Brocade® Fabric OS® Extension User Guide, 9.2.x

User Guide
April 28, 2023

Table of Contents

Introduction

This document describes the Fabric OS® Extension concepts and features, the supported platforms, the required configuration steps, and some troubleshooting tools.

This document also provides information on a CLI-based menu-driven interactive interface to simplify the configuration of the Brocade® 7810 Extension Switch, Brocade 7850 Extension Switch, and the Brocade SX6 Extension Blade. This interface is intended for those who need a simple configuration to build a tunnel and circuits to run traffic from the Fibre Channel side to the IP side.

Supported Hardware and Software

The following hardware platforms are supported by Brocade Fabric OS 9.2.x.

Brocade Gen 7 (64G) Fixed-Port Switches

- Brocade G720 Switch
- Brocade G730 Switch
- Brocade 7850 Extension Switch

Brocade Gen 7 (64G) Directors

- Brocade X7-4 Director
- Brocade X7-8 Director

Brocade Gen 6 (32G) Fixed-Port Switches

- Brocade G610 Switch
- Brocade G620 Switch
- Brocade G630 Switch
- Brocade 7810 Extension Switch
- Brocade G648 Blade Server SAN I/O Module
- Brocade MXG610 Blade Server SAN I/O Module

Brocade Gen 6 (32G) Directors

- Brocade X6-4 Director
- Brocade X6-8 Director

Contacting Technical Support for Your Brocade® Product

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

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

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

Online	Telephone
For nonurgent issues, the preferred method is to log on to the Support portal at support.broadcom.com. (You must initially register to gain access to the Support portal.) Once registered, log on and then select **Brocade Products**. You can now navigate to the following sites: • **Case Management** • **Software Downloads** • **Licensing** • **SAN Reports** • **Brocade Support Link** • **Training & Education**	For Severity 1 (critical) issues, call Brocade Fibre Channel Networking Global Support at one of the phone numbers listed at www.broadcom.com/support/fibre-channel-networking/contact-brocade-support.

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.

Brocade Extension Platforms

This section provides an overview of the Brocade Extension platforms and features.

The Brocade Extension features in the Fabric OS software are supported by the following platforms:

- Brocade 7810 Extension Switch
- Brocade 7850 Extension Switch
- Brocade SX6 Extension Blade (in the following chassis):
  - Brocade X7-4 Director
  - Brocade X7-8 Director
  - Brocade X6-4 Director
  - Brocade X6-8 Director

Some features might require additional Fabric OS licenses to operate. For information about available Fabric OS licenses, refer to the *Brocade Fabric OS Software Licensing User Guide*.

The Brocade 7810 Extension Switch, Brocade 7850 Extension Switch, and Brocade SX6 Extension Blade tunnels cannot connect to a Brocade 7840 Extension Switch. Usually, an extension tunnel requires the same version of Fabric OS on each end of the tunnel.

For more information about the tunnel interoperability requirements across extension platforms, refer to the *Brocade Fabric OS Release Notes*.

NOTE
Extension tunnels between the current extension platforms (Brocade 7810 Extension Switch, Brocade 7850 Extension Switch, and Brocade SX6 Extension Blade) and previous generation extension platforms (Brocade 7500 Extension Switch, Brocade 7800 Extension Switch, Brocade 7840 Extension Switch, Brocade FR4-18i Extension Blade, and Brocade FX8-24 Extension Blade) are not supported.

The following table provides details about the Brocade Extension platform features and its support in the Brocade Extension platforms:

Table 1: Brocade Extension Features

Capability	Brocade 7810 Extension Switch	Brocade 7850 Extension Switch	Brocade SX6 Extension Blade
Number of VE_Ports	4	6VE mode: 3 per data processor (DP) 18VE mode: 9 per DP	10VE mode: 5 per DP 20VE mode: 10 per DP
Max. DP wide area network (WAN) side speed for all virtual Ethernet (VE) ports in DP	2.5Gbps	50Gbps per DP	20Gbps per DP
Max. circuit speed (GbE interface allowing)	Base: 1Gbps Full: 2.5Gbps	25Gbps	10Gbps
Min. circuit speed	20Mb/s	50Mb/s	20Mb/s
Max. number of circuits per VE_Port	Base: 1 Full: 6	10	10
Max. number of circuits per DP	24	36	40
Max. IP addresses per DP	60	60	60

Capability	Brocade 7810 Extension Switch	Brocade 7850 Extension Switch	Brocade SX6 Extension Blade
Max. IP routes per DP	120	120	120
Max. IP routes per GbE interface	120	128	128
Max. DP FC/Fiber Connection (FICON) side speed Depends on data compressibility; otherwise, the max. speed equals the max. WAN side speed (1:1)	10Gbps with 4:1 compression	VE6 mode: 80Gbps with 2:1 fast-deflate VE18 mode: 32Gbps per port group with 2:1 fast-deflate	VE10 mode: FCIP only: 40Gbps with 2:1 fast-deflate Hybrid: 20Gbps with 2:1 fast-deflate VE20 mode: FCIP only: 20Gbps with 2:1 fast-deflate
Max. Local Area Network (LAN) side speed per DP	10Gbps with 4:1 compression	40Gbps with no compression	20Gbps with no compression
Brocade Extension Trunking (BET) support	Yes	Yes	Yes
Adaptive Rate Limiting (ARL) and Committed Information Rate (CIR) support	Yes	Yes	Yes
Fibre Channel over Internet Protocol (FCIP) support	Yes	Yes	Yes
IP Extension support	Yes	Yes	Yes
Synchronous Remote Data Replication (RDR) support	Yes	Yes	Yes
Number of GbE interfaces	6	16	16
Number of 10GbE interfaces	6	16	16
Number of 25GbE interfaces	No	16	No
Number of 40GbE interfaces	No	No	Yes
Number of 100GbE interfaces	No	Yes	No
Virtual LAN (VLAN) Tagging (802.1Q)	WAN side: Yes LAN Side: Yes	WAN side: Yes LAN Side: Yes	WAN side: Yes LAN Side: Yes
IPv4 to IPv4 support	Yes	Yes	Yes
IPv6 to IPv6 support	Yes	Yes	Yes
FC Ports (supported speeds)	12 (4, 8, 16, 32)	24 (8, 16, 32, 64)	16 (4, 8, 16, 32)
Compression	Deflate Aggressive deflate	Fast deflate Deflate Aggressive deflate	Fast deflate Deflate Aggressive deflate
Protocol Acceleration	FastWrite OSTP read/write	FastWrite OSTP read/write Advanced FICON Accelerator	FastWrite OSTP read/write Advanced FICON Accelerator
Quality of Service (QoS) Marking	DSCP L2Cos (802.1P)	DSCP L2Cos (802.1P)	DSCP L2Cos (802.1P)

Capability	Brocade 7810 Extension Switch	Brocade 7850 Extension Switch	Brocade SX6 Extension Blade
FICON Acceleration	N/A	IBM z/OS Global Mirror Tape read/write Teradata emulation	IBM z/OS Global Mirror Tape read/write Teradata emulation
IP Security (IPsec) Encryption Suite-B and CNSA compliant	ESP IKEv2 AES-256-GCM SHA-512 HMAC	ESP IKEv2 AES-256-GCM SHA-512 HMAC	ESP IKEv2 AES-256-GCM SHA-512 HMAC
Jumbo Frames (WAN side)	Minimum 1260 bytes Maximum 9216 bytes	Minimum 1260 bytes Maximum 9216 bytes	Minimum 1260 bytes Maximum 9216 bytes
Jumbo Frames (LAN side)	Minimum 1260 bytes Maximum 9216 bytes	Minimum 1260 bytes Maximum 9216 bytes	Minimum 1260 bytes Maximum 9216 bytes
Path MTU (PMTU)	Maximum is 9100 bytes	Maximum is 9100 bytes	Maximum is 9100 bytes
Link Layer Discovery Protocol (LLDP) LAN and WAN sides support	Yes	Yes	Yes
Number of IP Extension Link Aggregation Group (LAGs)	2	4	4
Number of Max. LANs supported in LAGs	2	4	4
Number of IP Extension Link Aggregation Control Protocol (LACP)	Yes	Yes	Yes
Number of IP Extension Max. Transmission Control Protocol (TCP) Sessions	256 per DP	512 per DP	512 per DP
Number of IP Extension Gateways	4 per DP	8 per DP	8 per DP
Extension Hot Code Load (eHCL) support	No	Yes	Yes
Circuit Failover/Failback support	Yes	Yes	Yes
Circuit Failover Groups support	Yes	Yes	Yes
Spillover support	Yes	Yes	Yes
WAN Test Tool (Wtool) support	Yes	Yes	Yes
Service Level Agreement (SLA) support	Yes	Yes	Yes
Third-Party WAN Optimization Support	No	No	No

Brocade 7810 Extension Switch

The Brocade 7810 Extension Switch is a Gen 6 32G FC product. The following illustration identifies the system Management port, serial port, LEDs, FC ports, and GE interfaces on the port-side of the Brocade 7810 Extension Switch.

Figure 1: Brocade 7810 Front FC Ports and GE Interfaces

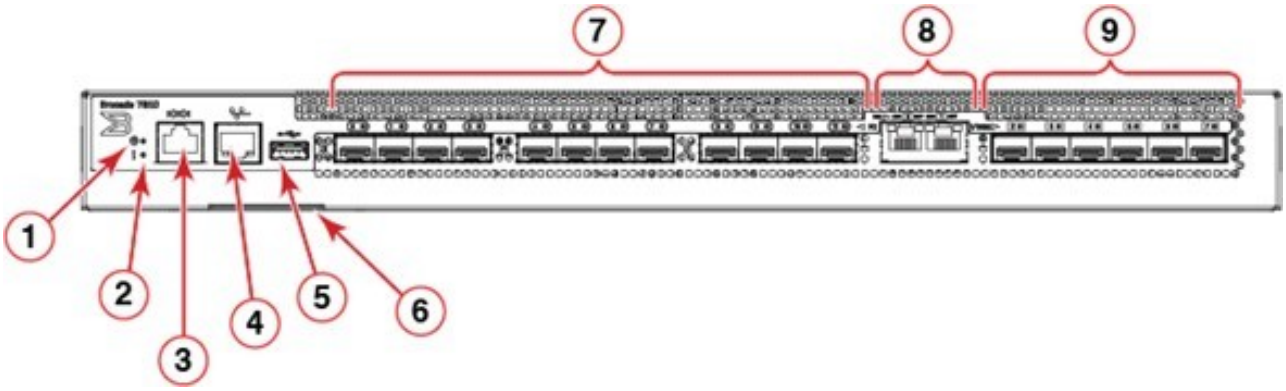

1. System power LED
2. System status LED
3. Serial console port (RJ-45)
4. Management Ethernet port (RJ-45)
5. USB port
6. Serial number pull-out tab
7. 32Gb/s enhanced small form-factor pluggable (SFP+) FC ports (FC0-11)
8. 1GbE copper (RJ-45) ports (GE0-1)
9. 1/10GbE SFP+ ports (GE2-7)

NOTE
All the FC ports are connected to one Brocade Gen 6 ASIC.

Brocade 7810 Extension Switch License Options

The Brocade 7810 Extension Switch has a single upgrade license. The Brocade 7810 Extension Switch has a base model that can be upgraded to a full model by applying the upgrade license. The base model and the upgraded (full) model support the functionality that is listed in the following table:

NOTE
The base model with the upgrade license is identical to the full model.

Table 2: Brocade 7810 Extension Switch License Options

Functionality	Base Brocade 7810 Extension Switch	Upgraded or Full Brocade 7810 Extension Switch
Number of RJ-45 interfaces (1Gb/s only)	2	2
Maximum WAN side bandwidth	1Gb/s	2.5Gb/s
Maximum FC side and IP side bandwidth	10Gb/s Requires compressibility of 10:1	10Gb/s Requires compressibility of 4:1
Number of enabled FC ports	4	12
Number of enabled VE ports	2	4
Number of enabled WAN side GE interfaces	2	6
Max. LAN side GE Interfaces	4	4
LAN side port channels	2 with 2 links each	2 with 2 links each

Functionality	Base Brocade 7810 Extension Switch	Upgraded or Full Brocade 7810 Extension Switch
Max. GE Interface speed setting	1Gb/s (SFP and RJ-45) 10Gb/s (SFP+)	1Gb/s (SFP and RJ-45) 10Gb/s (SFP+)
Max. WAN side circuit bandwidth	1Gb/s	2.5Gb/s
Max. IP Extension TCP sessions	256	256
Max. IP Extension UDP sessions	32	32
WAN side IPsec	AES 256, SHA 512	AES 256, SHA 512
Compression	Deflate, Aggressive Deflate	Deflate, Aggressive Deflate
Brocade Trunking (BT)	No	Yes
BET	No	Yes
Fabric Vision	No	Yes
ARL	Yes	Yes
FICON support	No	No
Virtual Fabrics	No	No
FC Routing (FCR)	No	Yes
Extended Fabrics	No	Yes

NOTE
On the Brocade 7850 Extension Switch and the Brocade SX6 Extension Blade, up to 64 UDP flows are optimized by batching and compressing, on the Brocade 7810 Extension Switch up to 32 UDP flows are optimized. Beyond those flows, UDP traffic is not dropped; it is sent as-is (non-optimized). Excessive UDP flows are processor constrained; therefore limited beyond the optimized flows.

Brocade 7810 Extension Switch FC Trunk Groups

To form a Brocade Trunk (BT), all inter-switch links (ISLs) must be in the same FC port group. The Brocade 7810 Extension Switch has two FC port groups as described in the following table:

Table 3: FC Port Groups on Brocade 7810 Extension Switch

Extension Platform	Port Group	FC Ports
Brocade 7810 Extension Switch	0	0-7
	1	8-11

Brocade 7810 Extension Switch VE_Port Distribution

On the Brocade 7810 Extension Switch, there is only one DP, which is DP0.

All four VE_Ports are on DP0.The VE_Port number is used when you configure a tunnel. The number of tunnels configurable per platform is determined by the number of VE_Ports.

Table 4: Brocade 7810 Extension Switch VE Port Distribution

Extension Platform	VE Mode	DP	VE_Ports
Brocade 7810 Extension Switch	N/A	DP0	12-15

Brocade 7810 Extension Switch GE Interface and Circuit Considerations

The Brocade 7810 Extension Switch has base and full/upgraded versions. A base unit can be upgraded to a full version by applying an upgrade license. The base Brocade 7810 Extension Switch has two enabled GE interfaces. An upgraded or full Brocade 7810 Extension Switch has six enabled GE interfaces. The copper RJ-45 interfaces (GE0-GE1) are 1GbE only.

Interface speed on GE2-GE7 can be set individually to either 1GbE or 10GbE on either the base or full versions.

The SFP+ interfaces (GE2-GE7) have a default speed of 1GbE.

Configuring tunnels, circuits, failover, and bandwidth requires consideration of GE interfaces. Each Brocade SX6 Extension Blade can support 10 circuits per VE, and each circuit can be up to 10Gbps, or a total of 20Gbps (10VE mode) or 10Gbps (20VE mode).

Brocade 7810 Extension Switch GE Interface Groups

The Brocade 7810 Extension Switch does not have GE groups. Each GE interface is its own group and all GE interfaces operate equally with any settings. However, the GE mode for the RJ-45 (GE0 and GE1) and SFP+ bays (GE2 and GE3) are configured as a set, GE mode cannot be configured per individual interface.

Brocade 7850 Extension Switch

The Brocade 7850 Extension Switch is a Gen 7 64G FC product. The following figure illustrates the Management port, serial port, LEDs, FC ports and GE interfaces on the port side of the Brocade 7850 Extension Switch.

Figure 2: Brocade 7850 Extension Switch Front FC Ports, GE Interfaces, and Trunk Groups

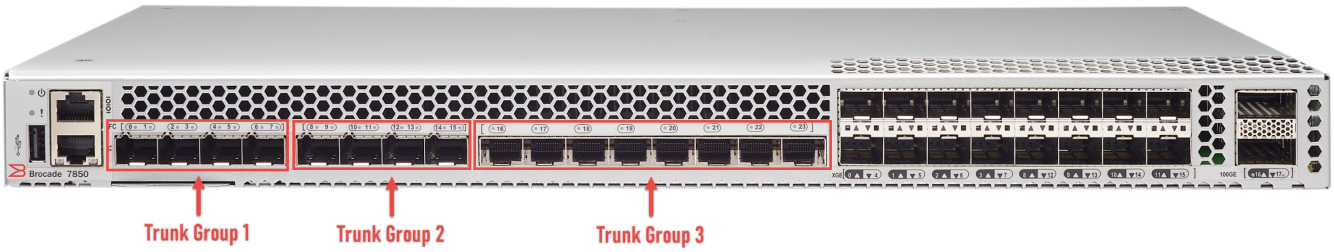

Figure 3: Brocade 7850 Extension Switch Front LEDs, Management Port, and USB

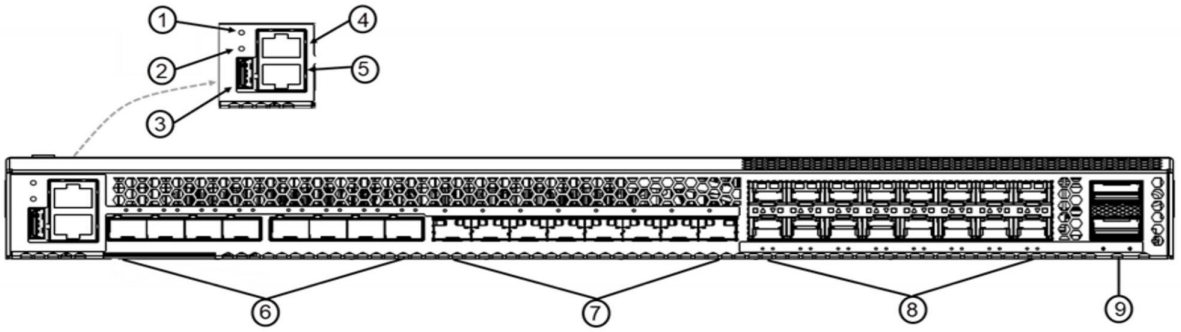

1.	System Power LED
2.	System Status LED
3.	USB Port
4.	UART RJ-45 Serial Console Port
5.	Management Ethernet Port
6.	8 SFP-DD FC ports (FC Ports 0–15)
7.	8 SFP+ FC ports (FC Ports 16–23)
8.	16 1/10/25G Ethernet ports (SFP Ethernet Ports 0–15)
9.	2 100GE QSFP ports (QSFP Ports 16–17)

Figure 4: Brocade 7850 Extension Switch Back Power Supplies and Fans

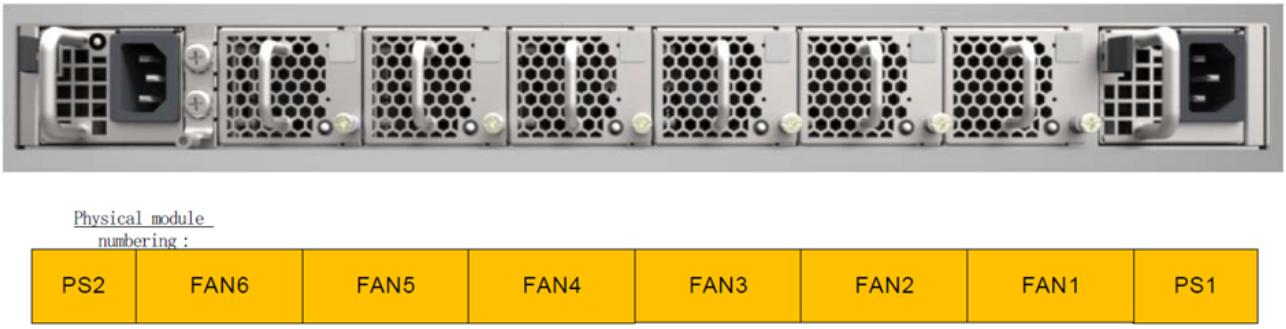

Brocade 7850 Extension Switch License Options

The Brocade 7850 Extension Switch does not have a license option. Platforms are available in two versions: open systems and mainframes. In the open systems model, all features are enabled except FICON, Advanced FICON Accelerator, and FICON Management Server (FMS) with the CUP support. A mainframe model includes all the features and functionality of an open systems model, and FICON-related features. In addition, optics are delivered differently in open systems compared to mainframe models (SWL MMF versus LWL SMF).

