from a list of allowed devices. 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. Use the **ag --adsdel** command to remove one or more devices from the list of allowed devices.

In the following example, two devices are removed from the list of allowed devices (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

Add specified WWNs to the list of devices allowed to log in to the specified F_Ports using the **ag --adsadd** command. Lists must be enclosed in 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.

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

For more details on this command and its operands, refer to the *Fabric OS Command Reference*.

The following steps show how to add devices to a list of allowed devices. 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. Use the **ag --adsadd** command with the appropriate options to add one or more new devices to the list of allowed devices.

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

The following steps show how to display 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. Use the **ag --adsshow** command.

For each F_Port, command output will show access for all devices, a list of device WWNs, or no access. For more details on this command and its output, refer to the *Fabric OS Command Reference*.

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.
- The ADS policy can be enabled or disabled independent of the status of other AG policies.
- The ADS policy is not supported with device mapping.

Upgrade and Downgrade Considerations for the ADS Policy

Downgrading to Fabric OS v6.4.0 or earlier is supported.

Downgrading from Fabric OS v7.1.0 to v6.4.0 or upgrading from Fabric OS v6.4.0 to v7.1.0 will not change the ADS policy settings.

NOTE

The Brocade G610 is supported in Fabric OS 8.1.0 and subsequent releases. The Brocade G620 is supported in Fabric OS 8.0.0 and subsequent releases. Downgrading to prior releases is not supported.

Automatic Port Configuration Policy

The Automatic Port Configuration (APC) policy 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 the APC Policy

The following steps show how to enable the APC policy on a switch.

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 (PG) policy.
5. Enter the **ag --policyenable auto** command to enable the APC policy.
6. At the command prompt, type **Y** to enable the policy.

The switch is ready. A reboot is not required.

Disabling the APC Policy

The following steps show how to disable the APC policy on a switch.

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 auto** command to disable the APC policy.
5. At the command prompt, type **Y** to disable the policy.

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

APC Policy Considerations

Following are the considerations for the Automatic Port Configuration (APC) 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 the APC policy is enabled.
- The APC policy applies to all ports on the switch.
- Enabling the APC policy is disruptive and erases all existing port 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 mapping revert back to the default factory configurations for that platform. It is recommended that before you either disable or enable APC policy, you save the current configuration file using the **configUpload** command in case you need this configuration again.
- The APC policy is disabled by default on the switch.
- The APC policy is not supported in cascaded Access Gateway configuration, either on core or edge AG devices.

Upgrade and Downgrade Considerations for the APC Policy

You can downgrade to a Fabric OS level that supports the APC policy. You can upgrade from Fabric OS v6.4.0 to Fabric OS v7.1.0, and the policy that was enabled in Fabric OS v6.4.0 will be maintained. Upgrading to Fabric OS v7.1.0 from Fabric OS releases prior to v6.4.0 is not supported.

The Brocade G610 is supported in Fabric OS 8.1.0 and subsequent releases. Downgrading to a prior release is not supported.

The Brocade G620 is supported in Fabric OS 8.0.0 and subsequent releases. Downgrading to a prior release is not supported.

Port Grouping Policy

Use the Port Grouping (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.


The following figure 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.

FIGURE 10 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.

FIGURE 11 Port group 1 (PG1) setup



Adding an N_Port to a Port Group

The following steps show how to add 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 options to add an N_Port to a specific port group. In the following example, N_Port 14 is added to port group 3.

If you add more than one N_Port, 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 the N_Port should be remapped before the 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 options 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 **ag --pgshow** command to verify the N_Port was deleted from the specified port group.

Removing a Port Group

The following steps show how to remove 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 the appropriate options to remove a port group. In the following example, port group 3 is removed.

```
switch:admin> ag --pgremove 3
```

Renaming a Port Group

The following steps show how to rename 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 the appropriate options to rename a port group. In the following example, port group 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 pg** command to disable the Port Grouping policy.

Port Grouping Policy Modes

You can enable and disable the Automatic Login Balancing and Managed Fabric Name Monitoring (MFNM) Port Grouping policy modes when you create port groups using the **pgcreate** command. Alternately, you can enable these modes using the **ag --pgsetmodes** command.

Automatic Login Balancing Mode

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 port group are redistributed to maintain a balanced N_Port-to-F_Port ratio.

Consider the following notes about Automatic Login Balancing mode:

- Automatic Login Balancing mode is disruptive. However, you can minimize disruption by disabling or enabling rebalancing of F_Ports on F_Port offline events or N_Port online events.
- 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 mode, 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 specified 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 Automatic Login Balancing mode. This is not the case for port groups not enabled for Automatic Login Balancing mode. 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 Automatic Login Balancing mode.

Managed Fabric Name Monitoring Mode

When enabled, Managed Fabric Name Monitoring (MFNM) mode queries the fabric name at a specific time period. If it detects an inconsistency, for example, all the N_Ports within a port group are not physically connected to the same physical or virtual fabric, it generates a RASLOG message. In “default” mode, a message is logged into RASLOG. In “managed” mode, automatic failover is disabled for all N_Ports within the N_Port group, and a message is logged into RASLOG about multiple fabrics.

Enable or disable MFNM mode on a port group using the steps under [Enabling MFNM Mode](#) on page 54 and [Disabling MFNM Mode](#) on page 55. In both default and managed mode, the system queries the fabric name once every 120 seconds. You can configure the monitoring timeout value to something other than 120 seconds using the steps under [Setting the Current MFNM Mode Timeout Value](#) on page 55.

Creating a Port Group and Enabling Automatic Login Balancing Mode

The following steps show how to create a port group and enable Automatic Load 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 the appropriate options 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. A table containing a port group with ID 3 and ports 1 and 3 should display.

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 Automatic Login Balancing mode 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 the appropriate options 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 the appropriate options to disable automatic login distribution of N_Ports for all port groups 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 the appropriate options 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. For more details on this command and its output, refer to the *Fabric OS Command Reference*.

Considerations When Disabling Automatic Login Balancing Mode

Consider the following when disabling Automatic Login Balancing mode:

- Be aware that modifying Automatic Login Balancing mode default settings using the **agautomapbalance** command may yield uneven distribution of F_Ports to N_Ports. In such cases, you might 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 54.

Enabling MFNM Mode

The following steps show how to enable MFNM mode.

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

2. Enter the **ag --pgsetmodes** command with the appropriate options to enable MFNM mode.

This command changes the monitoring mode from "default" to "managed." 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 MFNM Mode

The following steps show how to disable MFNM mode.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the **ag --pgdelmodes** command with the appropriate options 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
```

3. Use the **ag --pgshow** command to display the port group configuration. If the port group disabled, "mfnm" should not display under PG_Mode for port 3.

For more details on this command and its operands, refer to the *Fabric OS Command Reference*.

Displaying the Current MFNM Mode Timeout Value

The following steps show how to display the current MFNM mode timeout value.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the **ag --pgfnmtov** command to display the current MFNM timeout value.

```
switch:admin> ag --pgfnmtov
Fabric Name Monitoring TOV: 120 seconds
```

Setting the Current MFNM Mode Timeout Value

The following steps show how to set MFNM mode the current MFNM mode timeout value.

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

The following example sets the timeout value to 100 seconds.

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

Port Grouping Policy Considerations

Policy considerations for the Port Grouping policy include the following:

- 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 later. 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 Automatic Login Balancing mode is enabled or disabled for the port group, the N_Port maintains its original failover or fallback 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. The F_Ports may lose connectivity to the targets to which they were connected before the failover, thus causing I/O disruption, as shown in [How Port Groups Work](#) on page 50. Ensure that the port group mode is set to MFNM mode (refer to [Enabling MFNM Mode](#) on page 54). 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 [Managed Fabric Name Monitoring Mode](#) on page 53.

Upgrade and Downgrade Considerations for the Port Grouping Policy

The Brocade G610 is supported in Fabric OS 8.1.0 and subsequent releases. Downgrading to a prior release is not supported.

The Brocade G620 is supported in Fabric OS 8.0.0 and subsequent releases. Downgrading to a prior release is not supported.

Downgrading to Fabric OS 6.4.0 or earlier is supported.

Note the following considerations when upgrading to Fabric OS 7.1.0:

- When upgrading to Fabric OS 7.1.0 from 6.4.0, the PG policy that was enforced in Fabric OS 6.4.0 continues to be enforced in Fabric OS 7.1.0 and the port groups are retained. You should save the configuration file using the **configUpload** command in case you might need this configuration again.
- Upgrading to Fabric OS 7.1.0 from Fabric OS prior to 6.4.0 is not supported.

Device Load Balancing Policy

When the Device Load Balancing policy is enabled, devices mapped to a port group always log in to 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 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, it is recommended 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 to N_Ports take precedence over automatically created mappings. Refer to [Mapping Priority](#) on page 39 for details on connection priority for AG port mapping. For more information on device mapping, refer to [Device Mapping](#) on page 33.

Enabling the Device Load Balancing Policy

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. Because Fibre Channel devices are identified by their WWNs, CLI commands use device WWNs.

Disabling the Device Load Balancing Policy

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 disable the Device Load Balancing policy.
3. Enter the **ag --policyshow** command to determine the current status of the Device Load Balancing policy.

Device Load Balancing Policy Considerations

The following considerations apply to Device Load Balancing.

- The Device Load Balancing policy should be enabled on the edge AG of a cascaded AG configuration.
- The Device Load Balancing policy is not applicable on a port group when the APC policy or Automatic Login Balancing are 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 the 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 the Device Load Balancing 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

The Persistent ALPA policy is meant for host systems with operating systems that cannot handle different PID addresses across login sessions when booting over a SAN. The Persistent ALPA policy for switches in Access Gateway mode allows you to 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. Because the Arbitrated Port Loop Address (ALPA) field makes up a portion of the PID, the PID may change across switch or server power cycles. This policy, if enabled, helps reduce the chances of a different PID being issued for the same host.

The benefit of this policy is that it ensures that a host has the same ALPA on the F_Ports through the host power cycle. You can also achieve the same behavior and benefit by setting the same policy in the main (core) fabric. When this policy 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. The following modes deal with this situation:

- In "Flexible" mode, the AG logs an event that it did not receive the same (requested) ALPA from the core fabric and brings up the device with the ALPA assigned by the fabric. An unassigned ALPA value is assigned when the ALPA assigned to the device is taken by another host.
- In the "Stringent" mode, if the requested ALPA is not available, the server login will be rejected and the server port cannot log in to the fabric if assignment of the same ALPA is not possible.

Enabling the Persistent ALPA Policy

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]
```