An open systems version cannot be upgraded to a mainframe version and the reverse is also true. This requires a TruFOS License. Encryption is a licensed feature, if encryption functionality is permitted by law, encryption is included; otherwise, the license might not be included. The Brocade 7850 Extension Switch does not support FIPS Extension features.

Brocade 7850 Extension Switch FC Trunk Groups

To form a Brocade Trunk (BT), all ISLs must be in the same FC port group. The Brocade 7850 Extension Switch has three FC port groups as described in the following table:

Table 5: FC Port Groups on Brocade 7850 Extension Switch

Extension Platform	Port Group	FC Ports
Brocade 7850 Extension Switch	1	0-7
	2	8-15
	3	16-23

Brocade 7850 Extension Switch VE_Port Distribution

On the Brocade 7850 Extension Switch, depending on the VE mode, each DP (DP0 and DP1) supports either 3 or 9 VE_Ports, a total of 6 or 18 per platform.

The following table shows the VE_Ports supported in each VE mode. The VE_Port number is used when you configure a tunnel. The number of tunnels configurable per platform is determined, in part, by the VE mode and the number of VE_Ports supported on each DP.

Table 6: Brocade 7850 Extension Switch VE Port Distribution

Extension Platform	VE Mode	DP	VE_Ports
Brocade 7850 Extension Switch	6VE	DP0 (group 1)	24-26
		DP1 (group 1)	33-35
	18VE	DP0 (group 1)	24-26
		DP0 (group 2)	27-29
		DP0 (group 3)	30-32
		DP1 (group 1)	33-35
		DP1 (group 2)	36-38
		DP1 (group 3)	39-41

- Brocade 7850 Extension Switch in 6VE mode – A VE_Port can use all the FC bandwidth available to the DP up to a maximum of 50Gb/s.
- Brocade 7850 Extension Switch in 18VE mode – A VE_Port can use one third of the FC bandwidth available to the DP up to a maximum of 16Gb/s.

VE_Ports are used to connect multiple sites. As a best practice, use 6VE mode unless more than 6VE ports need to be connected.

Brocade 7850 Extension Switch GE Interface and Circuit Considerations

Configuring tunnels, circuits, failover, and bandwidth requires consideration of GE interfaces. Each VE_Port is limited to 10 circuits. The maximum VE_Port rate for a single circuit is 25Gb/s. The bandwidth of a circuit cannot exceed the bandwidth of the GE interface to which it has been assigned.

Brocade 7850 Extension Switch supports 1GbE, 10GbE, 25GbE, and 100GbE speeds. GE0-GE15 are SFP+ interfaces that support 1GbE, 10GbE, and 25GbE. GE16-GE17 are quad small form-factor pluggable (QSFP) interfaces that support 100GbE.

GE0-GE15 interfaces do not auto-negotiate to 1GbE or 10GbE speeds; the desired speed must be configured, and SFPs with the appropriate speed must be inserted. A 1GbE SFP cannot operate at 10GbE and a 10GbE SFP+ cannot operate at 1GbE. The SFP+ interfaces (GE0-GE15) have a default speed of 10GbE.

GE0-GE15 supports optics that can operate at 10GbE or 25GbE speeds.

The QSFP interfaces (GE16-GE17) support 100GbE speeds. The QSFP interfaces (GE16-GE17) can only be used on the WAN side and cannot be used on the IP Extension LAN side.

Brocade 7850 Extension Switch GE Interface Groups

This section describes the GE interfaces on a Brocade 7850 Extension Switch. There are 16 1GbE/10GbE/25GbE interfaces in eight groups. Specific recommendations apply to interfaces within a group to help alleviate blocking issues.

The Brocade Extension GE interfaces are numbered. On the Brocade 7850 Extension Switch, the 100Gb/s interfaces are GE16 and GE17. The 1GbE/10GbE/25GbE interfaces are GE0 through GE15.

The GE interface groups are shown in the following table:

Table 7: Brocade 7850 Extension Switch GE Interface Groups

GE Interface Number	GE Interface Group
0, 4	1
1, 5	2
2, 6	3
3, 7	4
8, 12	5
9, 13	6
10, 14	7
11, 15, 16, 17	8

Interface group 1 contains the two 100GbE interfaces (16 and 17) and two 25GbE interfaces (11 and 15). The remaining interface groups contain the 25GbE/10GbE/1GbE interfaces from GE0 to GE15. Use these interface groups with the following considerations:

- An interface can block any interface in its group, but it cannot block a port outside of its group.
- An interface can block another interface in its group due to interface speed differences (1Gb/s, 10Gb/s, or 25Gb/s) or if the interface is back-pressured from an external switch applying Ethernet pause. A slow-draining device or other congestion can result in a blocked port.

To avoid blocking, do not mix interface speeds within a group. Following are the recommendations for the port group:

- For port group 1, since the 100GbE ports are fixed at 100Gb/s, either (GE16 and GE17) or (GE11 and GE15) should be used, but not both. The bandwidth for GE11 and GE15 should either be 1Gb/s or 10Gb/s or 25Gb/s, not a combination.
- In port groups 2 through 8, which contain all 25GbE ports, configure the ports at either 25Gb/s or 10Gb/s or 1Gb/s. Do not use the combination of speeds within the port group.

> **NOTE**
> The previous table applies to interfaces configured in the WAN mode. If the ports are configured as LAN ports, grouping and blocking does not apply. You should allocate a LAN interface out of the same group as a WAN interface.

Brocade 7850 Extension Switch Forward Error Correction (FEC)

The Brocade 7850 Extension Switch has 25GbE and 100GbE interfaces that support FEC (RS-FEC IEEE 802.3), and FEC Clauses CL108 (25G) and CL91 (100G). By default FEC is enabled, but you can disable this feature.

Brocade SX6 Extension Blade

The Brocade SX6 Extension Blade is a Gen 6 32G FC product. The following figure illustrates the FC ports and GE interfaces on the Brocade SX6 Extension Blade.

Figure 5: Brocade SX6 Extension Blade Front FC Ports and GE Interfaces

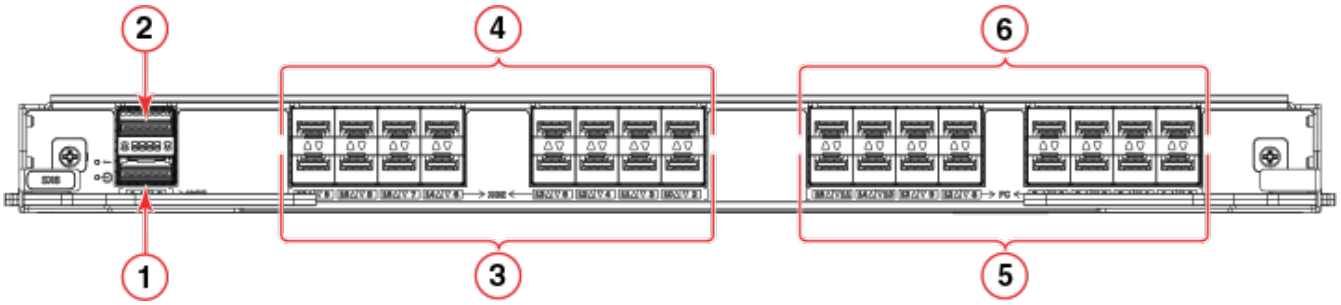

1. 40GE interface GE0
2. 40GE interface GE1
3. 1/10GE interface 2 through 9 (right to left)
4. 1/10GE interface 10 through 17 (right to left)
5. FC ports 0, 1, 2, 3, 8, 9, 10, 11 (right to left)
6. FC ports 4, 5, 6, 7, 12, 13, 14, 15 (right to left)

The Brocade SX6 Extension Blade has the following FC ports and optics:

- 16 FC SFP+ ports supporting F_Port, E_Port, and EX_Port port types
- FC SFP+ optics capable of either (32Gb/s, 16Gb/s, and 8Gb/s) or (16Gb/s, 8Gb/s, and 4Gb/s)
- Installed SFP+ optic that determines the maximum and minimum settable speed
- FC ports that can auto-negotiate speed within the SFP+'s capable range
- FC ports support 10Gb/s SFP+ optics, which are typically used for metro DWDM

The Brocade SX6 Extension Blade has the following GE interfaces and optics:

- Ethernet interfaces default to WAN side for tunnels, or can be set to LAN side for IP Extension
- Two 40GbE QSFP Ethernet interfaces
- 40GbE QSFP optics operate at only 40Gb/s
- 16 1/10GbE SFP/SFP+ Ethernet interfaces
- 1GbE SFPs and 10GbE SFP+s operate at 1Gb/s and 10Gb/s fixed speeds respectively
- 1/10GbE interface speed that is compliant with the installed transceiver
- 1/10GbE SFPs do not have multiple speeds and do not auto-negotiate speed; the transceivers only support one speed, and transceivers do not automatically adjust to the set interface speed

Brocade SX6 Extension Blade License Options

The Brocade SX6 Extension Blade is supported by both the Brocade X7 Director and the Brocade X6 Director. The following licenses are supported by these directors:

- Extended Fabrics (Long-Distance ISLs)
- BT
- FMS
- ICL
- FCR
- Fabric Vision and IO Insight

For more information about the licenses available for the Brocade X6 Director and Brocade X7 Director, refer to the *Brocade Fabric OS Software Licensing User Guide*.

Brocade SX6 Extension Blade FC Port Groups

To form a BT, all ISLs must be in the same FC port group. Two FC port groups are supported on the Brocade SX6 Extension Blade as described in the following table:

Table 8: FC Port Groups on Brocade SX6 Extension Blade

Extension Platform	Port Group	FC Ports
Brocade SX6 Extension Blade	0	0-7
	1	8-15

Brocade SX6 Extension Blade VE_Port Distribution

On a Brocade SX6 Extension Blade, depending on the VE mode, each DP (DP0 and DP1) supports either 5 or 10 VE_Ports, a total of 10 or 20 per blade.

The following table shows the VE_Ports supported in each VE mode. The VE_Port number is used when you configure a tunnel. The number of tunnels configurable per platform is partly determined by the VE mode and the number of VE_Ports supported on each DP.

Table 9: Brocade SX6 Extension Blade VE Mode Port Distribution

Extension Platform	VE Mode	DP	VE_Ports
Brocade SX6 Extension Blade	10VE	DP0	16-20
		DP1	26-30
	20VE	DP0	16-25
		DP1	26-35

- Brocade SX6 Extension Blade in 10VE mode – A VE_Port can use all the FC bandwidth available to the DP up to a maximum of 20Gb/s, if licensed.
- Brocade SX6 Extension Blade in 20VE mode – A VE_Port can use half the FC bandwidth available to the DP up to a maximum of 10Gb/s. 20VE mode allows more VE_Ports at a lower maximum bandwidth. Typically, the purpose of using more VE_Ports is to connect more sites. The best practice is to use 10VE mode unless more than 10 sites must be connected.

Brocade SX6 Extension Blade GE Interface and Circuit Considerations

There are GE interface considerations when configuring tunnels, circuits, failover, and bandwidth. Each VE_Port is limited to 10 circuits. The maximum VE_Port rate for a single circuit is 10Gb/s. A circuit cannot be configured with more bandwidth than the GE interface it has been assigned to.

Brocade SX6 Extension Blade supports 1GE, 10GE, and 40GE speeds.

GE0 and GE1 are 40GE QSFP interfaces that support only 40Gb/s and are WAN side only. The 40GE interfaces cannot be used for the IP Extension LAN side.

GE2-GE17 are SFP+ interfaces that support 1GE and 10GE. The interface does not auto-negotiate to 1GE or 10GE speed; the desired speed must be configured and the appropriate speed SFP must be inserted. A 1GE SFP cannot operate at 10GE and a 10GE SFP+ cannot operate at 1GE. GE2-GE17 have a default speed setting of 10GE.

Configuring tunnels, circuits, failover, and bandwidth requires consideration of GE interfaces. Each VE_Port is limited to six circuits. The maximum VE_Port rate for a single circuit is 2.5Gb/s. Circuits cannot be configured with more bandwidth than the Ethernet interface to which they are assigned.

Brocade SX6 Extension Blade GE Interface Groups

This section describes the GE interfaces on a Brocade SX6 Extension Blade. The Brocade SX6 Extension Blade has 18 1GbE/10GbE interfaces in eight groups. Specific recommendations apply to interfaces within a group to help alleviate blocking issues.

The Brocade Extension GE interfaces are numbered. On the Brocade SX6 Extension Blade, the 40Gb/s interfaces are GE0 and GE1. The 1GbE/10GbE interfaces are GE2 through GE17.

The following table shows the GE interface groups:

Table 10: Brocade SX6 Extension Blade GE Interface Groups

GE Interface Number	GE Interface Group
0, 1, 13, 17	1
2, 6	2
3, 7	3
4, 8	4
5, 9	5
10, 14	6
11, 15	7
12, 16	8

Interface group 1 contains the two 40GbE interfaces (0 and 1) and two 10GbE interfaces 13 and 17. The remaining interface groups contain the 1GbE/10GbE interfaces from GE2 to GE16. Use these interface groups with the following considerations:

- An interface can block any interface in its group, but it cannot block a port outside of its group.
- An interface can block another interface in its group due to interface speed differences (1Gb/s vs 10Gb/s) or if the interface is back-pressured from an external switch applying Ethernet pause. A slow-draining device or other congestion can result in a blocked port.

To avoid blocking, do not mix interface speeds within a group. Following are the recommendations for the port group:

- For port group 1, since the 40GbE ports are fixed at 40Gb/s, either (GE0 and GE1) or (GE13 and GE17) should be used, but not both. The bandwidth for GE13 and GE17 should either be 1Gb/s or 10Gb/s, not a combination of both.
- In port groups 2 through 8, which contain all 10GbE ports, configure the ports at either 10Gb/s or 1Gb/s. Do not use the combination of speeds within the port group.

NOTE
The previous table applies to interfaces configured in the WAN mode. If the ports are configured as LAN ports, grouping and blocking does not apply. You should allocate a LAN interface out of the same group as a WAN interface.

Brocade SX6 Extension Blade VE Modes and Internal Bandwidth

The following figure illustrates the Brocade SX6 Extension Blade's two VE modes of operation in FCIP mode, not hybrid mode. The diagrams show the DPs' internal connections when 10VE or 20VE mode is enabled. The difference is at the VE_Ports level. The 10VE mode has a 5 VE_Port block with 20Gb/s of connectivity. The 20VE mode has two 5 VE_Port blocks each with 10Gb/s of connectivity.

All 10Gb/s, 20Gb/s, and 40Gb/s connections shown in the illustration are full-duplex and internal to the platform.

Figure 6: DP Components and VE_Port Distribution in 10VE and 20VE Modes

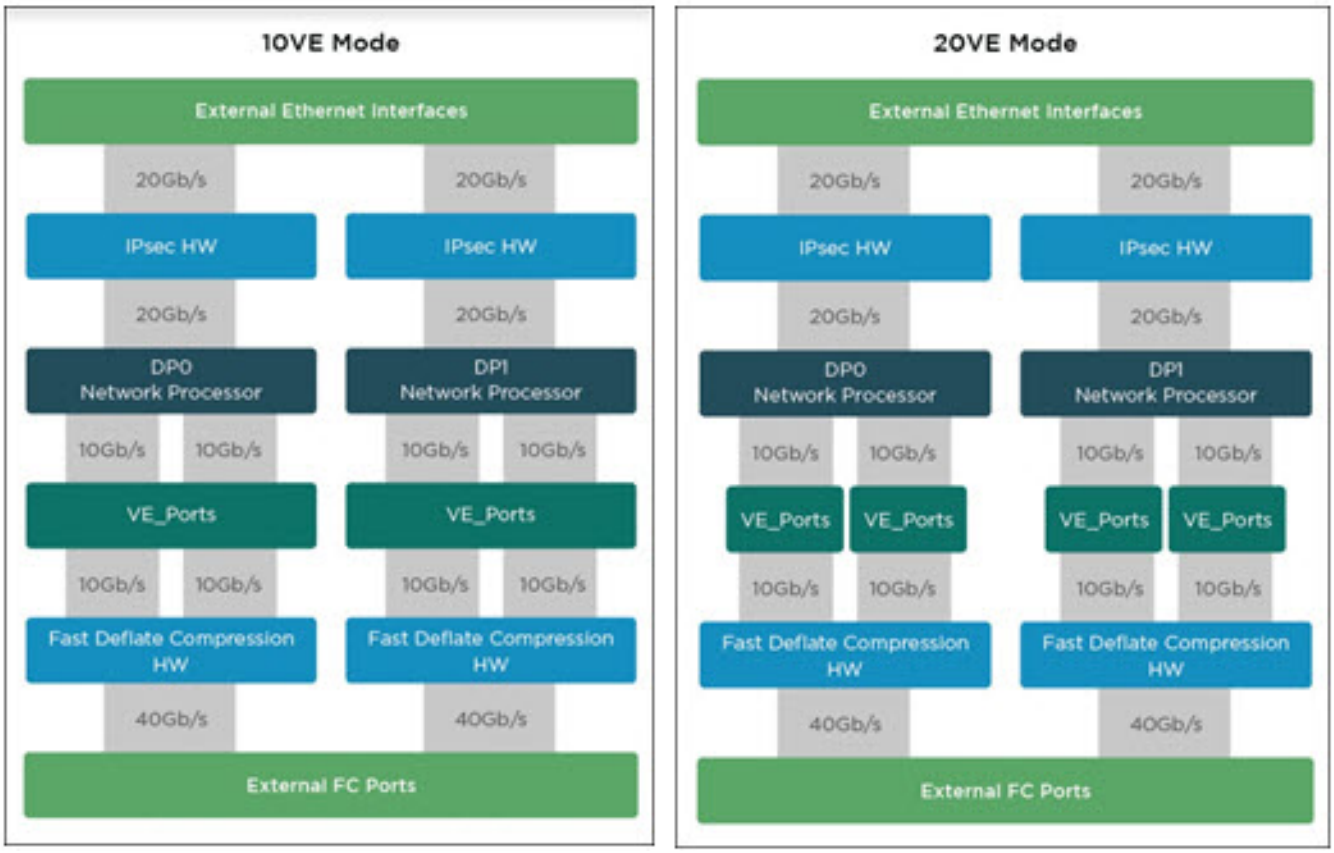

The following notes apply to the Brocade SX6 Extension Blade DP illustration:

- Each DP owns specific VE_Ports; selecting a VE_Port determines which DP is used.
- There is a 40Gb/s full-duplex connection between the external FC ports switching ASIC and each DP.
- FC frames can be compressed using the fast deflate compression hardware. Deflate and aggressive deflate compression options are also available, but do not use the hardware shown in the illustration.
- The maximum bandwidth size of any one tunnel across the internal connections can be no more than 10Gb/s.
- From the network processors, data can be encrypted by dedicated IPsec hardware using a high-speed, low-latency encryptor/decryptor protocols.
- Each DP can process a combined 20Gb/s of data to and from the WAN.
- In 10VE mode, each DP supports a total of five VE_Ports.
- The 10VE mode is required for hybrid mode.
- In 20VE mode, each DP supports a total of 10 VE_Ports.
- The 20VE mode requires FCIP mode.
- If a 2:1 compression ratio is achieved using fast deflate, 40Gb/s per DP is available to external FC ports.
- The ARL aggregate of all circuit maximum values on a DP cannot exceed 40Gb/s.
- The ARL aggregate of all circuit minimum values for a DP cannot exceed 20Gb/s.
- The term *all circuits* includes all circuits from all tunnels, not just all circuits from a single tunnel.

NOTE
Compression can achieve different compression ratios depending on the compressibility of data. Broadcom makes no promises, warranties, or guarantees to the achievable compression ratios for customer-specific data.