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 (-f) or stringent (-s) mode. The following example shows enabling the policy in flexible mode.

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

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

Disabling the Persistent ALPA Policy

When you disable this policy, 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 the instructions in [Removing Device Data from the Database](#) on page 58.

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 --deletepwnfromdb** command.

```
switch:admin> ag --deletepwnfromdb PWN
```

In the example, PWN is the port that you want to remove from the database.

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 option 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 option 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 device 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 policy is not supported in the following situations:

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

Failover Policy

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

NOTE

For port mapping, the Failover policy must be enabled on an N_Port for failover to occur. For device 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 Mapping

The AG device provides an option to specify a secondary failover N_Port for an F_Port. The Failover policy 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.

NOTE

If failover and fallback policy are disabled, an F_Port mapped to an N_Port will go offline when the N_Port goes offline and it will come online when the N_Port comes online.

Failover Configuration 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 the Port Grouping 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.

Failover Example

The following example shows the failover sequence of events 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 Example 1 in the following figure, 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 Example 2 in the following figure, 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. The following figure shows this example.

FIGURE 12 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, 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.

The following steps show how to add 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 --prefset** command with the "F_Port1[; F_Port2, ...]"N_Port options 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, as shown in the following 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 mode is enabled.

Deleting F_Ports from a Preferred Secondary N_Port

The following steps show how to delete 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 options 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 Mapping

Failover is handled similarly for port mapping and device 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 the Failover or Failback policies for N_Ports has no effect on device mapping. A device will always fail over to an online N_Port in the port group, regardless of whether the Failback policy is enabled for an N_Port or not. Whereas, with port mapping, if you disable the Failover or Failback policy on an N_Port, the F_Port will not fail over or fail back to other N_Ports.

Failover behavior is different if a device is mapped to a specific N_Port instead of to an 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, use the **ag --addwwnfailovermapping** command with the *N_Port* and **--all** options.

```
ag --addwwnfailovermapping 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, use 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, use 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 the Failover Policy on an N_Port

Use the following steps to enable or disable the 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 the Failover Policy for a Port Group

The Failover policy can be enabled on a port group. Use the following steps to enable or disable the 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 the Failover Policy

The Brocade G610 is supported in Fabric OS 8.1.0 and subsequent releases. Downgrading to a prior release is not supported.

The Brocade G620 is supported in Fabric OS 8.0.0 and subsequent releases. Downgrading to a prior release is not supported.

Consider the following when upgrading or downgrading Fabric OS versions:

- Downgrading to Fabric OS v6.4.0 or earlier is supported.
- Upgrading from Fabric OS v6.4.0 to v7.1.0 or downgrading from Fabric OS v7.1.0 to v6.4.0 will not change failover settings.

Fallback Policy

The Fallback 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. The Fallback 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 mapping, the Fallback policy must be enabled on an N_Port for failback to occur. For device mapping, the Fallback 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.

NOTE

If failover and failback policy are disabled, an F_Port mapped to an N_Port will go offline when the N_Port goes offline and it will come online when the N_Port comes online.

Fallback Policy Configurations in Access Gateway

The following sequence describes how a failback event occurs:

- When an N_Port comes back online, with the Fallback policy 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 fallback period is quite fast and rarely causes an I/O error at the application level.

Fallback Example

In the following figure, the Access Gateway N_1 remains disabled because the corresponding F_A1 port is offline. However, N_2 comes back online. Refer to [Failover Example](#) on page 60 for the original failover scenario.

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.

FIGURE 13 Fallback behavior



Enabling and Disabling the Fallback Policy on an N_Port

Use the following steps to enable or disable the Fallback policy 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 fallback setting.

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

3. Use the following commands to enable or disable the Fallback policy:

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

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

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

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

Enabling and Disabling the Fallback Policy for a Port Group

Use the following steps to enable or disable the Fallback 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 the Fallback policy for a port group:
 - Enter the **ag --failbackenable pg pgid** command to enable fallback on a port group.

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

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

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

Forcing Fallback on N_Ports

Forcing F_Ports to fail back and log in through their configured N_Port can help you recover from unexpected imbalances of F_Ports using N_Ports. Using the **ag --failbackforce** command, you can force fallback from a specific N_Port, all N_Ports in a port group, or all N_Ports on the switch, regardless of the set port group policy or port groups. The fallback action is immediate.

Consider the following restrictions and exceptions for this action:

- Only those F_Ports that are currently not using their configured N_Port are failed back.
- Fallback will not occur if the N_Port is in a port group with Login Balancing enabled as it will automatically map F_Ports to the least-utilized N_Ports.
- Fallback will not occur if fallback is disabled on the N_Port or if the F_Ports have been manually disabled.

Use the following steps to force fallback:

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

2. Use one of the following commands:

- Enter **ag --fallbackforce N_Port** to force failback for F_Ports configured on N_Port 16.