Brocade Trunks (BT)

Brocade Trunks are a Fibre Channel technology, not an extension technology. The Brocade Extension platform contains fully functional FC switches that support BTs when connected to other Brocade platforms. The following list contains the requirements for forming a BT:

- The ports must belong to the same ASIC port group to form a BT.
- The ports must belong to the same Brocade adapter if the F_Port trunking is used.
- Brocade platforms must be directly connected to each other, no active intermediate devices are supported.
- The use of passive devices is supported if they do not introduce a skew.
- Two to eight ports within a port group can form a BT.
- Port speed, encryption, compression, QoS, FEC, and long-distance settings must be the same within a BT.
- BTs are proprietary and not supported on M-EOS or third-party switches.

For more information about BT requirements and configuration, refer to the *Brocade Fabric OS Administration Guide*.

Brocade Extension Concepts and Features

Extension is a tunneling protocol to transmit Fibre Channel (FC) and IP Storage over a distance across standard IP networks.

Three Sides of Brocade Extension Platform

This section describes the three different sides of a Brocade Extension platform: FC and FICON, WAN, and LAN.

These three sides of extension platforms are associated with the types of connectivity a Brocade Extension platform is capable of. The Extension platform provides FC and IP storage connectivity over an IP WAN infrastructure.

The Extension platform is a part of a fast and reliable network infrastructure. It addresses storage area network (SAN) requirements, primarily to extend FC and IP storage beyond its natural effective native reach. FCIP and IP Extension protocols are implemented in a Brocade Extension platform. The Extension platform is used primarily for remote data replication (RDR), centralized tape backup, and data center data migration. The Extension platform is not used to connect local servers to a remote storage. Without an extension, the distance between an initiator and target is limited to a few kilometers. TCP has limited reach at large bandwidths before experiencing a drop. This drop is the drop of throughput over distance; moreover, it is reduced by packet drops.

The Brocade 7850 Extension Switch and Brocade SX6 Extension Blade support FC, FICON, and IP storage data flows. FC and IP storage are supported by the Brocade 7810 Extension Switch; however, FICON storage is not supported. Brocade Extension solutions maximize the replication and backup throughput over distance by using data compression, disk and tape protocol acceleration, and WAN-Optimized TCP. The Brocade Extension platform uses existing WAN infrastructure to connect FC end-devices and IP storage between distant endpoints that are either impractical or too costly using FC or native IP connectivity.

Figure 7: Extension Protocol Stack

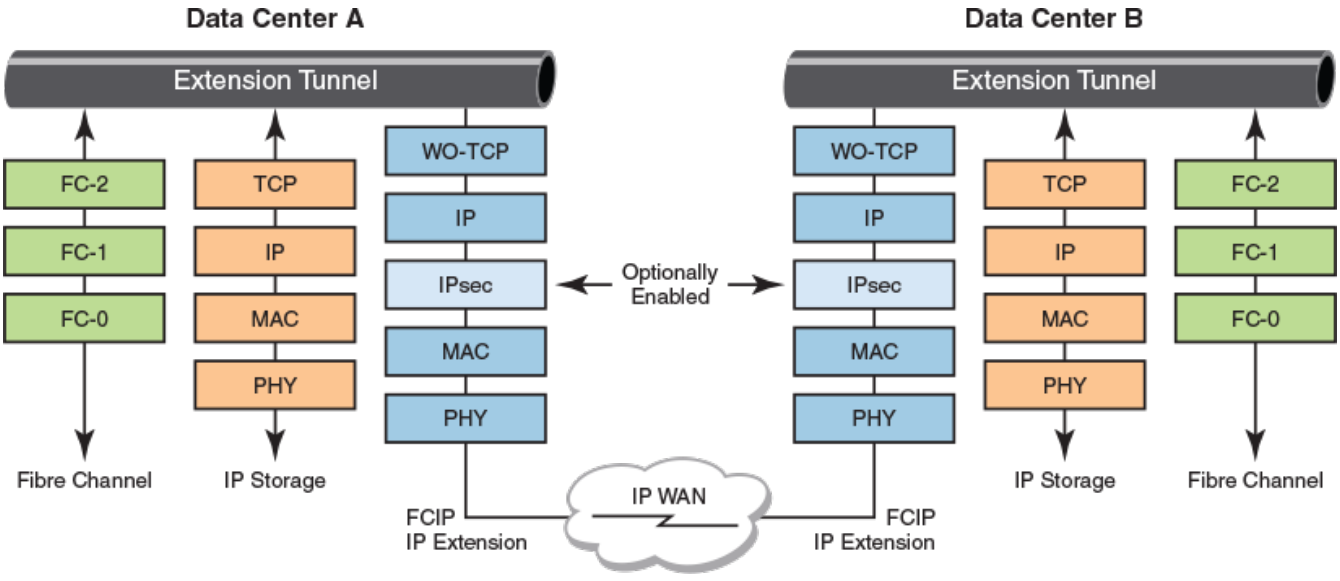

The FC and FICON Side

The FC switch in a Brocade Extension platform is a full-fledged FC switch. It is the same as in any other Brocade FC platform, and it runs the same FOS software. Also, the Brocade 7850 Extension Switch and Brocade SX6 Extension Blade platforms support FICON connectivity and features. Any FC or FICON-based traffic that flows into the switch can be sent across the VE_Port (Extension tunnel) to the remote VE_Port. VE_Ports are similar to an E_Port in that it forms an

ISL between logical switches or switches. If there are multiple VE_Ports that connect the same two domains, APTpolicy is used to determine which type of routing (EBR, DBR, and PBR) should be applied to the VE_Ports.

FC Considerations

While implementing and designing FC storage area networks (SANs), many considerations must be taken into account. In most cases, the replication fabrics, which include extension switches and blades, are separate from the production fabrics. The Brocade Extension Blades such as the Brocade SX6 Extension Blade can be placed in a logical switch within a chassis; however, it does not increase redundancy because it is within the same chassis. Typically, there are two replication fabrics, A and B, and the replication ports of the end devices are directly connected to the replication fabrics.

For more information, refer to the *Brocade SAN Design and Best Practices* document.

The WAN Side (Tunnel)

FC was not designed to traverse an IP WAN, but the Brocade Extension platform was designed to address the intricacies of an enterprise WAN. This document focuses mostly on the WAN side of the technology.

The WAN side consists of the following components:

- IP interfaces (IPIF)
- IP routes
- Tunnels
- Circuits
- BET
- Failover/failback (metrics)
- Encryption (IPsec)
- Compression
- ARL
- Extension Hot Code Load (eHCL)
- SLA
- WAN Test Tool (Wtool)

The WAN side connects to the IP network and has tunnels and circuits that extend across the WAN to a switch in a remote data center.

An extension tunnel is an ISL that transports data between extension platforms. The WAN side uses High-Efficiency Encapsulation (HEE), protocol optimization, encryption, and compression to transport data from the FC/FICON and LAN sides to the corresponding side at the remote end.

A tunnel is the basis of an extension connection, which is composed of one or more circuits between two extension endpoints. Extension tunnels allow FC and IP storage traffic to pass through a shared WAN optimally. An extension tunnel ensures lossless transmission and in-order delivery when multiple circuits are implemented. Whether FC or IP storage, end-devices do not experience the data reliability or latency issues that are often associated with WAN connections, and storage applications are unaware of the intermediate extension devices.

The circuits of an extension tunnel provide the following:

- High availability
- Load balancing across disparate network paths
- High-Efficiency Encapsulation (HEE)
- WAN Optimized TCP (WO-TCP)
- Adaptive Rate Limiting (ARL)
- Lossless link loss (LLL)

Brocade FCIP

Fibre Channel over IP (FCIP), is a tunneling protocol to link FC over distance on standard IP networks. Extension platforms enable FC connectivity over IP networks between FC devices or fabrics over IP networks. It is primarily used for remote data replication (RDR), tape backup, and migration. The extension link is an ISL that transports FC control and data frames between switches. FCIP is supported on all Brocade Extension platforms.

Brocade uses a proprietary tunnel protocol that is capable of moving both FC and IP storage across the same tunnel, and capable of supporting advanced features. Brocade FCIP is not interoperable with extension products from other vendors.

Brocade Extension platform provides the following features:

- High-speed performance
- Seamless use of disparate WAN paths
- Network resiliency using Brocade Extension Trunking (BET)
- Adaptive Rate Limiting (ARL)
- Lossless link loss (LLL)
- High-speed compression
- High-speed encryption
- Jumbo Frames (a maximum of 9216 bytes)
- TCP Acceleration with WAN optimized TCP (WO-TCP)
- Protocol acceleration for FICON, SCSI, and tape
- QoS with DSCP and/or L2CoS (802.1P) marking
- A deterministic path for a protocol acceleration
- Failover, failback, failover groups, failover metrics, and spillover
- Efficient use of VE_Ports
- Aggregation of circuit bandwidth
- In-order delivery (IOD)
- WAN bandwidth pooling—pool bandwidth from multiple links/providers
- Diagnostic and troubleshooting tools: SAN Health and WAN Tool
- Service Level Agreement (SLA)

NOTE
FICON is supported only on the Brocade 7850 Extension Switch and the Brocade SX6 Extension Blade. It is not supported on the Brocade 7810 Extension Switch.

Within the architecture of Brocade Extension platforms, there are Brocade application-specific integrated circuits (ASICs) that are FC-specific. The VE_Ports are logical representations of FC ports. VE_Ports are not part of the Brocade FC switching ASICs. Multiple internal FC ports feed a VE_Port, with a maximum FC rate of 10 Gb/s for a Brocade 7810 Extension Switch, a maximum rate of 100Gb/s for a Brocade 7850 Extension Switch, and a maximum rate of 40Gb/s for a Brocade SX6 Extension Blade. A VE is the transition point between FC and the FCIP tunnel within the extension platforms.

Table 11: FCIP Protocol

Network	WAN/MAN
Transport	FCIP/TCP/IP/Ethernet
Encapsulation	Brocade encapsulation is called HEE (High Efficiency Encapsulation). FC sequences become compressed byte streams that form batches. Batches fill TCP segments to their maximum segment size.
IP-routable	Yes

The LAN Side (IP Extension)

As the name implies, the LAN side is connected to the LAN of the data center to allow IP storage end-devices to access the IP Extension Gateway. The LAN side is specific to the IP Extension. FCIP does not use the LAN side. The LAN side is entirely processed in each DP, the IP Extension does not utilize the Condor ASIC or the FC switching portion of the platform. Although, the LAN side and FC/FICON sides share the WAN side, both can use the same tunnel at the same time. This section of the document describes the LAN side.

The IP Extension feature provides enterprise-class support for IP storage applications using existing IP-wide area network (WAN) infrastructure. IP Extension features are offered only on the Brocade 7850 Extension Switch, Brocade 7810 Extension Switch, and Brocade SX6 Extension Blade platforms.

The Brocade 7850 Extension Switch and Brocade SX6 Extension Blade support IP storage, FCP/SCSI, and FICON data flows. Brocade 7810 Extension Switch supports IP storage and FCP/SCSI data flows, but does not support FICON data flows. The IP storage data flows across a Brocade Extension tunnel is referred to as an IP Extension, which enhances IP storage transmissions across the enterprise WAN infrastructure. Also, the IP Extension gives you visibility into and control of flows using various diagnostic tools, IPsec, compression, ARL, QoS, BET, and lossless tunnel resiliency. IP Extension supports applications such as array-native IP-based remote data replication (RDR), IP-based centralized backup, and Grid replication. In addition, IP Extension supports host-based and database replication over IP, NAS replication between data centers, and data migration between data centers.

IP Extension data across the IP WAN uses WAN Optimized TCP (WO-TCP), a highly efficient and aggressive TCP stack for moving data between data centers. Brocade WO-TCP ensures in-order lossless transmission of IP Extension data. IP Extension establishes a proxy TCP endpoint for local devices. Local devices are unaware and unaffected by the latency and quality of the IP WAN. IP Extension accelerates end-device native TCP.

The IP Extension provides the following advantages:

- High-speed performance
- Unified support and management of FC, FICON, and IP storage
- Provision once and over time connect many devices
- WAN bandwidth pooling: pool bandwidth from multiple links/providers
- Lossless Link Loss (LLL)
- Adaptive Rate Limiting (ARL)
- Network resiliency using Brocade Extension Trunking (BET)
- Efficient protocol transport: negligible added overhead
- TCP Acceleration with WAN Optimized TCP
- Streams: Virtual windows on WAN optimized TCP to minimize the head of line blocking (HoLB)
- High-speed compression
- High-speed encryption
- Diagnostic and troubleshooting tools: SAN Health and WAN Tool
- QoS for FCIP and IP Extension
- DSCP and/or 802.1P marking
- Internal prioritization of FC traffic versus IP traffic with QoS distribution
- Jumbo Frames (9216 bytes)

Table 12: IP Extension Protocol

Network	WAN/MAN
Transport	IP Extension/TCP/IP/Ethernet
Encapsulation	Brocade encapsulates IP data flows (also called “streams”) into a byte stream that forms compressed batches. Those batches fill TCP segments to their maximum segment size and then form datagrams. The process is called HEE.
IP-routable	Yes

Configuring Brocade Extension

This section provides information on configuring Brocade Extension.

The configuration consists of two main phases. The first phase is the planning and preparation phase. You must log into the Brocade Extension platform or use Brocade SANnav to complete the following steps in the second phase of the configuration process:

- GE interfaces
- IP interfaces
- IP routes
- IPsec
- Tunnels
- Circuits

For more information about Brocade SANnav, refer to the *Brocade SANnav Management Portal User Guide*.

Brocade Extension configuration typically uses a sequence similar to the following list. All listed items are not applicable to all platforms or deployments.

Perform the following steps to configure the Brocade Extension platform:

1. Ensure that the extension platforms have been installed, powered, and cabled.
2. Complete the following basic Fabric OS configurations:
   a) Mgmt IP address
   b) Switch name
   c) Domain ID
   d) Default zoning
   e) Logical switches
   f) DNS
   g) NTP
   h) AAA
   i) SNMP
   j) IP filter
   k) Other
3. Configure the Application mode.
4. Configure the VE mode.
5. Configure the GE mode.
6. Configure the GE interface speeds.
7. Make sure that LLDP sees peer switches and expected peer ports, and conversely (data-link connectivity).
8. Configure the WAN side IP interfaces (IPIF).
9. Configure the WAN side HA IP interfaces (IPIF) for eHCL.
10. Configure the WAN side IP routes for all IPIF (Layer 3 deployment only).
11. Configure IPsec policies (optionally, PSK or PKI).
12. Create the extension tunnel.
    a) Enable IP Extension (if applicable).
    b) Specify the IPsec policy (if applicable).
    c) Specify the compression algorithm (default is none).
13. Create the Extension circuits.
14. Modify each circuit by adding the local and remote IP addresses.
15. Modify each circuit by adding the local and remote HA IP addresses (if applicable).
16. Modify each circuit by adding the min and max bandwidth.
17. Configure the SLA function (optional).
18. If you deploy IP Extension:
    a) Configure GE interfaces for LAN side.
    b) Configure LAN side port channel to data center LAN switch.
    c) Configure the LAN side IP interfaces (IPIF).
    d) Configure the LAN side IP routes, if applicable (only with Layer 3 deployment using PBR).
    e) Configure the TCL.
19. Validate the configuration.

Configuration Checklist

Before you begin, see Circuit Requirements and Tunnel Requirements and consider the following checklist:

Table 13: Configuration Checklist for a Simple Single FCIP Tunnel

Prerequisite	Value to Generate/Provide	Complete?
Have the Brocade platforms been installed into a rack?	Validation by SAN Admin	
Have the Brocade platforms been connected to a dual power?	Validation by SAN Admin	
Have the Brocade platforms been connected to the Mgmt network?	Ping validation by SAN Admin	
Have the Brocade platforms been cabled to the storage replication ports?	Validation by Storage Admin	
Have the Brocade platforms been cabled to the IP WAN network?	Ping validation by Network Admins	
The Brocade platform has undergone basic configuration for: • Switch name • Fabric name • Support Link • Domain ID • AAA (LDAP and RADIUS) • SSH access • DNS • NTP • IPfilter • syslog • default zoning • Any other security configurations	Validation by SAN Admin	
Has the WAN IP network been provisioned and is the IP network operational from end-to-end?	Validation by Network Admins	
Specific to each platform, determine the required Application mode, VE Mode, and GE Mode settings?	Application mode VE Mode GE Mode	
Determine the amount of bandwidth required to satisfy your replication needs?	Bandwidth in bps	
Determine the VE_Port you use for the tunnel? Typically, the first one is used, and the same number on both ends.	VE_Port number	

Prerequisite	Value to Generate/Provide	Complete?
Determine how many circuits that the tunnel needs for aggregate bandwidth and resiliency purposes? Will additional circuits be needed for failover (metric-0 and metric-1)? Often when using BET, failover circuits are not needed.	Number of circuits	
Determine the minimum and maximum bandwidth of each circuit? If ARL is not being used, setting the minimum = maximum rate disables ARL and enables Committed Information Rate (CIR) at the rate specified.	The min and max bandwidth of each circuit.	
Determine which GE interfaces will be used for each circuit?	List the GE interfaces and association to each circuit.	
Obtain the IP subnets, subnet masks, and end-to-end network MTU. Assign IP addresses from the subnets to IPIF of each circuit.	Subnets Subnet masks IP MTU	
Determine local and remote IP addresses for each circuit?	The local and remote IP addresses for each circuit	
Determine local and remote eHCL IP addresses for each circuit? Configuring eHCL is optional and the extension platform must support eHCL.	The local and remote HA IP addresses for each circuit.	
If the local and remote subnets are different, L3 network, determine the gateway IP address. Frequently, there is only one gateway; however, there could be different gateways for different remote subnets. The gateway is on the same subnet as the IP address of the IPIF that uses the gateway. The route also contains the remote subnet and its mask.	Gateway IP Addresses	

General Extension Configuration

Some aspects of Brocade Extension configuration are not specific to a particular side, and might apply to multiple sides.

Traffic Optimizer over FCIP

The Traffic Optimizer (TO) feature provides a mechanism to optimize traffic by automatically grouping FC flows into predefined Performance Groups (PG). TO is supported only between the Brocade 7850 Extension Switches, not across the Brocade 7810 Extension Switch or the Brocade SX6 Extension Blade platforms. On the Brocade 7850 Extension Switch, any FC traffic passing through a tunnel follows the configured tunnel bandwidth and QoS distribution and priority percentages. TO uses Extension QoS; therefore, if there are high-priority flows, DSCP or L2CoS markings must be used.

Tunnels detect whether each end supports TO, and if it does, TO is supported across the tunnel.

NOTE
TO over FCIP is only available between the Brocade 7850 Gen 7 platforms with Fabric OS 9.2.0 or later.

The TO as an Extension feature on the Brocade 7850 Extension Switch, the following three priorities are supported when the traffic goes through an extension tunnel:

- PG_QOS_HIGH uses high priority
- PG_SDDQ_LOW uses low priority
- Remaining PGs use medium priority

TO PG classification information is preserved across supported extension tunnels. PG classification is once again enforced at the remote end through the rest of the FC fabric to the end-device.

Zoning

Zoning is a fabric-based service that enables partitioning of FC storage ports into logical groups of devices that can access each other. Zoning protects FC storage ports from unauthorized and inadvertent access.

For example, a fabric can be partitioned into two zones, winzone and unixzone to prevent Windows servers and storage from inadvertently interfering with Unix servers and storage. Zones can be used to logically consolidate equipment for efficiency or to facilitate time-sensitive functions; for example, you can create a temporary zone to back up non-member devices.

A zoned device can communicate only with other connected devices within the same zone. Devices not included in the zone are not available to zone members. When the `defzone` command is set to `noaccess` , devices not belonging to a zone cannot communicate with other devices in the fabric. Best practice is to set the `defzone` command to `noaccess` .

By contrast, if the `defzone` command is set to `allacces` and no zones are defined, each device in the fabric can access every other device. In a small extension fabric with few replication ports, `allacces` might not be a problem. In large replication fabrics or extension fabrics that are connected to production fabrics, allowing unzoned devices access to other devices presents serious security issues, which is why zoning is recommended.

Zones can be configured dynamically. They can vary in size, depending on the number of fabric-connected devices, and devices can belong to more than one zone.

For more information about configuring zones, refer to the "Administering Advanced Zoning" chapter in the *Brocade Fabric OS Administration Guide*.

LLDP

Link Level Discovery Protocol (LLDP) is supported on the Brocade Extension platforms. It is a vendor-neutral open protocol, referred to as IEEE 802.1AB. LLDP is a Layer 2 protocol that is used to exchange device identity and port numbers with directly connected data center switches. LLDP applies to all extension GE interfaces on both the WAN and LAN sides. LLDP does not apply to the management port. Frequently, LLDP is used to ensure connectivity to the intended device and port; also, it is used in troubleshooting to ensure that the Ethernet link is functioning. If an LLDP exchange stops due to a physical layer problem (cable, optics, or configuration), the LLDP entry of a port will time out and will be removed from the listing. LLDP is enabled by default and preset global parameters are applied to all GE interfaces.

- The Brocade Extension platform supports LLDP on LAN and WAN side GE interfaces.
- LLDP is an Ethernet Layer 2 protocol and operates only on the data-link, it is not IP-based.
- LLDP advertises and receives network switch and port identities specific to the Ethernet link.

LLDP uses keepalives on the GE interfaces. These keepalives should not be confused with circuit Keepalive Timeout Values (KATOV). The following table shows the LLDP timeout values and transmission intervals:

Table 14: LLDP Timeout Values

Protocol	Default Hello Interval	Hello Interval Range	Minimum Timeout	Maximum Timeout	Multiplier (mx)
LLDP	30 Seconds	4 - 180 Seconds	8 Seconds (min hello x min mx)	1800 Seconds (max hello x max mx)	4 Range is 2 to 10 Equal to 120 Seconds at default hello interval.

NOTE
LLDP traffic is not impacted during the high availability (HA) reboot.

Perform the following steps to configure the LLDP global parameters:

1. Enable the LLDP protocol on the switch.
```
switch:admin> lldp --enable
```
2. Configure the LLDP global parameters.
```
switch:admin> lldp --config -sysname My_Ext_Switch

switch:admin> lldp --config -sysdesc 7810_on_default_vf
```
3. Enable the required TLVs (Type Length Values) using the `lldp --enable -tlv` command. Disable DCBX and FCoE TLVs such as dcbx, fcoe-app, and fcoe-lls, and these TLVs are not used with extension.
   The following TLVs are supported:
   - **dot1** – IEEE 802.1 organizationally specific
   - **dot3** – IEEE 802.3 organizationally specific
   - **mgmt-addr** – management address
   - **port-desc** – port description
   - **sys-name** – system name
   - **sys-desc** – system description
   - **sys-cap** – system capabilities

   The TLV port-desc string is in the following format:
   Switch Model Name (fixed-port switch)/Slot Model Name (case of chassis): Mode + Speed + Slot/Port
   For example: `Brocade 7810: WAN 10G ge3`
   The TLV `sys-desc` , the default system description string is advertised is in the following format:
```
Switch Model Name, Firmware Version
```
   For example: `Brocade 7810, Fabric OS Version 9.2.0`
4. Use the following command to display the LLDP neighbors and statistics:

```
switch:admin> lldp --show -nbr
Local Intf    Dead Interval   Remaining Life   Remote Intf   Chassis ID     Tx     Rx      System Name
ge2           120             111              ethernet1/1/6 684f.6451.067d 715577 714998  my-dmzs-sw01
ge3           120             99               ethernet1/1/6 684f.647a.68f5 715573 715167  my-dmzs-sw02
ge4           120             112              ethernet1/1/7 684f.6451.067d 714744 714168  my-dmzs-sw01
ge5           120             113              ethernet1/1/7 684f.647a.68f5 714728 714329  my-dmzs-sw02
```

Configuring and Activating an LLDP Profile for a Group of Ports

LLDP is enabled by default. Extension GE interfaces do not require a specific LLDP profile; global parameters are adequate. Multiple customized LLDP profiles with different parameters can be applied to specific groups of GE interfaces if required.

NOTE
You can have up to 512 LLDP profiles in a switch or chassis.

1. Create an LLDP profile using the `lldp --create -profile` command.
   ```
   switch#admin> lldp --create -profile lldp_profile_2
   ```
2. Configure the LLDP profile parameters using the `lldp --config` command.
   ```
   switch#admin> lldp --config -mx 4 -profile lldp_profile_2
   switch#admin> lldp --config -txintvl 40 -profile lldp_profile_2
   ```
3. Enable the required TLVs on the LLDP profile using the `lldp --enable -tlv` command.
   ```
   switch#admin> lldp --enable -tlv dot3 -profile lldp_profile_2
   ```
4. Use the `lldp --show -profile` command to display the configured LLDP profile parameters.
   ```
   switch#admin> lldp --show -profile lldp_profile_2
   ```
5. Enable the LLDP profile on a group of ports using the following command:
   ```
   switch#admin> lldp --enable -port ge13-14 -profile lldp_profile_2
   ```
   A group of ports can now be expressed as a range rather than as a string of ports that are separated by commas as shown in the following example:
   ```
   switch#admin> lldp --enable -port 3/40-56 -profile lldp_profile_1
   ```
6. Verify the LLDP profile details.
   ```
   switch:admin> lldp --show -profile

   Profile-name:lldp_profile1
     Enabled TLVs:dot1;mgmt-addr;
     Profile ports: ge13;ge14
   ========================================

   Profile-name:lldp_profile2
     Enabled TLVs:dot3;sys-desc;mgmt-addr;
     Profile ports: ge16;ge17
   ========================================

   Number of profile entries = 2
   ```
7. View the LLDP stats for a GE interface.
   ```
   switch:admin> lldp --show -stats ge4
   LLDP Interface statistics for ge4
   Frames transmitted: 979387
   Frames Aged out: 0
   Frames Discarded: 0
   Frames with Error: 979404
   Frames Recieved: 978578
   TLVs discarded: 5876469
   TLVs unrecognized: 979413
   ```

VE_Ports

VE_Ports are specific to extension and apply to all three sides; FC/FICON, WAN, and LAN. VE_Port is a port in an FC or FICON fabric or logical switch, and it is a tunnel endpoint on WAN and LAN sides. VE_Ports locally designate an extension tunnel. Because extension platforms support multiple VE_Ports, multiple tunnels can be created. The number of tunnels depends on the VE mode of the platforms (Brocade 7850 Extension Switch and Brocade SX6 Extension Blade) and applied licenses (Brocade 7810 Extension Switch). There is no requirement for VE_Port numbers to be identical at each end of a tunnel. If a VE_Port is disabled, the tunnel is brought down. Disabling and enabling a VE_Port bounces the tunnel.

Brocade 7850 Extension Switch 6VE and 18VE Mode

When operating in 6VE mode, all 3 VE ports available to each data processor (DP) are part of their own port group. When operating in 18VE mode; however, the VE ports are grouped into three ports each. There are three groups of three ports on each DP. Bandwidth caps are enforced on a per-port-group basis, similar to 20VE mode on the Brocade SX6 Extension Blade.

When operating in the 6VE mode the bandwidth limits for the tunnel are capped at a total of 50Gbps for the aggregate of all circuits under that tunnel. Each circuit has a maximum rate of 25Gbps. Therefore, to achieve the maximum bandwidth of a tunnel, at least two circuits must be used. This 50Gbps limit will also be the cap for the aggregate of all VEs for a given DP. This 50Gbps limit will also be the cap for the aggregate of all VEs for a given DP.

In 18VE mode, the bandwidth limit is 16Gbps for a tunnel, which is also the cap for the entire VE group. The maximum total throughput per DP when operating in 18VE mode is 48Gbps (3 groups x 16Gbps per group).

These bandwidth limits are enforced. If a user attempts to configure beyond the limits, an appropriate error message is displayed. These limits also are applicable to the maximum committed rate, which means the user cannot oversubscribe a DP beyond the 48Gbps bounds. This limit differs from the current implementation on the Brocade SX6 Extension Blade where the maximum DP WAN side rate is 20Gbps, and 2 to 1 oversubscription to 40Gbps is permitted.

Irrespective of the VE mode, the minimum bandwidth limit for a circuit increased to 50Mbps. When a Brocade 7850 Extension Switch tunnel is connected to a Brocade 7810 Extension Switch or Brocade SX6 Extension Blade, the circuit fails to come up if the peer is configured lower than 50Mbps, those platforms support a minimum rate of 20Mbps. If they are configured for any speed less than 50Mbps, the circuit fails to establish a connection and reports a configuration warning with the min-comm-rate flag set. A peer switch, such as a Brocade 7810 Extension Switch or Brocade SX6 Extension Blade, which attempts to modify the bandwidth of a tunnel below 50Mbps will be rejected with an error indicating that the bandwidth is too low and is not supported.

The following VEs are some examples of valid aggregate circuit configurations per tunnel:

- 6VE Mode (all VEs in a DP are in a single group):
  - DP0: VE24 - 20G
  - DP0: VE25 - 20G
  - DP0: VE26 - 10G
  - DP1: VE33 - 50G
- 18VE Mode (three groups of 3 VEs per group):
  - DP0-Group1: VE24 - 16G
  - DP0-Group2: VE27 - 16G
  - DP0-Group3: VE30 - 16G
  - DP1-Group1: VE33 - 16G
  - DP1-Group2: VE36 - 16G
  - DP1-Group3: VE39 - 16G

Table 15: Extension Platform VE Modes

Extension Platform	VE Mode	DP and VE_Ports
Brocade 7810 Extension Switch	Not applicable	DP0:12-15
Brocade 7850 Extension Switch	6VE • VE max is 50Gbps • Each DP max of 50Gbps • 1 group of 3 VEs per DP	DP0: 24-26 DP1: 33-35
	18VE • VE max is 16Gbps • Each VE group max of 16Gbps • 3 groups of 3 VEs per DP	DP0 Group 1: 24-26 DP0 Group 2: 27-29 DP0 Group 3: 30-32 DP1 Group 1: 33-35 DP1 Group 2: 36-38 DP1 Group 3: 39-41
Brocade SX6 Extension Blade	10VE	DP0: 16-20 DP1: 26-30
	20VE	DP0: 16-25 DP1: 26-35

FC and FICON Side VE_Port

FC communicates between platforms using inter-switch links (ISL). ISLs are supported across various port types, such as E_Port, EX_Port, VE_Port, and Inter-chassis links (ISLs). VE_Ports are like E_Ports, except they use TCP/IP as their transport instead of FC. Extension tunnel is a type of ISL that use VE_Ports. Through its VE_Port, FC frames enter into DP of an extension tunnel where they undergo the high-efficiency encapsulation (HEE) and are passed to the TCP layer. A DP of an extension platform handles the encapsulation, compression, encryption, and transmission of FC frames.

A VE_Port must be present in the LS where the storage devices are connected. Often, Logical Fabrics are used to isolate VE_Ports from each other to maintain deterministic paths for sequences of an exchange.

A FC Routed (FCR) port is referred to as EX_Port. In an edge fabric, EX_Ports represent the demarcation point between fabric services. Fabric services do not extend beyond an EX_Port.

The following table describes the types of VE_Port in a fabric:

Table 16: VE_Port Types

E_Port Types	No FCR	FCR
Native FC	E_Port	EX_Port
Extension	VE_Port	N/A (VEX_Port has been deprecated)

WAN Side VE_Port

On the WAN side, a VE_Port is an endpoint of a tunnel for one or more circuits. A BET logically forms a single tunnel due to the use of a single VE_Port at each end, even though the tunnel is made up of multiple circuits. Once a tunnel is configured and the circuits are up, a logical ISL is established between the switches.

LAN Side VE_Port

On the LAN side, the VE_Port is the target ID used by a TCL to match traffic. There is no concept of Virtual Fabrics Logical Switches on the LAN side.

NOTE
VE_Ports cannot be connected to the same domain simultaneously with E_Ports or EX_Ports.

Persistently Disabling and Enabling VE_Ports

NOTE
It is not necessary to disable the VE_Port during the extension configuration; however, if the VE_Port remains enabled, fabrics should not be connected to the extension platform because the tunnel forms an ISL with the remote side during configuration. When a configuration is initially altered, ISL connectivity might be unstable, and this instability might spread to the attached fabric.

By default, the VE_Port is enabled. VE_Ports can be manually changed from persistently enabled to persistently disabled. While tunnel configuration is in progress, change manually to prevent unwanted fabric merges. After a tunnel has been fully configured on both ends, persistently enable the VE_Ports. Persistent disabled ports remain disabled across power cycles, switch reboots, and switch enables.

NOTE
The `portcfgpersistentdisable` or `portcfgpersistentenable` commands cannot be used with FMS mode enabled. Use the `portdisable` and `portenable` commands instead during configuration.

Perform the following steps to enable or disable the VE_Ports:

1. Connect to the switch and log on as admin.
2. Use the `portcfgpersistentdisable` command to disable any VE_Ports that you use in the tunnel configuration. This example disables VE_Port 7/24. There is no additional response or confirmation after executing the command.

```
switch:admin> portcfgpersistentdisable 7/24
switch:admin>
```

3. Use the `portshow` command to show the status of a port. This example shows the status of VE_Port 7/24. You can see the status of persistent disable in the output.

```
switch:admin> portshow 7/24

portIndex: 24
portName: slot7 port24
portHealth: Not Monitored

Authentication: None
portDisableReason: Persistently disabled port
portCFlags: 0x0
portFlags: 0x4021    PRESENT VIRTUAL U_PORT DISABLED LED
LocalSwcFlags: 0x0
portType: 12.0
portState: Persistently Disabled
Protocol: FC
portPhys: 255 N/A    portScn:    2    Offline
port generation number:    24
state transition count:    20

portId:  0b1800
portIfId:     43020817
portWwn: 20:18:50:eb:1a:13:ad:16
portWwn of device(s) connected:
Distance: normal
Port part of other ADs: N/o
```

4. When the FMS mode is enabled for FICON, use the following steps to disable a port.
   a. Use the `ficoncupshow` command to show the FMS mode. The following example shows that the FMS mode is enabled.

```
switch:admin> ficoncupshow fmsmode
FMS_001(I) - FMSMODE for the switch: Enabled
```

   b. Use the `portdisable` command to disable a port when the FMS mode is enabled. This example disables port 7/24, a VE_Port in a Brocade SX6 Extension Blade.

```
switch:admin> portdisable 7/24
switch:admin>
```

   There is no additional response or confirmation after entering the command.
5. Use the `portshow` command to show the port status.
   The following example shows that the port is offline and disabled:

```
switch:admin> portshow 7/24 portIndex: 24
portName: slot7 port24
portHealth: Not Monitored
Authentication: None
portDisableReason: None
portCFlags: 0x0
portFlags: 0x4021    PRESENT VIRTUAL U_PORT DISABLED LED
LocalSwcFlags: 0x0
portType: 12.0
portState: 2    Offline
Protocol: FC
portPhys: 255 N/A portScn: 2 Offline
port generation number: 26
state transition count: 22

portId: 0b1800
portIfId: 43020817
portWwn: 20:18:50:eb:1a:13:ad:16
portWwn of device(s) connected:
```

A persistently enabled or disabled VE_Port remains enabled or disabled following a reboot. To configure persistently enabled and disabled VE_Ports, use the commands `portcfgpersistentenable` and `portcfgpersistentdisable` respectively.

To perform the configuration, follow these steps:

1. Connect to the switch and log on as admin.
2. To show the status of a port, use the `portshow` or `switchshow` command. This example shows the status of VE37. You can see the persistent disable status in the output.

```
switch:admin> portshow 37

portIndex:  37
portName: port37
portHealth: Not Monitored

Authentication: None
portDisableReason: Persistently disabled port
portCFlags: 0x0
portFlags: 0x4021        PRESENT VIRTUAL U_PORT DISABLED LED
LocalSwcFlags: 0x0
```

```
portType:  12.0
portState: Persistently Disabled
Protocol: FC
portPhys:  255  N/A     portScn:   2    Offline
port generation number:    0
state transition count:    1

portId:    502500
portIfId:    43020830
portWwn:   20:25:88:94:71:a5:c4:30
portWwn of device(s) connected:
16b Area list:
Distance:  normal
```

3. Use the `portcfgpersistentenable` command to enable the VE_Port used in a tunnel configuration; this example enables VE12.

```
switch:admin> portcfgpersistentenable 12
switch:admin>
```

   After entering the command, there is no additional response or confirmation.

4. Use the `portshow` command to verify the status of a port.

   This example shows the status of VE37. You can see that the port is offline and no longer disabled. If the port was configured, it would be in an online state.

```
switch:admin> portshow 37

portIndex:  37
portName: port37
portHealth: Not Monitored

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x4001        PRESENT VIRTUAL U_PORT LED
LocalSwcFlags: 0x0
portType:  12.0
portState: 2    Offline
Protocol: FC
portPhys:  255  N/A     portScn:   2    Offline
port generation number:    0
state transition count:    1

portId:    502500
portIfId:    43020830
portWwn:   20:25:88:94:71:a5:c4:30
portWwn of device(s) connected:
16b Area list:
Distance:  normal
```

5. If the platform is in the FMS mode, use the `portenable` command to enable or disable a VE_Port. The following example enables port 7/24, a VE_Port on a Brocade SX6 Extension Blade. There is no additional response or confirmation after entering the command.

```
switch:admin> portenable 7/24
switch:admin>
```

Gigabit Ethernet Interface Speed

On the Brocade Extension platforms, the gigabit Ethernet interfaces are abstracted from the VE_Ports, which means that the GE interfaces are not fixed to any set of VE_Ports. Depending on the platform, you have access to a combination of 1GbE, 10GbE, 25GbE, 40GbE, and 100GbE interfaces.

A gigabit Ethernet interface can be configured for the WAN side or the LAN side. A GE interface is used for WAN operation when it functions as a circuit endpoint. A WAN side IPIF IP address is a circuit endpoint, and assigned to the desired GE interface and the DP of the VE_Port to be used.

GE ports are mapped to circuits using IP interfaces (IPIFs). By designating the VE_Port to which the circuit belongs, the circuit is grouped into an extension trunk. Tunnels and trunks are identified by their VE_Port number. Once a tunnel exists, you can create circuits and add them to the tunnel.

Each WAN-side Ethernet interface on the Brocade Extension platforms can be used by circuits from multiple VE_Ports. Multiple VEs in different FIDs can share a GE port if the GE interface remains in the default logical switch (FID 128). If a GE port is moved to a non-default switch, the GE interface can only be used by VE_Ports in that FID. Best practice is to keep the GbE interfaces in the default logical switch (FID 128).

Each VE_Port can have multiple circuit endpoints, which are IP interfaces (IPIF). Each circuit in a VE_Port is associated with an IPIF and GE interface. Each VE_Port can use multiple Ethernet interfaces. A circuit cannot span multiple Ethernet interfaces.

For BET, the DP owning the VE_Port controls all member circuits. There is no distributed processing, load sharing, or LLL across different DPs. FCIP failover from one DP to another is done at the FC level by the Brocade FC switching ASIC, provided there is more than one DP and the configuration permits it. The only components that can be shared are the Ethernet interfaces.

IP networks can route circuits over separate paths based on the destination IP addresses, VLAN tagging, QoS marking, and other Layer 2 and Layer 3 attributes. Ethernet interfaces on the Brocade Extension platforms provide connectivity to the network for one or many circuits.

The Brocade Extension platform provides the following GE interfaces:

- The Brocade 7810 Extension Switch has up to six GE/10GE SFP/SFP+ interfaces or two 1GbE copper interfaces and four GE/10GE SFP/SFP+ interfaces
- The Brocade 7850 Extension Switch has 16 GE/10GE/25GE and two 100GE interfaces
- The Brocade SX6 Extension Blade has 16 GE/10GE and two 40GE interfaces

There is no difference in speed configuration whether the GE interface is used for WAN or LAN. Interface speed is set using the `portcfgge --set -speed <speed>` command. The following steps are examples of configuring interface speed on a Brocade SX6 Extension Blade.