```
switch:admin> ag --fallbackforce 16
Fallback action has been forced.
```

- Enter **ag --fallbackforce -pg pgid** to force all F_Ports configured on all N_Ports in port group 0 to fail back to their respective N_Ports.

```
switch:admin> ag --fallbackforce -pg 0
Fallback action has been forced.
```

- Enter **ag --fallbackforce -all** to force all F_Ports on the switch to fail back to their configured N_Port.

```
switch:admin> ag --fallbackforce -all
Fallback action has been forced.
```

Upgrade and Downgrade Considerations for the Fallback Policy

- Downgrading to Fabric OS 6.3.0 or earlier is supported.
- Upgrading from Fabric OS 6.4.0 is supported.

The Brocade G610 is supported in Fabric OS 8.1.0 and subsequent releases. Downgrading to a prior release is not supported.

The Brocade G620 is supported in Fabric OS 8.0.0 and subsequent releases. Downgrading to a prior release is not supported.

Fallback Policy Disabled on Unreliable Links (N_Port Monitoring)

Links from all N_Ports are monitored for the number of online and offline static change notifications (SCNs) that occur during a set time period (5 minutes). If the number of SCNs on a link exceeds a set threshold, the link is considered unreliable, and failback is disabled for that N_Port. Failover continues for the port as needed. Once the number of SCNs drops below the set threshold, the port is deemed reliable again and failback is re-enabled. If the link from a preferred secondary N_Port for an F_Port becomes unreliable, failover will not occur to that N_Port.

The default threshold is 25 SCNs per 5 minutes. You can modify the SCN threshold counter using the following command.

```
ag --reliabilitycounterset "count"
```

Specify 0 for the count as shown in the following command to disable reliability. When disabled, any reliable N_Ports are recovered to be reliable again.

```
ag --reliabilitycounterset 0
```

You can view counter settings using the following command.

```
ag --reliabilitycountershows
```

You can display the current reliable state of N_Ports using the **ag --reliabilityshow N_Port**.

```
ag --reliabilityshow 17
N_Port 17 is Unreliable
```

If an N_Port is not specified (**ag --reliabilityshow**), reliability is displayed for all N_Ports, regardless of port group. The output flags reliable ports with 1 and unreliable ports with 0.

Considerations for Fallback Policy Disabled on Unreliable Links

Consider the following when an N_Port link becomes reliable again after being unreliable:

- Preferred N_Port settings are enforced.
- If fallback is enabled, configured F_Ports will fail back to the N_Port.
- If the configured F_Ports are offline, they will go back online.
- If Device Load Balancing is enabled, rebalancing occurs.

Trunking in Access Gateway Mode

The 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 fallback 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.

ch. To enable this feature, you must install the Brocade ISL Trunking 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 a 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.

NOTE

You can also enable N_Ports on the Access Gateway switch for trunking using the **portcfgtrunkport** command. However, do not create an F_port trunk if the connecting Edge switch has no trunking capability (no license or license disabled). Doing so can cause the Edge switch ports to not link up.

NOTE

N_Port trunking is only supported for HBAs connected to switches running in Native mode. It is not supported for HBAs connected to switches running in Access Gateway mode.

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. With trunking, all but one link within a trunk group can go offline or become disabled, but the trunk remains fully functional and no reconfiguration is required. Trunking prevents reassignments of the port ID when N_Ports go offline.

Configuring Trunking on the Edge Switch

Because AG trunking configuration is mostly on the Edge switch, information in this section is applicable to the Edge switch and not the AG device. On the AG device, 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.

In the **nsshow** command, end devices that log in to the edge fabric through an F_Port trunk will show the F_Port trunk speed and not necessarily the actual device speed. For example, when a host using an 8-Gbps interface is attached to the AG device and logs in to the fabric through a 32-Gbps F_Port trunk, the device link speed will show 32 Gbps (the trunk speed) and not 8 Gbps (the device speed).

NOTE

Do not create an F-port trunk if the connecting Edge switch has no trunking capability (no license or trunk port configuration disabled). Doing so can cause the Edge switch ports to not link up.

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 port IDs (PIDs). 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 device. Verify that this feature has not been disabled on N_Ports that are part of a 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 trunking groups are based on the user port number, with eight contiguous 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

Use the following steps to set up trunking.

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 [slot /]port, 1** command.
4. Ensure that the ports within a trunk have the same speed.
5. Ensure that the 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 Area

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.

Trunk Area Assignment Example

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 (TA) 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 **porttrunkarea --help** or refer to the *Fabric OS Command Reference*.

TABLE 11 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							

The following steps show how to assign a trunk area.

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 a TA for the appropriate ports. The following example enables a TA for ports 13 and 14 on slot 10 with port index of 125.