Perform the following steps to configure interface speed on a Brocade SX6 Extension Blade:

> **NOTE**
> Disregard the slot designation for the Brocade 7850 Extension Switch and Brocade 7810 Extension Switch platforms.

1. Connect to the switch and logon as admin.
2. Perform one of the following steps:
   - Enter the following command to set the speed at 1Gb/s for interface 7/ge4 (slot7 GE4):
     ```
     switch:admin> portcfgge 7/ge4 --set -speed 1G
     ```
   - Enter the following command to set the speed at 25Gb/s for interface ge0:
     ```
     switch:admin> portcfgge ge0 --set -speed 25G
     ```
   - Enter the following command to disable autonegotiation on slot 7 interface ge4:

```
switch:admin> portcfgge 7/ge4 --disable -autoneg
```

– Enter the following command to enable autonegotiation on slot 7 interface ge4:

```
switch:admin> portcfgge 7/ge4 --enable -autoneg
```

3. Enter the following to display the current speed for interface ge0:

```
switch:admin> portcfgge ge0 --show
Port    Speed    Flags    Lag ID    Channel

---------------------------------------
ge4    25G    ----    ----    N/A
---------------------------------------

Flags: A:Auto-Negotiation Enabled C:Copper Media Type
L:LAN Port G: LAG Member

switch:admin> portshow ge0
Eth Mac Address: 00.05.33.65.83.57
Port State: 1    Online
Port Phys: 6    In_Sync
Port Flags: 0x3 PRESENT ACTIVE
Port Speed: 25G
```

1GE Interface Auto-Negotiation

Auto-negotiation applies only to 1GE interfaces and is standards based for 1GbE PHY. It is not a speed negotiation. The Brocade GE interfaces can be set to either 1GE, 10GE, or 25GE mode and requires the corresponding SFP (1GbE) or SFP+ (10GbE). An interface set to auto-negotiate is determining with its peer switch if communications are full-duplex or half-duplex, and if pause frames (802.3X) are used, these parameters do not need to match on each end. However, the interface does not come up if there is a mismatch in the auto-neg setting itself, meaning that one side has auto-neg enabled and the other side does not.

By default, 1GE interfaces have auto-negotiation enabled. Auto-negotiation is not supported at speeds higher than 1GbE. If the data center LAN switch needs auto-negotiation disabled, the `portcfgge ge# --disable -autoneg` command disables it.

NOTE

- The Brocade 7810 Extension Switch GE0 and GE1 are equipped with copper interfaces (RJ-45) that support only 1GbE speeds.
- Brocade 7850 Extension Switches do not support auto-negotiation at 1GbE speeds.
- Changing the speed of an interface with active circuits is disruptive.

VM Insight

VM Insight is a feature that allows flows between virtual machines (VMs) to be tracked and identified with a greater granularity. VM instances can include an optional header that contains a VMID that identifies the specific VM instance that is initiating the I/O. The VMID header is supported with all emulated FCP I/O sequences. When the server adds the optional VMID header, the header is used between the FCIP complex and the target device. It is also used in FCP sequences from the target device to the initiator.

Drivers for HBAs and storage vendors must support the VMID frame-tagging method, which identifies each subflow through the fabric to support the VM Insight feature.

VM Insight supports the following features:

Table 17: VM Insight Support

Tunnel Type	Brocade 7810 Extension Switch	Brocade 7850 Extension Switch	Brocade SX6 Extension Blade
Non-emulated tunnel	Supported	Supported	Supported
FastWrite and OSTP	Supported	Supported	Supported
Advanced FICON Accelerator	Not supported	Not supported	Not supported

For more information, refer to the *Brocade Fabric OS Administration Guide* and the *Brocade Fabric OS MAPS User Guide*.

NVMe over Extension

NVMe (Non-Volatile Memory Express) is supported on all Brocade FC and Extension platforms. No special configuration tasks are required.

IPv6 Addressing

The Brocade Extension WAN side implementation of IPv6 uses unicast addresses for the circuit IP interfaces. The link-local unicast address is automatically configured on the interface; however, using the link-local address space for a circuit endpoint is not permitted. Site-local unicast addresses are not allowed for circuit endpoints.

> **NOTE**
> IPv6 addresses can exist with IPv4 addresses on the same interface, but each circuit must be configured as IPv6-to-IPv6 or IPv4-to-IPv4. IPv6-to-IPv4 circuits are not supported.

The following considerations should be taken into account when addressing IPv6:

- IPv6 IPsec is supported.
- IPv6 PMTU (Path MTU) discovery is supported.
- IPv6 interfaces have unique unicast addresses.
- IPv6 addresses and routes must be statically configured by the user.
- IPv6 interfaces cannot be configured with anycast addresses.
- IPv6 interfaces cannot be configured with multicast addresses.
- The IPv6 8-bit Traffic Class field is defined by the configured Differentiated Services field for IPv6 (RFC 2474).
- The DSCP (Differentiated Services Code Point) marking configuration is done per circuit using the portcfg fcipcircuit DSCP arguments.
- IPv6 optional Extension Headers are not supported.
- IPv6 router advertisements and stateless address auto-configuration (RFC 2462) are not supported.
- IPv6 Neighbor Discovery (ND) and RFC 4443 ICMPv6 message types are supported.
- Brocade Extension platforms use portions of the Neighbor Discovery protocol (RFC 4861).
  - Hop limits, such as Time to Live (TTL), are learned from Neighbor Advertisements.
  - Neighbor's link-local addresses are learned from Neighbor Advertisements.
  - Each GE interface’s IPv6 link-local address is automatically set at startup and advertised to neighbors.
  - The user does not configure an interface’s link-local address.

IPv6 with Embedded IPv4 Addresses

When using IPv6 within an IPv4 network, only IPv4-compatible IPv6 addresses are supported. Only the low-order 32 bits of the address can be used as an IPv4 address (the high-order 96 bits must be all zeros). This allows IPv6 addresses to be used on an IPv4 routing infrastructure that supports IPv6 tunneling over the network. Both endpoints of the circuit must be configured with IPv4-compatible IPv6 addresses. IPv4-to-IPv6 connections are not supported on Brocade Extension

platforms. IPv4-mapped IPv6 addresses are not supported because they are intended for nodes that support IPv4 only when mapped to an IPv6 node.

Configuring Extension Platform Modes

The available modes depend on the platform and include Application mode (FCIP only or Hybrid), VE mode, and GE mode. The platform modes depend on the platform and its specific modes:

- **App mode**: The Brocade SX6 Extension Blade is the only platform with Application mode (FCIP or Hybrid)
- **VE mode**: The Brocade 7850 Extension Switch and Brocade SX6 Extension Blade have VE modes, and the Brocade 7810 Extension Switch has no VE mode
- **GE mode**: The Brocade 7810 Extension Switch is the only platform with GE mode

NOTE
When you remove a Brocade SX6 Extension Blade or change its slot, configuration information for the blade must be cleared.

Brocade SX6 Extension Blade

The Brocade SX6 Extension Blade contains the following modes:

- **App mode**: FCIP Only or Hybrid
- **VE mode**: 10VE or 20VE
- **GE mode**: Brocade SX6 Extension Blade has no GE mode

Brocade 7850 Extension Switch

The Brocade 7850 Extension Switch contains the following modes:

- **App mode**: Hybrid mode only, the Brocade 7850 Extension Switch has no Application mode
- **VE mode**: 6VE or 18VE
- **GE mode**: The Brocade 7850 Extension Switch has no GE mode
  The Brocade 7850 Extension Switch does not have an Application or GE mode. It supports two VE modes, which controls the total available DP bandwidth relative to the number of VE_Ports.

Brocade 7810 Extension Switch

The following modes are available in the Brocade 7810 Extension Switch:

- **App mode**: Hybrid mode only, Brocade 7810 Extension Switch has no Application mode
- **VE mode**: Brocade 7810 Extension Switch has no VE mode
- **GE mode**: (ge0-ge1 RJ-45) or (ge2-ge3 optics)
- **Copper mode**: (ge0-ge1 RJ45 + ge4-ge7 SFP)
- **Optical mode**: (ge2 - ge7 SFP)

Brocade 7810 Extension Switch has no Application or VE mode. The base unit has two VE_Ports and the full/upgraded unit has four VE_Ports.

Configuring the WAN Side (Tunnel)

This section describes how to configure extension GE interfaces, IP interfaces, IP routes, tunnels, circuits, and various optional features.

GE Interfaces

Ensure that the GE interface speed configuration is complete before creating IPIF, tunnels, and circuits.

NOTE

- An extension GE interface can be set for either WAN or LAN. GE interfaces default to FCIP, which is a legacy designation for WAN. Even though the GE interface might show FCIP, it is a WAN side interface that supports both FCIP and IP Extension. Output from a `switchshow` command shows FCIP as the protocol when a GE interface is configured for WAN.
- GE interfaces are not associated with any particular VE_Ports.

GE interfaces per platform:

Table 18: GE Interfaces and Speeds per Platform

Platform	GE (SFP)	10GE (SFP+)	25GE (SFP+)	40GE (QSFP)	100GE (QSFP)
Brocade 7810 Extension Switch*	2: ge0-ge1 (RJ-45) 6: ge2-ge7	6: ge2-ge7	—	—	—
Brocade 7850 Extension Switch	16: ge0-ge15	16: ge0-ge15	16: ge0-ge15	—	2: ge16-ge17
Brocade SX6 Extension Blade	16: ge2-ge17	16: ge2-ge7	—	2: ge0-ge1	—

* Brocade 7810 Extension Switch ge0-ge1 and ge2-ge3 are enabled as a set, and the two sets are mutually exclusive.

Perform the following steps to configure if a GE interface is not already set for WAN:

1. Enter the `portcfgge` command with the `--help` option:

```
switch:admin> portcfgge --help

Usage: portCfgGe [<slot>/][<port>] [<args>]

Port Format:
   ge#

Args:
   --enable -autoneg         - Enable the auto-negotiate mode of the 1 GE ports only.
   --disable -autoneg        - Disable the auto-negotiate mode of the 1 GE ports only.
   --set -speed <speed>      - Set the port speed of the GE ports.
                               Allowable speeds are [1G] [10G]
   --set -lan                - Set the port as lan.
   --set -wan                - Set the port as wan.
   --set -channel <channel>  - Set the port tunable SFP channel ID of the 10 GE ports only.
                               Allowable channel ID Range [1] - [102]
```

```
  --show                    - Show the current GE port configurations.
  --show -lmac              - Show the Local MAC addresses for GE/LAN ports.
```

2. Show which GE interfaces are set for speed and WAN or LAN. LAN is indicated, no indication of LAN means it is a WAN side GE interface.

```
switch:admin> portcfgge --show

 Port         Speed     Flags     Channel  Lag Name
-------------------------------------------------------------------------
 ge0          1G        AC--      N/A      -
 ge1          1G        AC--      N/A      -
 ge2          10G       ----      N/A      -
 ge3          10G       ----      N/A      -
 ge4          10G       ----      N/A      -
 ge5          10G       ----      N/A      -
 ge6          10G       --L-      N/A      -
 ge7          10G       --L-      N/A      -
-------------------------------------------------------------------------

Flags: A:Auto-Negotiation Enabled  C:Copper Media Type
L:LAN Port G: LAG Member
```

3. Set the GE interface to be a WAN side interface.

```
switch:admin> portcfgge ge6 --set -wan
Operation Succeeded.
```

4. Enter the `portCfgGe --show -lmac` command to view the MAC addresses for the GE interfaces:

```
switch:admin> portcfgge ge6 --show -lmac

 Port       dpid    MAC Address
-------------------------------------------------------------------------
 ge0        dp0     88:94:71:60:99:7f
 ge1        dp0     88:94:71:60:99:80
 ge2        dp0     88:94:71:60:99:81
 ge3        dp0     88:94:71:60:99:82
 ge4        dp0     88:94:71:60:99:83
 ge5        dp0     88:94:71:60:99:84
 ge6        dp0     88:94:71:60:99:85
 ge7        dp0     88:94:71:60:99:86
 lan        dp0     88:94:71:60:99:87
-------------------------------------------------------------------------
```

IP Interfaces (IPIF)

An IP Interface (IPIF) is the endpoint of an extension circuit. For each circuit that you want to create, an IPIF must be configured. The IPIF is created before circuits are configured, so plan out the circuits and their associated IPIFs in advance. If you intend to use eHCL, you might also must configure the eHCL IPIFs.

Each IPIF is associated with a GE interface and DP. The Brocade 7850 Extension Switch and Brocade SX6 Extension Blade have DP0 and DP1, while the Brocade 7810 Extension Switch has only DP0.

Part of configuring an IP address to an IPIF is assigning the GE interface that it operates on, and in turn, the interface the circuit uses. The importance of selecting proper GE interfaces is that circuits physically connect to a specific data center LAN switch port. Port, optic, cable, and data center switch redundancy and resiliency are important considerations for circuits and physical connectivity. Circuits can be assigned to any physical connection. Circuits cannot span more

than one physical connection. Tunnels span multiple physical connections by using multiple circuits assigned to different connections.

Local IP and local HA IP addresses are configured using the `portcfg ipif` command on each platform. When configuring a circuit, either with the `fciptunnel` command or the `fcipcircuit` command, the local and remote IP addresses of the circuit are specified, and optionally the local HA IP and remote HA IP. IPIFs cannot be assigned to circuits until they are configured as IPIFs; otherwise, the CLI displays the error message '*Object Does Not Exist*'.

To configure an IPIF, use the `portCfg ipif create` command. The IPIF consists of an IP address, a netmask, an IP MTU size, and other options depending on the extension switch or blade. To modify an existing IP address, use the `portCfg ipif modify` command.

The maximum number of IPIF supported on each platform:

Table 19: Maximum Number of IPIF per DP and GE

Extension Platform	Maximum IPIF per DP	Maximum IPIF per GE
Brocade 7810 Extension Switch	60	60
Brocade 7850 Extension Switch	60	64
Brocade SX6 Extension Blade	60	64

NOTE
Each IP address pair (local and remote) must be unique. For example, the following address pairs use the same source address in each pair, but the destination addresses differ and each address pair is unique.

- --local-ip 10.0.1.10 --remote-ip 10.1.1.10
- --local-ip 10.0.1.10 --remote-ip 10.1.1.11

IPIF configuration requires the following requirements and options:

- A tunnel can have a mixture of IPv4 and IPv6 circuits.
- A circuit must be either IPv4 or IPv6 at both ends; endpoints cannot be a mix of IPv4 and IPv6.
- Classless Inter-Domain Routing (CIDR) notation is supported for IPv4 and IPv6 addresses.
- Specifying an IP MTU is optional; if not specified, the default is 1500 bytes.
  - The maximum supported MTU is 9216 bytes.
  - The minimum supported MTU is 1280 bytes.
- MTU can be set to manual or AUTO.
- Auto uses a Path MTU (PMTU) discovery to determine the MTU.
- CIDR subnet mask of 31 is supported.
  - An IPv4 address with a 31-bit subnet mask has no network and broadcast addresses.

31-bit subnet mask example, there is a one-bit difference in the following two addresses:

- Address A: 192.0.2.0/31
- Address B: 192.0.2.1/31

Using 31-bit subnet mask IP addresses, as defined in RFC 3021, reduces the number of IP addresses consumed by networking devices for a point-to-point circuit connectivity. 31-bit subnets allow the all-zeros and all-ones IP addresses to be assigned as working addresses on point-to-point connections. Prior to RFC 3021, the longest subnet mask in common use on point-to-point links was 30-bits. This length meant that the all-zeros and all-ones IP addresses were wasted, so 50% of IP addresses went to waste.

1. Connect to the switch and log on as admin.
2. Use the `portCfg ipif create` command to create an IPIF.

The following example shows an IPIF created on a Brocade SX6 Extension Blade in slot 4 and GE14 of a Brocade X7 Director. You can use either a CIDR length or a netmask value for IPv4 subnets masking.

Any MTU larger than the default 1500 bytes is considered a jumbo frame.

```
switch:admin> portcfg ipif 4/ge14.dp0 create 192.1.2.2/24
Operation Succeeded.
```

The following is an example of a Brocade 7810 Extension Switch or Brocade 7850 Extension Switch:

```
switch:admin> portcfg ipif ge3.dp0 create 192.168.10.10/24

Operation Succeeded.
```

3. To delete an IP interface, use the `portCfg ipif delete` command. To delete an IPIF, the IP address cannot be in use on a circuit. The following is an example of the `delete` command on the Brocade SX6 Extension Blade:

```
switch:admin> portcfg ipif 8/ge6.dp1 delete 192.1.2.2
```

NOTE
An IPIF cannot be deleted if a circuit is configured using the IP address. The I flag indicates that the IPIF is in use by a circuit.

4. Enter the following command to show the configured IPIFs:

```
switch:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.10.1                  / 30   9216  0     U R M I
 ge2.dp0      10.0.0.5                      / 24   1500  0     U R M
 ge3.dp0      192.168.10.10                 / 24   1320  0     U R M I
 lan.dp0      10.10.10.123                  / 24   1500  0     U R M
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

IPIF VLAN Tagging (802.1Q)

When a VLAN ID (802.1Q) is configured on an IPIF, all traffic is tagged with that VLAN ID (802.1Q). If no VLAN ID is configured on the IPIF, circuit traffic is not tagged. The peer Ethernet interface of the data center must be configured the same; it is either tagging (a trunk port) or not tagging (an access port). Access ports are the most common method of connecting extension GE interfaces to data center switches.

The Brocade Extension platforms use the `portcfg ipif` command to add, delete, or change a VLAN ID to an existing IPIF. Also, it can be added at the time of creation.

1. Connect to the switch and log on as admin.
2. Use the `portcfg ipif` command to create an IPIF with a VLAN ID information to an IPIF. The following example creates an IPIF with VLAN 200:

```
switch:admin> portcfg ipif ge2.dp0 create 192.0.2.0/24 vlan 200
Operation Succeeded
```

3. Use the `portshow ipif` command to display the IPIF information.

```
switch:admin> portshow ipif

Port          IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
ge2.dp0       192.0.2.0                     / 24   1500  200   U R M
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

4. Modify an IPIF to remove a VLAN ID.

```
switch:admin> portcfg ipif ge2.dp0 modify 192.0.2.0/24 vlan none
Operation Succeeded
```

IPIF MTU and PMTU

The Brocade Extension platforms support Jumbo Frames, which refer to the IP datagram maximum transmission unit (MTU) size. The smallest supported MTU is 1280 bytes and the largest supported MTU is 9216 bytes.

The Brocade Extension platforms support path maximum transmission unit (PMTU). PMTU is the process of sending a set of various known size Internet Control Message Protocol (ICMP) datagrams across the IP network to determine the maximum size that was successfully received. The ICMP datagrams are marked Do Not Fragment. Based on the largest ICMP Echo Reply received, the PMTU discovery process sets the IP MTU for that the IP interface (ipif) of the circuit.

PMTU does not discover an MTU greater than 9100 bytes. If the MTU is larger than 9100 bytes, best practice is to manually configure the MTU. If the MTU of the IP network is known, best practice is to set it manually in the `portcfg ipif` command, which avoids values that are determined by PMTU discovery that might be less optimal than the exact network MTU and the additional time that is required to bring up a circuit.

If a circuit bounces, the PMTU discovery process is initiated when attempting to reestablish the circuit. Each circuit initiates the PMTU discovery process before coming online, which is required because circuits might have gone offline due to an IP network failure and rerouted to a new path with different network characteristics. The PMTU discovery process requires additional time when bringing a circuit online.

If ICMP is not permitted on the network, the circuit is established using the smallest MTU (1280).

Enable the PMTU discovery by setting the MTU value to *auto* when configuring the ipif for a circuit using the `portcfg ipif` command. Use the `portcfg ipif` command to display the command arguments and `portshow fcipcircuit --detail` command to display the discovered PMTU value in use. You can also initiate PMTU discovery using the `portcmd --pmtu` command. This runs through the discovery process. However, it runs on the side and does not affect the circuit processing or the discovery of MTUs for circuits.

Perform the following steps to discover the PMTU:

1. Connect to the switch and log on as admin.
2. To modify an IPIF MTU, use the `portCfg ipif modify` command.
   The following example shows how to modify the MTU of an IPIF. Also, the VLAN ID can be modified by adding, deleting, or changing it. If the IPIF is in use, the circuit is bounced when the modifications are made. IP addresses and netmasks cannot be modified, in which case, it is necessary to delete and recreate the IPIF, which is disruptive to the circuit. The tunnel and other circuits remain online.

```
switch:admin> portcfg ipif ge3.dp0 modify 192.0.2.0 mtu 9216

!!!! WARNING !!!!
Modify operation can disrupt the traffic on any tunnel using this IP address. This operation may bring the
 existing tunnel down (if tunnel is up) before applying new configuration.

Note: This operation can take a long time depending on the size or complexity of the configuration. It may
 take up to 3 minutes in some cases.

Continue with Modification (Y,y,N,n): [ n]    y
Operation Succeeded.
```

3. Use the `portshow ipif` command to display the IPIF information.

```
switch:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
```

```
-------------------------------------------------------------------------------
 ge2.dp0      192.0.2.0                    / 30   9216  0     U R M I
 ge2.dp0      192.0.0.5                    / 24   1500  0     U R M
 ge3.dp0      192.0.2.1                    / 24   9216  0     U R M I
 lan.dp0      192.0.2.3                    / 24   1500  0     U R M
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

4. To modify an IPIF MTU to enable PMTU, use the `portCfg ipif modify` command.

```
switch:admin> portcfg ipif ge2.dp0 modify 192.0.2.0/24 mtu auto

!!!! WARNING !!!!
Modify operation can disrupt the traffic on any tunnel using this IP address. This operation may bring the
 existing tunnel down (if tunnel is up) before applying new configuration.

Note: This operation can take a long time depending on the size or complexity of the configuration. It may
 take up to 3 minutes in some cases.

Continue with Modification (Y,y,N,n): [ n]      y
Operation Succeeded.
```

5. Use the `portshow ipif` command to display the IPIF information.

```
switch:admin> portshow ipif

 Port         IP Address                   / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.0.2.1                    / 30   9216  0     U R M I
 ge2.dp0      192.0.0.5                    / 24   AUTO  0     U R M
 lan.dp0      192.0.2.2                    / 24   1500  0     U R M
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

IP Routes

A WAN side IP route is determined by the remote IP address of an extension circuit. If the remote address is not on the same subnet as the local IP address, an IP route must be configured to reach the local gateway. The IP route consists of the destination subnet and the local IP gateway. Local IP gateways, as opposed to IP Extension Gateways, are located on the same subnet as the IPIF addresses on the local WAN side. Specify the destination IP subnet with mask length and the gateway address when creating a route. The gateway is a router interface that forwards traffic towards the destination.

When configuring IP routes, you should consider the following limitations:

- **Brocade SX6 Extension Blade**: Up to 128 routes per GE interface and a maximum of 120 routes per DP can be created.
- **Brocade 7850 Extension Switch**: Up to 128 routes per GE interface and a maximum of 120 routes per DP can be created.
- **Brocade 7810 Extension Switch**: Up to 120 routes per GE interface can be created.

The following example shows an IP route configuration for a Layer 3 routed network. A layer 3 routed network means that the local and remote subnets are different. Notice that on each end, the IPIF IP address and the gateway are in the same subnet. IPIFs communicate to the WAN through a gateway of a router.

Both the GE interface and the DP must be specified when creating IP routes. Circuits are configured for a specific VE_Port and with the configured IP addresses, which in turn puts the circuit on the specific GE interface. DPs own certain VE_Ports. The DP configured for the IPIF, the IP route, and the VE_Port used must all align.

Both the GE interface and the DP must be specified when creating IP routes. For example, specifying the `ge2.dp0` command puts the IP route on ge2 and DP0. The source IP address and VE_Port are not specified during the IP route creation. Therefore, the GE interface and DP must be provided. The exact same route is added to other GE interfaces and DPs, which is necessary when a single subnet is used for multiple tunnels, circuits, and eHCL.

For example, on the Brocade 7850 Extension Switch, using VE24, it is owned by DP0. Therefore, the IPIF created must be the desired GE interface and assigned to DP0, for example, the `portcfg ipif ge2.dp0 create` command. The IP route must be assigned to the same GE interface and DP0, for example, the `portcfg iproute ge2.dp0 create` command. The exception to this is eHCL, where the IPIF and IP route are configured on the opposite DP, for example, the `portcfg ipif ge2.dp1 create` and `portcfg iproute ge2.dp1 create` commands.

Example:

- Side A: Subnet 10.0.0.0/24, IPIF 10.0.0.5, Gateway 10.0.0.1
- Side B: Subnet 192.0.2.0/24, IPIF 192.0.2.3, Gateway 192.0.2.1

Figure 8: Configuring an IP Route

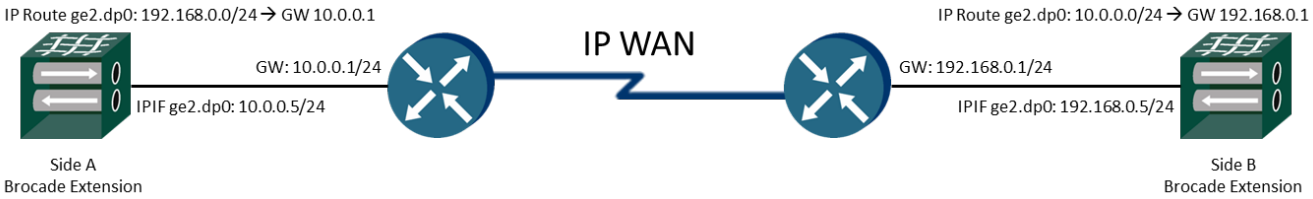

1. Connect to the switch and log on as an admin.
2. Use the `portcfg iproute <slot/ge#.dp#> create` command to create a route to the destination subnet using the gateway of a local router. Omit the `slot/` if the platform is not a director.

   On Side A, the following command creates an IP route to destination subnet 192.0.2.0/24 on interface ge2 for DP0. The route is through local gateway 10.0.0.1. After the destination address, either specify a /pfx (prefix length - shown here) or the word netmask and the subnet mask (that is, netmask 255.255.255.0).

   ```
   switch:admin> portcfg iproute ge2.dp0 create 192.0.2.0/24 10.0.0.1
   Operation Succeeded
   ```
3. Use the `portcfg iproute <ge#.dp#> create` command to create a route on the Brocade 7850 Extension Switch or the Brocade 7810 Extension Switch to the destination subnet using the gateway.

   On Side B, the following command creates an IP route to destination subnet 10.0.0.0/24 for interface ge2 on DP0. From the perspective of the switch, the route uses the local gateway 192.0.2.0.

   ```
   switch:admin> portcfg iproute ge2.dp0 create 10.0.0.0/24 192.168.0.1
   Operation Succeeded
   ```
4. Use the `portcfgshow iproute` command displays the user configured IP routes on Side A.

   The `portshow iproute` command displays more routes than you configured, which might be confusing. Fabric OS adds additional routes for internal use. The user configured routes have the S flag, which denotes the *Static* option. To display the user configured routes, use the `portcfgshow iproute` command.

   ```
   switch:admin> portcfgshow iproute

   Port         IP Address                / Pfx  Gateway                Flags
   ------------------------------------------------------------------------------
   ge2.dp0      192.0.2.0                 / 24   10.0.0.1                S
   ge2.dp1      192.0.2.0                 / 24   10.0.0.1                S
   ge3.dp0      192.0.2.0                 / 24   10.0.0.1                S
   ```

```
ge3.dp1      192.0.2.0               / 24   10.0.0.1              S
lan.dp0      10.85.0.0               / 16   10.10.10.1            S
lan.dp1      10.85.0.0               / 16   10.10.10.1            S
-------------------------------------------------------------------------------
Flags: S=Static X=Crossport
```

NOTE

The lan.dp# routes are not used to forward traffic to the WAN side. They are IP Extension Gateways that are used to forward traffic either directly to end devices (L2) or to an intermediate router (L3) reaching end devices on various different subnets.

5. The `portcfg iproute <ge#.dp#> modify` command cannot alter the destination subnet. To change an IP route subnet or subnet mask, the IP route must be deleted and recreated, which is disruptive to all circuits on the specific GE interface using that IP route. Circuits on other GE interfaces are not affected. If the tunnel has multiple circuits across other GE interfaces, it continues to operate. You cannot use the `modify` option to change the length of the subnet or network mask. Only the gateway address can be altered.

```
switch:admin> portcfg iproute ge2.dp0 modify 192.0.2.0/24 10.0.0.2
Operation Succeeded
```

6. Use the `portcfg iproute <ge#.dp#> delete` command to delete an IP route. Notice that only the destination subnet is required. The gateway at the end is not entered.

```
switch:admin> portcfg iproute ge2.dp0 delete 192.0.2.0/24
Operation Succeeded
```

IPsec Encryption

Internet Protocol Security (IPsec) of in-flight data employs cryptographic security to ensure private and secure communications over IP networks. IPsec supports network-level data integrity, data confidentiality, data origin authentication, and anti-replay protection. It helps secure Extension traffic against network-based attacks from untrusted entities.

The IPsec protocol is invoked as a result of the following sequence of events:

- IPsec encryption has a tunnel level scope, not per circuit.
- IPsec and IKEv2 (Internet Key Exchange) policies are created on both ends of the tunnel.
- IKE exchanges policy information on each end of the connection. If the policy information does not match, the connection does not come online. Some of the exchanged security parameters include the encryption and authentication algorithms, Diffie-Hellman key exchange, and the SAs.
- Data is transferred between IPsec peers based on the IPsec parameters and keys stored in the SA database.
- When authentication and IKE negotiation complete, the IPsec SA is ready for the data transmission.
- SA lifetimes terminate through deletion or expiration. SA lifetime approximately equates the following values:
  - 2 billion frames of data traffic passed through the SA
  - 4-hours expiration time
- When the SA is about to expire, an IKE rekey occurs to create a set of SAs on both ends, after which, data starts using the new SA. IKE and SA rekeys are non-disruptive.

Fabric OS supports a cryptographic suite that is CNSA and Suite B compliant. This enables the configuration of the Elliptic Curve Diffie-Hellman (ECDH) for key agreement and Elliptic Curve Digital Signature Algorithm (ECDSA) for the peer authentication. CNSA and Suite B configuration requires the generation and importation of Suite B compliant X.509 end-entity certificates, and the issuing CA certificates, used to sign them.

The following table shows the algorithm selection for a Suite B compliant configuration as against the configurations supported in pre-Fabric OS 9.2.x releases:

Table 20: IPsec CNSA and Suite B Attributes

Attribute	CNSA and Suite B in Fabric OS 9.2.0
Authentication	ECDSA-P384
Diffie-Hellman	ECDH-P384
Integrity	HMAC-384-192
Encryption	AES-256-CBC (IKE connection) AES-256-GCM (Data connection)
PRF	PRF-HMAC-384

As a part of CSNA and Suite B support, you can use digital certificates and an Enterprise or Public certificate authority (CA) to verify the identity. This requires each end of a secure connection to have access or a means to look up the CA certificate for verification purposes. X.509 Certificate Revocation Lists (CRLs) are not supported.

When you define an IPsec policy, you can select between two profiles. The PSK (preshared key) profile allows you to specify the key when IPsec is configured. The PKI (public-key infrastructure) profile uses certificates.

NOTE
The PKI profile uses certificates. The PKI profile is compliant with the Commercial National Security Algorithm (CNSA) and Suite B.

AES-GCM-ESP: AES (Advanced Encryption Standard), GCM (Galois Counter Mode), ESP (Encapsulating Security Payload) is used as a single, predefined mode of operation for protecting all TCP traffic over a tunnel. AES-GCM-ESP is described in RFC 4106. The following list contains the key features of AES-GCM-ESP:

- Encryption is provided by AES with 256-bit keys.
- Data integrity is provided by GCM with a 128-bit integrity check value.
- IKEv2 key exchange is used by peer platforms for the key agreement.
- IKEv2 uses UDP port 500 to communicate between platforms.
- IKEv2 traffic is protected using AES-256 encryption.
- HMAC (hash message authentication code) is a 192-bit or 128-bit GCM used to check the data integrity and man in the middle tampering.
- PRF (pseudorandom function) is used to generate multiple security keys using 384-bit or 512-bit HMAC.
- MODP (modular exponential) is a 2048-bit or 384-bit Elliptic Curve Diffie-Hellman group that is used for IKEv2 and IPsec key generation.
- No encryption key is used for more than four hours or 2 billion frames, whichever comes first.
  - When an SA expires a new key is generated limiting the amount of time or quantity of data an attacker has to decipher a key.
  - Depending on the length of time expired or the number of frames being transferred, parts of a message might be protected by different keys generated as the SA life expires.
- When the key lifetime of 4 hours or 2 billion frames has been reached, new keys are generated and the old keys are destroyed; the process of rekeying is non-disruptive.
- ESP (Encapsulating Security Payload) is used in transport mode.
- ESP uses a hash algorithm to calculate and verify an authentication value, and it encrypts only the TCP header and its payload.
- Circuits from a non-secure tunnel can use the same Ethernet interfaces as circuits from a secure tunnel.
- Brocade IPsec is a hardware implementation that does not degrade or impact performance.
- Brocade IPsec does not preclude the use of compression or QoS.
- Unlike AES-GCM, AES-CBC does not have an integrated integrity algorithm; therefore, HMAC-384-192 provides integrity.

IPsec Limitations

The usage of IPsec has the following limitations:

- Network Address Translation (NAT) is not supported.
- Authentication Header (AH) is not supported.
- IPsec supports IPv6 and IPv4-based tunnels.
- IPsec is not allowed with the `--connection-type tunnel` option set to anything other than the default.

Creating and using IPsec policies is recommended for securing data transmission across a network, especially a network beyond the confines of a secure data center. IPsec is enabled at the tunnel level, not the circuit level, which means all of a circuit of a tunnel is encrypted and use the same IPsec settings. Different tunnels can have unique IPsec settings. IPsec uses Internet Key Exchange (IKE) to set up the security association. The key exchange can be through a pre-shared key (PSK) or a public-key infrastructure (PKI).

When running IPsec, both sides of the extension tunnel must be running the same Fabric OS version. When you use a PSK, both ends of the secure tunnel must be configured with the same key string. If both ends are not configured with the same key, the IKE session will not start and the extension tunnel will not be able to be established.

The pre-shared key must be a string of 16 to 64 alpha-numeric characters. The following are the requirements for PKI:

- X.509 certificates are supported.
- ECDSA certificates are supported only on the Brocade Extension platforms.
- Non-ECDSA certificates are not supported.
- PKI support is restricted to key-size P384 and hash-type SHA384.

> **NOTE**
> For more information about configuring and managing certificates using the `secCertMgmt` command, refer to "SSL Configuration Overview" in the *Brocade Fabric OS Administration Guide*.

An IPsec policy must be defined before IPsec can be enabled on a tunnel. Multiple IPsec policies can be defined; however, only one policy can be applied to a tunnel. All circuits in the tunnel use the same IPsec policy.

An IPsec policy can be modified while the policy is assigned to a tunnel or WAN Tool session. Sometimes, the local and remote sides become out of sync, preventing the tunnel from coming up and displaying an authentication error.

For more information on how to restart IKE authentication, see IPsec IKE Authentication Failures.

When you use the PSK, the IPsec policy must be configured with the same PSK on each end of the tunnel. The policy name can be different at each end, but the key must be the same.

Use the following steps to create, enable, or disable the IPsec policy:

1. Connect to the switch and log on as an admin.
2. Enter the `ipsec-policy --help` command. For example, `portcfg ipsec-policy --help`.

```
switch:FID128:admin> portcfg ipsec-policy --help

Usage:   portCfg ipsec-policy <name> { create [<args>] | modify [<args>] | delete | restart | --help }

Name Format:
    <string> - The IPSec Policy name(Min 1 character, Max 31 characters).
               Cannot contain the following special characters:
               ;$!#`/\><&'"=,?.*^{}()

Option:    create  - Create the specified IPSec Policy
           modify  - Modify the specified IPSec Policy
           delete  - Delete the specified IPSec Policy
           restart - Restart all inactive IKE sessions for this policy
```

```
            help    - Show this usage message

    Optional Arguments:
      -p,--profile { preshared | pki } -
                                  -  Set the IPSec-Profile.
      -k,--preshared-key <16-64>  -
                                  -  String value for preshared key (for
                                     authentication method "SHARED_KEY").
      -K,--keypair <keypair name> -  Name of the keypair. Max 31 Chars (for
                                     authentication method "ECDSA_P384").
      -h,--help                   -  Show the IPSec-Policy configuration usage
                                     statement.

    Example:
      portcfg ipsec-policy myPolicy create --preshared-key 1234567890abcdef
```

3. Use the `portcfg ipsec-policy create` command to define a policy. The pre-shared key must be 16 through 64 characters long. The following example creates an IPsec policy with the name myPolicy1.

```
switch:admin> portcfg ipsec-policy myPolicy1 create -k "123ashorttestkey"
Operation Succeeded.
```

4. After creating the IPsec policy, apply it to the tunnel. Use the `portcfg fciptunnel modify` command to enable a policy on a tunnel.

   NOTE

   Applying an IPsec policy to a tunnel is disruptive when modifying an online tunnel.

   The following example uses the `portcfg fciptunnel <ve> modify` command to enable the policy myPolicy1 for an existing tunnel. IPsec must be enabled on both ends of the tunnel.

```
switch:admin> portcfg fciptunnel 24 modify --ipsec myPolicy1

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fciptunnel specified for a brief period of time. This
 operation will bring the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n]y
Operation Succeeded
```

5. Use the `portshow ipsec-policy` command to display the available IPsec policies. The following example displays the IPsec policy name and policy key. You must use the `--password` option to display the key; otherwise, it is represented as a string of asterisks.

```
switch:admin> portshow ipsec-policy --password

 IPSec Policy    Flg Authentication data
-------------------------------------------------------------------------------
 MyIPsec         S-- CwYQBFJUAo87zGRApVvIWxiINtmAZJtn
 MyIPsec2        S-- abcdefghijklmnopqrstuvwxyz1234567890
-------------------------------------------------------------------------------
Flags: *=Name Truncated. Use "portshow ipsec-policy -d for details".
       P=PKI Profile  S=Shared-Key Profile
       X=Expired Cert  M=Hash Mismatch
```

   The following example displays IKE information on a tunnel with IPsec enabled. Notice that the `--password` option is not used.

```
switch:admin> portshow ipsec-policy --ike

IPSec Policy Flg Authentication data
```

```
IKE-ID Oper  Flg Local-Addr   Remote-Addr  IKE Rekey      ESP Rekey
-------------------------------------------------------------------
myPolicy1    S-- ****************
dp0.0  UP    I   192.168.1.2  192.168.5.2  5h59m51s  303  3h20m10s  1080
dp0.1  UP    R   192.168.1.2  192.168.5.12 -         -    -         -
dp1.0  UP    R   192.168.1.12 192.168.5.2  -         -    -         -
-------------------------------------------------------------------
Flags: *=Name Truncated. Use "portshow ipsec-policy -d for details".
       P=PKI Profile  S=Shared-Key Profile
       X=Expired Cert  M=Hash Mismatch
       I=Initiator R=Responder
```

The following example displays additional detail information on a tunnel with IPsec enabled:

```
switch:admin> portshow ipsec-policy --detail

 IPSec-policy: MyIPsec
-----------------------------------------------
  Preshared-Key:          ********************************
  Profile:                preshared
  Authentication:         SHARED_KEY
  Encryption:             AES_256_GCM
  Integrity:              NONE
  Diffie Hellman:         MODP_2048
  Pseudo Random Function: HMAC_512
  Num IKE Sessions:       1
```

6. To disable an IPsec policy on a tunnel, use the `portcfg fciptunnel <ve> modify` command. The following example disables the IPsec policy on tunnel 24:

```
switch:admin> portcfg fciptunnel 24 modify --ipsec none

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fciptunnel specified for a brief period of time. This
 operation will bring the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n]    y
Operation Succeeded
```

7. To delete an IPsec policy, use the `portcfg ipsec-policy delete` command. The following example deletes the IPsec policy, myPolicy1. You cannot delete a policy that is in use.

```
switch:admin> portcfg ipsec-policy myPolicy1 delete
Operation Succeeded
```

8. To create an IPsec policy using public-key infrastructure (PKI), use the `portcfg ipsec-policy policy1 create --profile pki` command.