```
switch:admin> porttrunkarea --enable 10/13-14 -index 125
```

4. Display the TA port configuration (ports still disabled) using the **porttrunkarea --show enabled** command.
5. Enable the ports specified in step 3 using the **portenable** command.

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

6. Show the TA port configuration after enabling the ports using the **porttrunkarea --show enabled** command. The ports that you enabled should appear in the output.

Enabling the DCC Policy on a Trunk

After you assign a Trunk Area, the **porttrunkarea** command checks whether there are any active Device Connection Control (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.

Use the following steps to enable the DCC policy on a trunk.

1. Add the WWN of all the devices to the DCC policy against the TA.
2. 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.

3. 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

The following steps show how to enable 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 the appropriate options to form a trunk group for the desired ports. For example, if ports 36-39 were disabled in step 2, then the following example command forms a trunk group for ports 36-39 with index 37. These will be connected to N_Ports on an AG device.

```
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** command.

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

If an error occurs as in the previous example, disable each port using the **portdisable port** command, and then reissue the command.

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

Monitoring Trunking

For F_Port masterless trunking, you must install monitors on the F_Port trunk port. If the master port changes, you need not move the monitors to the new master port. For example, if a master port goes down, a new master port is selected from the remaining slave ports and becomes active. The new master port monitoring starts automatically because the existing F_Port trunk ports have monitors installed. When you add a monitor to a slave port, it is automatically added to the master port.

For more information about installing monitors on F_Port trunks, refer to the *Brocade Flow Vision Administrator's Guide* publication.

AG Trunking Considerations for the Edge Switch

Refer to the following table for information about trunking considerations for the Edge switch.

TABLE 12 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 an Edge switch or blade. The static trunk area you assign must be one of the port's default areas of the trunk group.
Authentication	<p>Authentication occurs only on the F_Port trunk master port and only once per the entire trunk. This behavior is the 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.</p> <p>NOTE Authentication is also supported on switches configured in AG mode.</p>
Management Server	Registered Node ID (RNID), Link Incident Record Registration (LIRR), and Query Security Attributes (QSA) Extended Link Service Requests (ELSS) are not supported on F_Port trunks.
Trunk area	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. You cannot assign a Trunk Area to ports if the standby CP is running a firmware version earlier than Fabric OS 6.2.0.
PWWN	The entire Trunk Area trunk group shares 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 13 .
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 6.1.0.</p>
Upgrade	No limitations on upgrade to Fabric OS 7.1.0 and subsequent releases 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.
portCfgTrunkPort [slot/] port , 0	portCfgTrunkPort [slot/] port , 0 will fail if a Trunk Area is enabled on a port. The port must be Trunk Area-disabled first.
switchCfgTrunk 0	switchCfgTrunk 0 will fail if a port has TA enabled. All ports on a switch must be TA disabled first.

TABLE 12 Access Gateway trunking considerations for the Edge switch (continued)

Category	Description
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 is permitted 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 fast-write is enabled on a port, 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.
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 Fast Write.</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 request 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.
D,I Zoning (D,I) AD (D, I) DCC and (PWWN, I) DCC	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".

TABLE 12 Access Gateway trunking considerations for the Edge switch (continued)

Category	Description
	<p>This means that Domain,Index (D,I), which refers 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.

Refer to the following table for information about PWWN format for F_Port and N_Port trunk ports.

TABLE 13 PWWN format for F_Port and N_Port trunk ports

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

Trunking Considerations for Access Gateway Mode

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 the 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 to Fabric OS v7.1.0 and downgrading to Fabric OS v6.4.0 and earlier is supported.

The Brocade G610 is supported in Fabric OS v8.1.0 and subsequent releases. Downgrading to a prior release is not supported.

The Brocade G620 is supported in Fabric OS v8.0.0 and subsequent releases. Downgrading to a prior release is not supported.

Adaptive Networking on Access Gateway

Adaptive Networking (AN) services ensure 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.

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

QoS: Ingress Rate Limiting

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 "Ingress Limiting" in the *Fabric OS Administration 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 Administration Guide*.

The following figure shows the starting point for QoS in various Brocade and non-Brocade configurations.

FIGURE 14 Starting point for QoS



Upgrade and Downgrade Considerations for Adaptive Networking in AG Mode

The Brocade G610 is supported in Fabric OS v8.1.0 and subsequent releases. Downgrading to a prior release is not supported.

The Brocade G620 is supported in Fabric OS v8.0.0 and subsequent releases. Downgrading to a prior release is not supported.