   NOTE
   A certificate must have previously been generated either by self-signing or through a CSR from a CA, refer to the `seccertmgmt` command for more information.

   The following example creates a PKI policy:

```
switch:admin> portcfg ipsec-policy policy1 create --profile pki --key-pair MyKeyPair
Operation Succeeded
```

9. Use the `portshow ipsec-policy --detail` command to display the details. The following example shows IPsec with active IKE sessions. The summary info for the IKE data includes the remote certificate that is requested and an indicator if the hash matches or not.

```
switch:admin> portshow ipsec-policy -i
```

```
IPSec Policy  Flg Authentication data
IKE-ID  Oper  Flg Local-Addr   Remote-Addr  IKE Rekey  ESP Rekey
-------------------------------------------------------------------
ec_pol2           Loc Cert: ven60.pem Hash: Matched
dp0.0   UP    R   79.196.8.10  78.195.8.10  -   -      -   -
                  Rem Cert: sb65.pem Hash:  Matched
dp0.1   UP    R   79.196.8.10  78.195.8.11  -   -      -   -
                  Rem Cert: sb125.pem Hash: Matched
-------------------------------------------------------------------
Flags: *=Name Truncated. Use "portshow ipsec-policy -d for details".
I=Initiator R=Responder
```

10. To modify a PSK policy, only the pre-shared key must be modified.

```
switch:admin> portcfg ipsec-policy MyPSKPolicy modify --preshared-key asdf1234asdf1234

!!!! WARNING !!!!
Modify operation can disrupt the traffic on any tunnel using this IPSec policy. This operation may bring
 the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n]    y
Operation Succeeded
```

11. To modify a PKI policy, modify the profile and the authentication data. Both actions must occur at the same time.

```
switch:admin> portcfg ipsec-policy MyPKIPolicy modify --profile pki --keypair MyKeyPair

!!!! WARNING !!!!
Modify operation can disrupt the traffic on any tunnel using this IPSec policy. This operation may bring
 the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n]    y
Operation Succeeded
switch:admin> portshow ipsec-policy --detail

IPSec-policy: policy1
-------------------------------------------------------
Profile:          PKI
Encryption:       AES-256-CBC
Pseudo-Random:    PRF-HMAC-384
Integrity:        HMAC-SHA-384-192
Diffie-Hellman:   ECDH-P384
Authentication:   ECDSA-P384
Key-Pair:         MyKeyPair
Certificate:      MyKeyPair_cert.pem
Certificate Hash: aff6fea1b19d81ea43aa72f4275a9cf550edadc0
Num IKE Session:  0
```

IPsec IKE Authentication Failures

A Brocade Extension platform might require user intervention to resolve an IKE authentication error. This error occurs if there is an IKE session parameter mismatch. The IKE session is placed in a fault state when such an error occurs, and it remains there until it is manually corrected, thus preventing an intruder from using multiple authentication attempts. For example, if the preshared-key is incorrect during the initial IKE authentication exchange, it triggers an IKE authentication error generates a RASLog XTUN-2012 message. For more information, refer to the *Brocade Fabric OS Message*

Reference Manual. The IKE Operational Status is listed as FAULT on the output of the `portshow ipsec- policy --ike` command.

To recover from an IKE authentication error, the IKE authentication exchange must be restarted. Restart IKE authentication with the `portcfg ipsec-policy <policy_name> restart` command.

1. Connect to the switch and log on as admin.
2. Restart the IKE sessions that are in a FAULT state.

The following command displays an IPsec policy with IKE sessions in a FAULT state, which requires restarting the IKE authentication for the IPsec policy:

```
switch:admin> portshow ipsec-policy --ike

IPSec Policy    Authentication data
IKE-ID    Oper    Flg Local-Addr    Remote-Addr    IKE Rekey    ESP Rekey
-------------------------------------------------------------------------------
pol1        asdf12345678asdf
dp0.0    FAULT    I 192.168.4.20    192.168.4.10    49710d6h    0    0s    0
dp0.2    FAULT    I 192.168.4.20    192.168.4.11    49710d6h    0    0s    0
dp0.1    FAULT    I 192.168.5.20    192.168.5.10    49710d6h    0    0s    0
dp0.3    FAULT    I 192.168.5.20    192.168.5.11    49710d6h    0    0s    0
dp1.0    FAULT    I 192.168.4.21    192.168.4.10    49710d6h    0    0s    0
dp1.1    FAULT    I 192.168.5.21    192.168.5.10    49710d6h    0    0s    0
-------------------------------------------------------------------------------
Flags: *=Name Truncated. Use "portshow ipsec-policy -d for details".
       I=Initiator R=Responder

switch:admin> portcfg ipsec-policy MyPolicy restart
Operation Succeeded

switch:admin> portshow ipsec-policy --ike

IPSec Policy    Authentication data
IKE-ID    Oper Flg Local-Addr    Remote-Addr    IKE Rekey    ESP Rekey
-------------------------------------------------------------------------------------
MyPolicy            asdf12345678asdf
dp0.0    UP    R    192.168.4.20    192.168.4.10    -    -    -    -
dp0.1    UP    R    192.168.4.20    192.168.4.11    -    -    -    -
dp0.2    UP    R    192.168.5.20    192.168.5.10    -    -    -    -
dp0.3    UP    R    192.168.5.20    192.168.5.11    -    -    -    -
dp1.0    UP    I    192.168.4.21    192.168.4.10    5h59m51s    0    3h13m0s    0
dp1.1    UP    I    192.168.5.21    192.168.5.10    5h59m51s    0    3h8m17s    0
-------------------------------------------------------------------------------------
Flags: *=Name Truncated. Use "portshow ipsec-policy -d for details".
       I=Initiator R=Responder
```

Extension Tunnels

Ensure that the following features have been configured before configuring WAN side tunnels:

- App mode: Brocade SX6 Extension Blade only: FCIP is default, or Hybrid mode if deploying IP Extension
- VE mode: Brocade 7850 Extension Switch on 6VE/18VE and Brocade SX6 Extension Blade on 10VE/20VE
- GE mode: Brocade 7810 Extension Switch only: RJ-45 interfaces or first two optical bays
- GE interface speeds (GE, 10GE, or 25GE)
- GE interfaces needed are set to the WAN or LAN side, WAN side is the default
- WAN side IP interfaces (IPIF ge#.dp#)
  - VLAN
  - MTU
- WAN side IP routes, if a layer 3 WAN network
- IPsec policy
- Service Level Agreement (SLA)

You must verify the IP connectivity before configuring tunnels; otherwise, the tunnel does not come up. Use the `portcmd --ping` command. If there is a local WAN side gateway, first ping the gateway. If the ping is successful with the local gateway, ping the remote gateway. If ping is successful with the remote gateway, ping the remote IPIF IP address.

When you configure a tunnel, you must configure the local and remote sides before it comes up. A tunnel consists of one or more circuits. A tunnel has two endpoints identified by Virtual E_Ports (VE_Ports) used. A particular VE_Port belongs to its corresponding DP. For example, on a Brocade 7850 Extension Switch, tunnels 24 through 26 are established using VE_Ports 24 through 26, which live on DP0. Circuit endpoints are IPIFs and identified by its IP address. When an IPIF is configured, it is associated to a particular DP and Ethernet interface. Therefore, using an IP address from an already configured local IPIF while adding a circuit to a tunnel, specifies the Ethernet interface. The DP of the VE_Port used to create the tunnel and the DP specified when creating the IPIF must match; otherwise, an *Object does not exist* error is generated. The Brocade 7810 Extension Switch only has DP0. The Brocade 7850 Extension Switch and Brocade SX6 Extension Blade have DP0 and DP1.

The circuit IP addresses are the same on each end, except the local and remote entries are flipped.

It is not necessary to use the same VE_Port or DP number on the local and remote sides. For example, deployment of the Brocade 7850 Extension Switch, one side can use VE24 on DP0 and the other side can use VE33 on DP1. Use consistent numbers on both ends for the ease of deployment, management, and troubleshooting.

An extension tunnel is a point-to-point object between two extension platforms. A tunnel is a conduit containing one or more circuits. When a tunnel contains multiple circuits, it becomes a BET because multiple circuits are trunked together. An extension tunnel is a single ISL. Certain features operate within the tunnel scope, such as IPsec, compression, FastWrite, OSTP, FICON Acceleration, protocol distribution percentages, and high/medium/low QoS bandwidth percentages.

To understand tunnels, you must understand the relationship between a tunnel and its VE_Port. Because an extension tunnel is an ISL, each end of a tunnel terminates at a VE_Port. For example, the Brocade SX6 Extension Blade supports VE_Ports that are numbered 16 through 35. Each tunnel endpoint is identified by the VE_Port number of the platform. A VE_Port number of the tunnel does not need to be the same at each end. Each extension platform (Brocade 7850 Extension Switch, Brocade 7810 Extension Switch, and Brocade SX6 Extension Blade) has a different VE_Port numbering scheme.

NOTE
In this document, *source* or *local* is the switch that you are configuring, whereas *destination* or *remote* is the switch on the other end. The local switch and the remote switch depend on which side of the tunnel that you configure.

Staging Configuration Method

The Extension configuration can be made simpler by performing the configuration in steps. Fabric OS facilitates staging by allowing common pieces of the configuration to be entered at a time. For example, first create the tunnel with the VE_Port, IP Extension, IPsec, and Compression settings as applicable using the `portcfg fciptunnel <ve> create` command. Then create and add the circuits, with no settings, using the `portcfg fcipcircuit <ve> create`

<cir_num> command. For the circuits, change create to modify and add the --min and --max bandwidth values. Then add the --local-ip and --remote-ip addresses. Then add the --local-ha-ip and --remote-ha-ip address for eHCL. Each step has a clear objective toward the final configuration, making it easier to follow and less prone to errors. Staging allows for a smaller, more manageable, and readable input from the CLI. As a result, long and unwieldy entries are avoided, which are difficult to read. The cut and paste and *up-arrow* tools are useful when working with staging.

The following is an example configuration of a tunnel with two circuits. Use the portcfg fciptunnel command to create a tunnel with no circuits. Use the portcfg fcipcircuit command to create circuits and then modify those circuits. A tunnel must exist before circuits can be added. On most CLI, the *up-arrow* easily brings up a recent command, which can be altered to save time.

NOTE
Configuration values provided in the following example are not applicable to all environments and are intended only for general configuration. Based on the specific needs of each environment, different values might apply.

```
switch:admin> portcfg fciptunnel 24 create

switch:admin> portcfg fciptunnel 24 modify --ipext enable
switch:admin> portcfg fciptunnel 24 modify --ipsec MyIPsecPolicy
switch:admin> portcfg fciptunnel 24 modify --ip-compression deflate

switch:admin> portcfg fcipcircuit 24 create 0
switch:admin> portcfg fcipcircuit 24 create 1

switch:admin> portcfg fcipcircuit 24 modify 0 --min 5000000 --max 10000000
switch:admin> portcfg fcipcircuit 24 modify 1 --min 5000000 --max 10000000

switch:admin> portcfg fcipcircuit 24 modify 0 --local-ip 172.16.1.3 --remote-ip 172.16.2.3
switch:admin> portcfg fcipcircuit 24 modify 1 --local-ip 172.16.1.4 --remote-ip 172.16.2.4

switch:admin> portcfg fcipcircuit 24 modify 0 --local-ha-ip 172.16.1.5 --remote-ha-ip 172.16.2.5
switch:admin> portcfg fcipcircuit 24 modify 1 --local-ha-ip 172.16.1.6 --remote-ha-ip 172.16.2.6

switch:admin> portcfg fcipcircuit 24 modify 0 --keepalive 1000
switch:admin> portcfg fcipcircuit 24 modify 1 --keepalive 1000
```

Maximum number of tunnels:

- Brocade 7810 Extension Switch:
  - Base: 2 (12VE-13VE)
  - Upgraded/Full: 4 (12VE-15VE)
- Brocade 7850 Extension Switch:
  - 6VE mode: 6 (DP0: 24VE-26VE and DP1:33VE-35VE)
  - 18VE mode: 18 (DP0: 24VE-32VE and DP1:33VE-41VE)
- Brocade SX6 Extension Blade:
  - 10VE mode: 10 (DP0: 24VE-28VE and DP1: 34VE-38VE)
  - 20VE mode: 20 (DP0: 24VE-33VE and DP1: 34VE-43VE)

Tunnels can be created between the Brocade 7810, Brocade 7850, and Brocade SX6 Extension platforms; however, there are Fabric OS version and bandwidth limitations for each platform.

Tunnels are created between the VE_Ports of a platform . IPIFs on each platform identify the local and remote endpoints of a circuit. You configure each end of a circuit so that it points to the opposite end.

A tunnel can be created initially with no circuits, which is called staging. Staging is recommended to simplify and better organize configuration tasks. Circuits are added in a subsequent configuration step making configuration more stepwise logical and each CLI command shorter and concise with less chance of error.

If a tunnel is created with the `portcfg fciptunnel <ve> create --local-ip <local_ip> --remote-ip <remote_ip>` command, circuit-0 is created in addition to the tunnel. To add additional circuits, use the `portcfg fcipcircuit <ve> create <cir_num>` command.

NOTE
Options specified when creating a tunnel must match at each end, for example, compression algorithm and min/max bandwidth values. VE_Port numbers and IPsec policy names can be different. Refer to the *Brocade Fabric OS Command Reference Manual* for additional information about the options available with the `portcfg fciptunnel` command.

Tunnel Requirements

A Brocade Extension platform has the following tunnel requirements:

- The difference between a minimum and maximum circuit bandwidth of a tunnel cannot exceed a factor of 5:1.
- VE_Ports are not associated with any particular GE interface.
- Tunnel settings must be identical on both ends.
  - If there is a mismatch in IPsec, compression, FastWrite, OSTP, and FICON settings, the tunnel fails to come up.
- Best practice is to not implement multiple VE_Ports between the same domains.
  - With each application having unique requirements, it should be confined to a particular logical fabric, which will have a designated domain of its own. Each logical fabric should have a single tunnel (one VE_Port) with multiple circuits between the two domains.
- VE_Ports cannot connect between the same two domains while E_Ports or EX_Ports are connected in parallel.
  - It is not possible to have simultaneous VE_Ports and E_Ports between the same two domains. For example, Extension FCIP plus native FC over DWDM between the same two switches.

Configuring a Tunnel

Ensure that the following configuration is completed on the WAN side to configure a tunnel:

- The IPIF IP address endpoints are configured for each circuit on the extension platforms.
- The IPIF VLAN ID is configured (optional) if the Ethernet traffic is tagged.
- The IPIF MTU is set properly, 1500 bytes is the default.
- The IP routes are configured, provided the local and remote endpoints are on different subnets.
- The IPsec policy is configured (optional, but highly recommended).
- IP connectivity to local gateways, remote gateways, and remote WAN side IPIFs are verified.

NOTE
Configuration values that are provided in the example are not applicable to all environments and are intended only for a general configuration. Based on the specific needs of each environment, different values might apply.

The example above shows the following:

- Tunnel-24 is created
- IP Extension, IPsec, and IP compression are added
- Circuit-0 and Circuit-1 are created for Tunnel-24
- ARL minimum is set to 5Gb/s
- ARL maximum is set to 10Gb/s
- Local and remote IP addresses are added
- Local and remote HA IP addresses are added
- KATOV is changed to 3 seconds (default = 6 seconds), which decreases the link loss recovery time
- To configure the remote side, most commands are the same except for the local and remote IP addresses are swapped

If the settings are not the same on both sides of the tunnel, Op Status displays UpWarn as shown in the following example. The max rate is set to 5000Mbps instead of 10000Mbps, as outlined in the example.

```
switch:admin> portshow fciptunnel -c

Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt  Met/G
-------------------------------------------------------------------------------
24     -        Degrad  --if---PI   2d54m    0.00    0.00    1      -      -
24     0 ge0    UpWarn  ----a--i4   2d54m    0.00    0.00    1  5000/5000  0/-
24     1        InComp  ----a----     10s    0.00    0.00    0      -      0/-
-------------------------------------------------------------------------------
Flags (tunnel): l=Legacy QOS Mode
                i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                I=IP-Ext
      (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                 ARL a=Auto r=Reset s=StepDown t=TimedStepDown  S=SLA
```

1. Connect to the local and remote switch and log on as admin.
2. Create a tunnel. Use the `portcfg fciptunnel <ve> create` command on the local and remote platforms.
   a. Create the local tunnel. Using staging, the following example creates a tunnel on the local Brocade 7850 Extension Switch.

      The Brocade SX6 Extension Blade has slots, add `<slot#>/<ge#>` (for example, 7/24) when specifying the GE interface on a director chassis.

      Optionally, IP Extension can be enabled. If IP Extension is enabled, compression is specified either generally using `--compression` (not shown below) or specifically for FC and IP (as shown below). If not deploying IP Extension, do not add `--ipext enable` .

      IPsec is enabled in this example; the `ipsec-policy` must have previously been configured.

      ```
      SideA_switch:admin> portcfg fciptunnel 24 create
      SideA_switch:admin> portcfg fciptunnel 24 modify --ipext enable
      SideA_switch:admin> portcfg fciptunnel 24 modify --ipsec MyIPsecPolicy
      SideA_switch:admin> portcfg fciptunnel 24 modify --fc-compression fast-deflate
      SideA_switch:admin> portcfg fciptunnel 24 modify --ip-compression aggr-deflate
      ```

   b. Create the remote tunnel.

      The following example creates a tunnel on the remote Brocade SX6 Extension Blade. The Brocade SX6 Extension Blade has slots, and the other platforms do not. Optionally, IP Extension can be enabled. If IP Extension is enabled, compression is specified either generally using `--compression` (not shown below) or specifically for FC and IP (as shown below). IPsec is enabled in this example; the `ipsec-policy` must have already been configured. Configuration is simplified because usually copy and paste work.

      ```
      SideB_switch:admin> portcfg fciptunnel 24 create
      SideB_switch:admin> portcfg fciptunnel 24 modify --ipext enable
      ```

```
SideB_switch:admin> portcfg fciptunnel 24 modify --ipsec MyIPsecPolicy
SideB_switch:admin> portcfg fciptunnel 24 modify --fc-compression fast-deflate
SideB_switch:admin> portcfg fciptunnel 24 modify --ip-compression aggr-deflate
```

Extension Circuits

Extension circuits are individual extension connections within a tunnel, and data flows between the source and destination VE_Ports. Each tunnel can contain a single circuit or multiple circuits. A circuit is a connection between a pair of IP interfaces (IPIF). Each IPIF must have a unique IP address. Each circuit has its own unique pair of source and destination IP addresses. The local and remote IP addresses of a circuit can be from the same subnet, which is referred as Layer 2 (L2), or from different subnets, which is referred as Layer 3 (L3). Both L2 and L3 architectures are supported with both FCIP and IP Extension.

After you configure a tunnel, create its circuits. Start with circuit-0 and add each circuit incrementing by 1. You must create a circuit for each unique path through the IP WAN infrastructure. Circuits require identical settings at each end.

NOTE
The WAN side subnets have no restrictions, they can be layer 2 or 3, as described above. The WAN side subnets are flexible to the environment of the organization. The subnets do not have to be different (L2), but can be (L3). On the LAN side, the local and remote subnets must be different.

When the remote IPIF is not on the same subnet as the local IPIF, an IP route to the remote subnet through the local gateway must be configured for the interface DP pair.