Upgrading to Fabric OS v7.1.0 from Fabric OS v6.4.0 is supported. Note the following considerations when upgrading to Fabric OS v7.1.0 from Fabric OS v6.2.X and earlier and downgrading from Fabric OS v7.1.0 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 v7.1.0 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 ISLs.

Adaptive Networking on Access Gateway Considerations

- QoS is configured in the fabric, as normal, and not on the AG device.
- QoS on Access Gateway is only supported on Fabric OS 6.3 and later.
- You should disable HBA QoS if connecting to an AG device running Fabric OS 6.2 or prior release.

- 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, and 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
- For QoS to function in AG mode, the QoS mode state must either be ON (enabled) or AE (auto-enable) for both F_Ports and N_Ports. If the QoS state is OFF, the ports will only carry traffic with medium VC priority.
- Ingress rate limiting is not enforced on trunked ports.

Per-Port NPIV Login Limit

The Per-Port NPIV login limit feature allows you to set a specific maximum NPIV login limit on individual ports. This feature works in both Native and Access Gateway modes. 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:

- 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. The F_Port must then be re-mapped to its N_Port after using **portCfgDefault**.
- 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 both Fabric OS Native mode and Access Gateway mode. For details, refer to the *Fabric OS Command Reference*.

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 **portcfgnpiport --setloginlimit [Slot/]Port loginlimit [Slot/] Port loginlimit** command to set the login limit. For example, the following example sets the login limit on port 12 to 200.

```
portcfgnpiport --setloginlimit 12 200
```

Duplicate PWWN Handling during Device Login

Handling of logins from two devices using the same PWWN follows standards for all Fabric OS switches. Having two devices with the same PWWN logged into the fabric at the same time may not be desirable, as there have been cases when ports coming online get stuck in G_Port state in the AG switch. You can configure three policies:

- Default: The Access Gateway switch will not accept a second login request from a PWWN when a device with the same PWWN has already logged into the fabric. The second request with similar PWWN will be rejected and the port will be persistently disabled with reason as “Duplicate Port WWN detected.”
- Enforced login: The second login request will have precedence over the existing login and Access Gateway will accept the login.
- Mixed: This option takes port type into consideration. The second login request will have precedence over the existing login in case of a duplicate entry exit on the F_Port with an NPIV device logged in. The first login takes precedence over the second login request in case of a duplicate entry exit on the F_Port without any NPIV device logged in.

You can configure different handling of duplicate PWWNs other than the default operation using the `configure` command through the F_Port login parameters. For more information on configuration and handling of duplicate PWWNs during device login, refer to the “Duplicate PWWN handling during device login” section in the *Fabric OS Administration Guide*.

This feature is supported in the following configurations:

- AG switch connected to AG switch in cascaded configuration.
- AG switch connected to Brocade fabric switch.
- AG switch connected to a Host Bus Adapter (HBA).

Performance Monitoring

Performance monitoring is available through the following Flow Vision features:

- Flow Monitor feature for platforms using Fabric OS 7.2 and later.
- Flow Mirror feature for platforms using Fabric OS 8.1.0 and later.

In Fabric OS 7.4.0 and later, legacy Advanced Performance Monitoring features are no longer supported.

Flow Monitor

Flow Vision flow monitoring is a licensed feature supported on platforms using Fabric OS 7.2.0 and later. It provides a unified platform to manage traffic-related applications on Fabric OS devices. Storage administrators can use this platform to simulate, monitor, and capture the network's traffic patterns and to make capacity-planning decisions based on the collected statistical data.

To access Flow Vision, the Fabric Vision (FV) license must be installed on the switch.

The Flow Monitor component of Flow Vision is supported by Access Gateway. Flow Monitor allows you to monitor the network's traffic pattern and provides statistics to make capacity planning decisions based on the collected data. Flow Monitor provides a single interface to manage flows and unifies different performance monitoring features such as end-to-end monitors and frame monitors.

Flow Monitor expands on basic performance monitoring by allowing you to monitor any hardware-supported flow parameters and define your own flows using combinations of source and destination devices, source and destination IDs, LUN IDs, CSCTL values, and frame types as parameters.

Following are examples of monitors that you can replicate using Flow Monitor:

- End to End Monitor: This measures the traffic in terms of word count between a pair of ports (host and target). Use this to view end-to end traffic values on an Access Gateway.

- Frame monitoring: Set up an SCSI monitor so that you can view the **-frametype** frame and byte count for a flow.
- LUN monitoring: This allows you to view the **-frametype** frame and byte count for a specific LUN.
- Flow learning monitor: This learns flows going through a specific F_Port in the Access Gateway.
- DST monitor: This provides the legacy support equivalent to a frame monitor.

Access Gateway switches support flow monitors on F_Ports only, and only the ingress port parameter is supported. For more information on using Flow Monitor features for Access Gateway in Flow Vision, refer to the *Flow Vision Administrator's Guide*.

Flow Mirror

Flow Mirror is supported in Access Gateway mode on Gen 5 and Gen 6 platforms using Fabric OS 8.1.0 and later. It is supported for N_Ports and F_Ports as ingress or egress ports in the flow definition.

Flow Mirror, is a Flow Vision feature that provides a non-disruptive diagnostic tool for analyzing fabric performance issues. Using Flow Mirror, you can perform the following tasks:

- Non-disruptively create copies of application flows that can optionally be captured for deeper analysis.
- Define a traffic pattern and create a real-time copy of this traffic, allowing you to analyze a live system without disturbing existing connections. You can also use this feature to view traffic passing through a port.
- Use a MAPS logical port group for either an ingress or egress port. Note that logical port groups are not supported on CPU flow mirroring (CFM) or logical flow mirroring (LFM).
- Use a predefined flow with the flow --show command to mirror a SCSI command, first response, status and ABTS frames from all the Gen 5 and Gen 6 F_Ports on a switch.


Flow Mirror duplicates the specified frames in a user-defined flow and sends them to a sink. This sink could be either:

- The local switch control central processor unit (CPU); this form is called "CPU flow mirroring" (CFM), and has a limit of 256 frames per second.
- An analyzer/packet sniffer connected through a port in the metaSAN. The bandwidth limit for a flow using this type of link is the bandwidth of the mirror destination port. This form is called "Local flow mirroring" (LFM), and mirrors the flow to a port on the same physical switch. This requires that a loopback SFP be plugged in at the other side of the analyzer.

LFM and CFM flows are supported in AG mode. Remote Flow Mirroring (RFM) is not supported. Since all forms of mirroring are mutually exclusive, only one form of mirroring (LFM or CFM) can be active at a time.

The following illustration illustrates a typical application of Flow Mirror on an Access Gateway platform:

FIGURE 15 Flow Mirror application on Access Gateway platform



Following are considerations and limitations for using Flow Mirror for Access Gateway:

- A port should be in disabled state before configuring it as mirror port. The mirror port configuration will be blocked if port is enabled.
- You do not need to clear the N_Port to F_Port mappings to configure a mirror port.
- The mirror port can be specified by its port or index number as both are identical in AG mode. Display ports and their index numbers using the **switchShow** command.
- AG mode enable and disable will be blocked if flows are configured on the switch.
- You can only configure external ports as mirror ports. You cannot configure internal ports on SAN I/O modules as mirror ports.
- If all ports that are part of the flow definition are online, Flow Mirror will be enforced.
- If an offline port is part of the flow definition, Flow Mirror will be activated, but not enforced.
- For firmware upgrades and downgrades:
 - If a Flow Mirror flow exists, the downgrade will be blocked. You must delete all the Flow Mirror flows before proceeding with the downgrade.
 - If a mirror port configuration exists, the downgrade shall be blocked. You must clear all mirror port configurations before proceeding with the downgrade.
- During a high availability failover, all flows will be deleted and then reinitiated after failover. If any Flow Mirror flow is active, then mirroring shall resume after the failover.

For more information on using and configuring Flow Mirror features, refer to the *Flow Vision Administrator's Guide*.

Considerations for the Brocade Switches

The Brocade 6505, Brocade 6510, Brocade G610, and Brocade G620 can function in either Fabric OS Native mode or Brocade Access Gateway mode. These switches are shipped in Fabric OS Native mode. They are also supported in Access Gateway cascaded configurations. All PoD licenses must be present to support Access Gateway for all releases prior to Fabric OS 7.3.0. However, starting with Fabric OS 7.3.0, all PoD licenses are not required.

NOTE

The Brocade G620 is supported in Fabric OS 8.0.1 and later. The Brocade G610 is supported in Fabric OS 8.1.0 and later.

SAN Configuration with Access Gateway

• Connectivity of Multiple Devices Overview.....	82
• Direct Target Attachment.....	82
• Target Aggregation.....	84
• Access Gateway Cascading.....	85
• Fabric and Edge Switch Configuration.....	87
• Connectivity to Cisco Fabrics.....	88
• Rejoining Fabric OS Switches to a Fabric.....	89

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. Switches in AG mode can connect to third-party fabrics with the following firmware versions:

- Fabrics operating with M-EOSc v9.6.2 or later and M-EOSn v9.6 or later.
- Cisco fabrics that support standards-based NPIV.

Considerations for Connecting Multiple Devices


Consider the following points when connecting multiple devices to a switch in AG mode:

- AG does not support daisy chaining when two AG devices are connected to each other in a loop configuration.
- 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, as illustrated in the following figure.

FIGURE 16 Direct target attachment to switch operating in AG mode



Although target devices can be connected directly to AG ports, it is recommended that the switch operating in AG mode be connected to the core fabric.

Considerations for Direct Target Attachment

Consider the following points for direct target attachment:


- 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 must be established by the core fabric.
- A recommended practice is that you do not map hosts and targets to the same N_Port.

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

Target Aggregation

Access Gateway mode is normally used as host aggregation. In other words, a switch 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, as shown in the following figure. Target aggregation 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.

FIGURE 17 Target aggregation




Access Gateway Cascading

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

Access Gateway cascading allows you to link two Access Gateway (AG) switches 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 18 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.

Access Gateway Cascading Considerations

Note the following configuration considerations when cascading Access Gateways:

- Only one level of cascading is supported. Note that several Edge AG switches can connect into a single Core AG switch to support an even 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 Device Security (ADS) policy 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.
- You can view the devices with active connections to an F_Port using the **agshow** command as follows. Note that the **agshow** command is not supported in AG mode.
 - Issuing the **agshow** command without operands displays a summary of F_Ports on all AG Core and Edge switches attached to the fabric, such as the AG devices that are directly connected to fabric.
 - Issuing the **agshow --name** command displays details, such as Access Gateway and attached F_Port information, for a specific AG Core and Edge switch attached to the fabric.
 - Issuing the **agshow --all** command displays details, such as Access Gateway and attached F_Port information, for all AG Core and Edge switches connected to the fabric.
- Due to high subscription ratios that could occur when cascading AG switches, ensure there is enough bandwidth for all servers when creating such configurations. The subscription ratio becomes more acute in a virtual environment.

- In Fabric OS 7.3.0 and later, the registration and de-registration of FDMI devices connected to an AG switch or cascaded AG switch is supported, and the **fdmishow** command on AG will display the local FDMI devices connected to the AG. However, remote FDMI devices will not be displayed.

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 in the switch's Access Control List (ACL). 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

The following steps show how to verify 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
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
-----
```

Refer to [Table 6](#) on page 25 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

The following steps show how to enable NPIV on M-EOS switches.

1. Connect to the switch and log in as admin on the M-EOS switch.
2. Enable Open Systems Management Server (OSMS) services by using the following command.
For the Mi10K switch, enter the following command.

```
fc osmsState vfid state
```

In the command, *vfid* is the virtual fabric identification number. The *state* variable can be **enable** for the enabled state or **disable** for the disabled state.

The *osmsState* variable 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. Use 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 use 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, you need to make sure that NPIV is enabled on the connecting switch and that Fabric OS 3.1.0 or later is used.

Enabling NPIV on a Cisco Switch

The following steps show how to enable NPIV on a Cisco switch.

1. Log in as admin on the Cisco MDS switch.
2. Enter the **show version** command to determine if you are using the correct SAN operating system version and if NPIV is enabled on the switch.
3. Enter the **configure terminal** command to begin configuration mode on the MDS switch.
4. Enter the **npiw enable** command to enable NPIV.
5. Press **Ctrl-Z** to exit configuration mode.
6. Enter the **copy** command to save the configuration.
7. Enter the **run** command to load the configuration.
8. Enter the **start** command to restart the switch with the new configuration.

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 rejoin a switch to the fabric:

- If you saved a Fabric OS configuration before switch reboot, download the configuration using the **configDownload** command.
- If you want to rejoin the switch to the fabric using the fabric configuration, use the following procedure.

To rejoin the Fabric OS switch to a fabric, perform the following steps:

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 Default zone changes.
5. Enter the **switchEnable** command to enable the switch and allow it to merge with the fabric.

The switch automatically rejoins the fabric.

Reverting to a Previous Configuration

The following steps show how to revert 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.

The switch automatically joins the fabric.

Troubleshooting

The following table provides troubleshooting information for Fabric OS switches in AG mode.

TABLE 14 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 option 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 option.</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 option.</p>
LUNs are not visible	Zoning on fabric switch is incorrect. Port mapping on Access Gateway mode switch is incorrect. Cabling not properly connected.	<p>Verify zoning on the Edge switch.</p> <p>Verify that F_Ports are mapped to an online N_Port. Refer to Table 8 on page 28.</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.</p> <p>Ensure the F_Port on AG module is enabled and active.</p>
Failover is not working	Failover disabled on N_Port.	<p>Verify that the failover and fallback policies are enabled, as follows:</p> <p>Enter the ag --failoverShow command with the port_number option.</p> <p>Enter the ag --fallbackShow command with the port_number option.</p>

TABLE 14 Troubleshooting (continued)

Problem	Cause	Solution
		<p>Command returns "Fallback (or Failover) on N_Port <i>port_number</i> is supported."</p> <p>If it returns, "Fallback (or Failover) on N_Port <i>port_number</i> is not supported." Refer to Adding a Preferred Secondary N_Port (Optional) on page 61.</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 8470 in AG mode.	The CNA host is retrying a login before the switch has finished processing a previous fabric logout (LOGO) attempt.	Messages display as designed. After the switch has completed LOGO processing, it will accept another login.

Revision History

FOS-821-access-AG101; September 28, 2018

Updated the [Slow-Drain Device Quarantining Support](#) on page 19.

FOS-821-access-AG100; August 28, 2018

- Modified the table in [Fabric OS Features in Access Gateway Mode](#) on page 11 to support Slow-Drain Device Quarantining (SDDQ).
- Added the section [Slow-Drain Device Quarantining Support](#) on page 19.