Local and remote, tunnel, and circuit settings must match. To form a point-to-point circuit, the source and destination IP addresses are the same on each end, which is swapped. For example, if you configure IPsec on a tunnel, each end of the tunnel must be configured to use the same IPsec parameters, such as the PSK (pre-shared key). Other circuit parameters must be consistent, such as MTU, min/max bandwidth, distribution and QoS percentages, and keepalive timeout (KATOV) values.

Circuit Requirements

The following requirements and limitations are applicable for extension circuits:

- Each circuit has a unique pair of source and destination IP addresses.
- An IP interface (IPIF) defines an IP address associated with a circuit endpoint and GE interface.
- The maximum rate of a circuit cannot exceed the GE interface speed.
- On a GE interface, defining multiple IPIFs allows multiple circuits on that interface.
- A VE_Port with multiple circuits is a Brocade Extension Trunk (BET).
- If the source and destination IP addresses of a circuit are not on the same subnet, an IP route must be configured on both ends.
- Metric-0 circuits are active under normal operating conditions within the failover group.
- Metric-1 circuits are active when all metric-0 circuits within the failover group have gone offline.
- A failover group control which metric-1 circuits become active when metric-0 circuits in the group go offline.
- When a spillover is configured, failover groups are not used.
- When a spillover is configured, metric-1 circuits are used when the capacity of the metric-0 circuit is exceeded.
- 4:1 Rule: Within a tunnel, consider the ARL minimum rates, the difference between the slowest min rate and the fastest min rate can be no greater than 4x.
  - For example, minimum rates are Cir0=1Gb/s and Cir1=4Gb/s is supported; however, Cir0=1Gb/s and Cir1=5Gb/s is not.
  - If circuits are configured with rates greater than 4x apart, entire bandwidth of the BET might not be fully utilized.
- Minimum supported bandwidth as follows:

  - Brocade 7810 Extension Switch is 20Mb/s.
  - Brocade 7850 Extension Switch is 50Mb/s.
  - Brocade SX6 Extension Blade is 20Mb/s.
- Circuit settings should be identical on both ends including minimum and maximum bandwidth values and KATOV. There are some settings that must match on both ends, otherwise the circuit will not function.

WAN Limits

The maximum supported RTT of 250 ms @ 1.0% packet loss with a KATOV configuration of 2 seconds or more.

The RTT of 200 ms @ 0.1% packet loss is supported on all KATOV configurations.

These limits are the same on all supported platforms.

Circuit Bandwidth Syntax

The `portcfg fcipcircuit` committed rate configuration has been enhanced in Fabric OS 9.2.0. The operands now accept a new syntax of `<mbps>M` and `<gbps>G` in addition to the existing `<kbps>` values. The new syntax allows for an easier entry when dealing with large bandwidths. Only whole number values can be entered, decimals, and floating point values are not permitted. The `portshow fciptunnel -c` output is shown in <Mbps>.

The following example shows the committed rate configuration with the `<mbps>M` and `<gbps>G` values:

```
switch:admin> portshow fcipcircuit

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt  Met/G
-------------------------------------------------------------------------------
 24    0 ge4     InProg   --S-a---4   41m39s    0.00    0.00   2 10000/10000 0/0
 24    1 ge4     InProg   --S-a---4   41m39s    0.00    0.00   2 10000/10000 1/0
 24    2 ge4     InProg   --S-a---4   41m39s    0.00    0.00   2 10000/10000 1/0
-------------------------------------------------------------------------------
 Flags (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                  ARL a=Auto r=Reset s=StepDown t=TimedStepDown  S=SLA

switch:admin> portcfg fcipcircuit 24 modify 0 --min-comm-rate 2500000 --max-comm-rate 2500000
Operation Succeeded

switch:admin> portcfg fcipcircuit 24 modify 1 --min-comm-rate 500M --max-comm-rate 500M
Operation Succeeded

switch:admin> portcfg fcipcircuit 24 modify 2 --min-comm-rate 500M --max-comm-rate 2G
Operation Succeeded

switch:admin> portshow fcipcircuit

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
-------------------------------------------------------------------------------
 24    0 ge4     InProg   --S-a---4   41m39s    0.00    0.00   2  2500/2500 0/0
 24    1 ge4     InProg   --S-a---4   41m39s    0.00    0.00   2   500/500  1/0
 24    2 ge4     InProg   --S-a---4   41m39s    0.00    0.00   2   500/2000 1/0
-------------------------------------------------------------------------------
 Flags (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                  ARL a=Auto r=Reset s=StepDown t=TimedStepDown  S=SLA
```

Configuring Extension Circuits

1. Create two circuits (Cir0 and Cir1) for VE24 on both the local and remote platforms. Create two circuits as follows:

```
SideA_switch:admin> portcfg fcipcircuit 24 create 0
SideA_switch:admin> portcfg fcipcircuit 24 create 1

SideB_switch:admin> portcfg fcipcircuit 24 create 0
SideB_switch:admin> portcfg fcipcircuit 24 create 1
```

2. Add circuit IP addresses to the local and remote platforms.
   Notice that the `create` option was changed to the `modify` option in the command. The local and remote IP addresses are swapped on each end, they point to each other.

```
SideA_switch:admin> portcfg fcipcircuit 24 modify 0 --local-ip 10.0.0.5 --remote-ip 192.168.0.5
SideA_switch:admin> portcfg fcipcircuit 24 modify 1 --local-ip 10.0.0.6 --remote-ip 192.168.0.6

SideB_switch:admin> portcfg fcipcircuit 24 modify 0 --local-ip 192.168.0.5 --remote-ip 10.0.0.5
SideB_switch:admin> portcfg fcipcircuit 24 modify 1 --local-ip 192.168.0.6 --remote-ip 10.0.0.6
```

3. Add eHCL circuit IP addresses to the local and remote Brocade 7850 platforms.

```
SideA_switch:admin> portcfg fcipcircuit 24 modify 0 --local-ha-ip 10.0.0.7 --remote-ha-ip 192.168.0.7
SideA_switch:admin> portcfg fcipcircuit 24 modify 1 --local-ha-ip 10.0.0.8 --remote-ha-ip 192.168.0.8

SideB_switch:admin> portcfg fcipcircuit 24 modify 0 --local-ha-ip 192.168.0.7 --remote-ha-ip 10.0.0.7
SideB_switch:admin> portcfg fcipcircuit 24 modify 1 --local-ha-ip 192.168.0.8 --remote-ha-ip 10.0.0.8
```

4. Add minimum and maximum bandwidth to each circuit. The circuits do not come online without the `min` and `max` set. In this example, min is set to 10Gb/s and max is set to 20Gb/s, which is possible on the Brocade 7850 Extension Switch. The minimum must be min < max, and in this case it is referred to as Adaptive Rate Limiting (ARL). Minimum can equal maximum (min = max), in this case, it is referred to as Committed Information Rate (CIR). For example, `--min 10G --max 10G`.

```
SideA_switch:admin> portcfg fcipcircuit 24 modify 0 --min 10G --max 20G
SideA_switch:admin> portcfg fcipcircuit 24 modify 1 --min 10G --max 20G

SideB_switch:admin> portcfg fcipcircuit 24 modify 0 --min 10G --max 20G
SideB_switch:admin> portcfg fcipcircuit 24 modify 1 --min 10G --max 20G
```

5. Use the `portshow fciptunnel -c` command to display the tunnel and circuit information.
   The following example displays the tunnel and circuit (-c) information on the local switch. OpStatus shows **Up**, which indicates that both sides of the tunnel are configured correctly. Various arguments exist for the `portshow fciptunnel` command for showing various tunnel and circuit settings: `--perf, --config, --qos, --tcp, --hcl-status, --ipsec, --ha`.

```
SideA_switch:admin> portshow fciptunnel -c

Tunnel Circuit OpStatus Flags    Uptime        TxMBps RxMBps ConnCnt CommRt Met/G
--------------------------------------------------------------------------------
24     -       Up       --i----PI   18h3m  0.00   0.00  0.00    5       -       -
24     0 ge16  Up       ----ah-i4   18h3m  0.00   0.00  0.00    5  10000/20000 0/-
24     1 ge17  Up       ----ah-i4   18h3m  0.00   0.00  0.00    5  10000/20000 0/-
--------------------------------------------------------------------------------
Flags (tunnel): l=Legacy QOS Mode
                i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                I=IP-Ext
     (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                ARL a=Auto r=Reset s=StepDown t=TimedStepDown S=SLA
```

Tunnel-24 contains two circuits, circuit-0 on ge16 and circuit-1 on ge17.

6. Use the `portshow fcipcircuit <ve>` command to display the circuit information for a specific tunnel. The following example shows the circuit information for tunnel-24 on the local platform:

```
SideA_switch:admin> portshow fcipcircuit 24

Circuit 24.0 (DP0)
==================================================
Admin/Oper State     : Enabled / Online
Flags                : 0x00000000
IP Addr (L/R)        : 10.0.0.5 ge16 <-> 192.168.0.5
HA IP Addr (L/R)     : 10.0.0.7 ge16 <-> 192.168.0.7
Configured Comm Rates: 10000000 / 20000000 kbps
Peer Comm Rates      : 10000000 / 20000000 kbps
Actual Comm Rates    : 10000000 / 20000000 kbps
Keepalive (Cfg/Peer) : 6000 (6000 / 6000) ms
Metric               : 0
Connection Type      : Default
ARL-Type             : Auto
PMTU                 : Disabled
SLA                  : (none)
Failover Group       : 0
VLAN-ID              : NONE
L2Cos (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
L2Cos (IP:h/m/l)     : 0 / 0 / 0
DSCP (FC:h/m/l)      : 0 / 0 / 0 (Ctrl:0)
DSCP (IP:h/m/l)      : 0 / 0 / 0
Cfgmask              : 0x40000000 0x01e13def
Flow Status          : 0
ConCount/Duration    : 7 / 72d18h
Uptime               : 55m11s
Stats Duration       : 55m11s
Receiver Stats       : 75576 bytes / 480 pkts / 23.00 Bps Avg
Sender Stats         : 63576 bytes / 477 pkts / 15.00 Bps Avg
TCP Bytes In/Out     : 64196432680 / 65120353832
ReTx/OOO/SloSt/DupAck: 465 / 180 / 14 / 0
RTT (min/avg/max)    : 1 / 1 / 1 ms
Wan Util             : 0.0%
```

Verifying Tunnel and Circuit Configuration

After creating the local and remote configurations, verify that the tunnel and circuit parameters are correct using the `portshow fciptunnel` command. For more information, refer to the *Brocade Fabric OS Command Reference Manual*.

By comparing the command output for each end of the tunnel, configuration issues can be identified. Also, the output often indicates if there is a parameter mismatch.

1. Use the `portshow fciptunnel` command to display information about the tunnel. This example shows basic use of the `portshow fciptunnel -c` command, which shows the tunnel and its circuits. The tunnel is up but has a warning.

```
SideA_switch:admin> portshow fciptunnel -c

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt    Met/G
-------------------------------------------------------------------------------
```

```
 24    -         UpWarn  --i----PI   73d13h    0.00    0.00   1      -          -
 24    0 ge16    Up      ----ah-i4   73d13h    0.00    0.00   1 10000/20000    0/-
 24    1 ge17    Up      ----ah-i4   73d13h    0.00    0.00   1 10000/20000    0/-
-------------------------------------------------------------------------------
```

2. Use the `portshow fciptunnel <ve> -c` command to show detailed information about a specific tunnel (or VE_Port). This example shows the `portshow fciptunnel slot/port -c` command, which expands the basic information of the `portshow fciptunnel -c` command. The tunnel is up and shows a warning with a reason. The ARL bandwidth values must be the same on both ends of a circuit, in this case circuit 0 minimum is set to 110000 kbps and the remote side minimum is set to 100000. Circuit 1 is online without any warnings. Min and Max values are the scope of a circuit. The protocol distribution and QoS priority percentages must be consistent on both ends of the tunnel. QoS distribution and priority percentages are the scope of a tunnel.

   The CLI output is flagged with **<<<<<** shows the warning; it does not appear in the actual display.

```
SideA_switch:admin> portshow fciptunnel 24 -c

 Tunnel: VE-Port:24 (idx:0, DP0)
 ================================================
  Oper State            : Online Warning         <<<<<
  TID                   : 24
  Flags                 : 0x00000000
  IP-Extension          : Enabled
  Compression           : None
  FC-Compression        : None (Inherited)
  IP-Compression        : Aggressive Deflate (Override)
  QoS Distribution      : Protocol (FC:50% / IP:50%)
  FC QoS BW Ratio       : 50% / 30% / 20%
  IP QoS BW Ratio       : 50% / 30% / 20%
  Fastwrite             : Disabled
  Tape Pipelining       : Disabled
  IPSec                 : Enabled
  IPSec-Policy          : MyIPsecPolicy
  Legacy QOS Mode       : Disabled
  Load-Level (Cfg/Peer): Failover (Failover / Failover)
  Local WWN             : 10:00:88:94:71:a5:c4:30
  Peer WWN              : 10:00:c4:f5:7c:bc:c1:50
  RemWWN (config)       : 00:00:00:00:00:00:00:00
  cfgmask               : 0x0080001f 0x4001024c
  Uncomp/Comp Bytes     : 289154186 / 289154186 / 1.00 : 1
  Uncomp/Comp Byte(30s): 1274 / 1274 / 1.00 : 1
  Flow Status           : 0
  ConCount/Duration     : 9 / 210d18h
  Uptime                : 73d13h
  Stats Duration        : 73d13h
  Receiver Stats        : 984209944 bytes / 4796400 pkts /  192.00 Bps Avg
  Sender Stats          : 694942606 bytes / 3206864 pkts /  143.00 Bps Avg
  TCP Bytes In/Out      : 153049690740 / 190897474128
  ReTx/OOO/SloSt/DupAck: 2885998 / 5845943 / 2775590 / 108251
  RTT (min/avg/max)     : 145 / 149 / 263 ms
  Wan Util              : 0.5%
  TxQ Util              : 0.0%

 Circuit 24.0 (DP0)
 ================================================
```

```
 Admin/Oper State     : Enabled / Online Warning  <<<<<
 Flags                : 0x00000000
 IP Addr (L/R)        : 10.0.0.5 ge16 <-> 192.168.0.5
 HA IP Addr (L/R)     : 10.0.0.7 ge16 <-> 192.168.0.7
 Configured Comm Rates: 10000000 / 20000000 kbps  <<<<<
 Peer Comm Rates      : 11000000 / 20000000 kbps  <<<<<
 Actual Comm Rates    : 10000000 / 20000000 kbps
 Keepalive (Cfg/Peer) : 6000 (6000 / 6000) ms
 Metric               : 0
 Connection Type      : Default
 ARL-Type             : Auto
 PMTU                 : Disabled
 HA PMTU              : Disabled
 SLA                  : (none)
 Failover Group       : 0
 VLAN-ID              : NONE
 L2Cos (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
 L2Cos (IP:h/m/l)     : 0 / 0 / 0
 DSCP (FC:h/m/l)      : 0 / 0 / 0 (Ctrl:0)
 DSCP (IP:h/m/l)      : 0 / 0 / 0
 cfgmask              : 0x40000000 0x00003c2f
 Configuration Warnings:
   Min-comm-rate / QoS-Ratio / Dist-Ratio         <<<<<
 Flow Status          : 0
 ConCount/Duration    : 12 / 210d18h
 Uptime               : 73d13h
 Stats Duration       : 73d13h
 Receiver Stats       : 37069992 bytes / 165410 pkts /   73.00 Bps Avg
 Sender Stats         : 28639790 bytes / 119041 pkts /  117.00 Bps Avg
 TCP Bytes In/Out     : 214605889958 / 265968025192
 ReTx/OOO/SloSt/DupAck: 1439707 / 2918941 / 1384226 / 54617
 RTT (min/avg/max)    : 145 / 149 / 226 ms
 Wan Util             : 0.5%

Circuit 24.1 (DP0)
==================================================
 Admin/Oper State     : Enabled / Online
 Flags                : 0x00000000
 IP Addr (L/R)        : 10.0.0.6 ge17 <-> 192.168.0.6
 HA IP Addr (L/R)     : 10.0.0.8 ge17 <-> 192.168.0.8
 Configured Comm Rates: 10000000 / 20000000 kbps
 Peer Comm Rates      : 10000000 / 20000000 kbps
 Actual Comm Rates    : 10000000 / 20000000 kbps
 Keepalive (Cfg/Peer) : 6000 (6000 / 6000) ms
 Metric               : 0
 Connection Type      : Default
 ARL-Type             : Auto
 PMTU                 : Disabled
 HA PMTU              : Disabled
 SLA                  : (none)
 Failover Group       : 0
 VLAN-ID              : NONE
 L2Cos (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
```

```
L2Cos (IP:h/m/l)      : 0 / 0 / 0
DSCP (FC:h/m/l)       : 0 / 0 / 0 (Ctrl:0)
DSCP (IP:h/m/l)       : 0 / 0 / 0
cfgmask               : 0x40000000 0x00003c2f
Flow Status           : 0
ConCount/Duration     : 9 / 210d18h
Uptime                : 73d13h
Stats Duration        : 73d13h
Receiver Stats        : 491904122 bytes / 2395683 pkts /  118.00 Bps Avg
Sender Stats          : 347677445 bytes / 1602336 pkts /   26.00 Bps Avg
TCP Bytes In/Out      : 214558978572 / 265959389436
ReTx/OOO/SloSt/DupAck: 1446291 / 2927002 / 1391364 / 53634
RTT (min/avg/max)     : 145 / 149 / 263 ms
Wan Util              : 0.5%
```

3. Use the `portshow fcipcircuit <ve> <cir_num>` command to display the same information as above specific to a circuit without having the output of the tunnel and other circuits.

Verifying WAN side Connectivity

To verify the IP connectivity, use the `portcmd --ping` command to confirm that the local IPIF IP address can communicate with the local gateway (verifies that L2 is working), and the local and remote IP addresses can communicate (verifies that L3 is working, if it is an L3 network). IP connectivity must exist between circuit endpoints before a circuit can come up. Use the `portcmd --ping` command to verify the connectivity.

The following example shows a successful ping operation, which confirms the network connectivity between the source and destination addresses. The source and destination are on different subnets; therefore, the ping test verifies L3 proper operation in the network.

```
switch:admin> portcmd --ping ge2.dp0 -s 192.168.5.2 -d 192.168.1.2

PING 192.168.1.2 (192.168.5.2) with 64 bytes of data.
64 bytes from 192.168.1.2: icmp_seq=1 ttl=64 time=1 ms
64 bytes from 192.168.1.2: icmp_seq=2 ttl=64 time=1 ms
64 bytes from 192.168.1.2: icmp_seq=3 ttl=64 time=1 ms
64 bytes from 192.168.1.2: icmp_seq=4 ttl=64 time=1 ms
--- 192.168.1.2 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 714 ms rtt min/avg/max = 1/1/1 ms
```

BET

BET is an advanced Brocade Extension WAN side feature that benefits both FCIP and IP Extension. The BET feature enables bandwidth aggregation, it is logically a single ISL, it performs Lossless Link Loss (LLL), guarantees in-order delivery, and lossless failover/failback for the non-disruptive resiliency over the WAN. Also, it provides redundant paths, manages load balancing, and protects against a transmission loss due to WAN or IP network failures.

BET is enabled by creating multiple circuits within a tunnel. The tunnel uses multiple circuits to carry data between a source and destination data center. BET does not apply to the LAN side.

BET Example

The following figure shows an example of two extension tunnels. There are two tunnels, one with four circuits (VE_Port 16 to VE_Port 24, Circuits 0-3), and another with two circuits (V_Port 17 to VE_Port 25, Circuits 0-1).

Each circuit is assigned a unique IPIF address. IPIFs are, in turn, assigned to Ethernet interfaces. In the following figure, each IPIF is assigned to a different Ethernet interface, which is not required but provides port redundancy for the tunnel.

Ethernet interface assignment is flexible depending on the needs of the environment, and assignments can be made as desired. For instance, multiple IP interfaces can be assigned to a single Ethernet interface. A circuit flows from an IP interface to an IP interface through the assigned Ethernet interfaces.

Figure 9: Brocade Extension Trunking

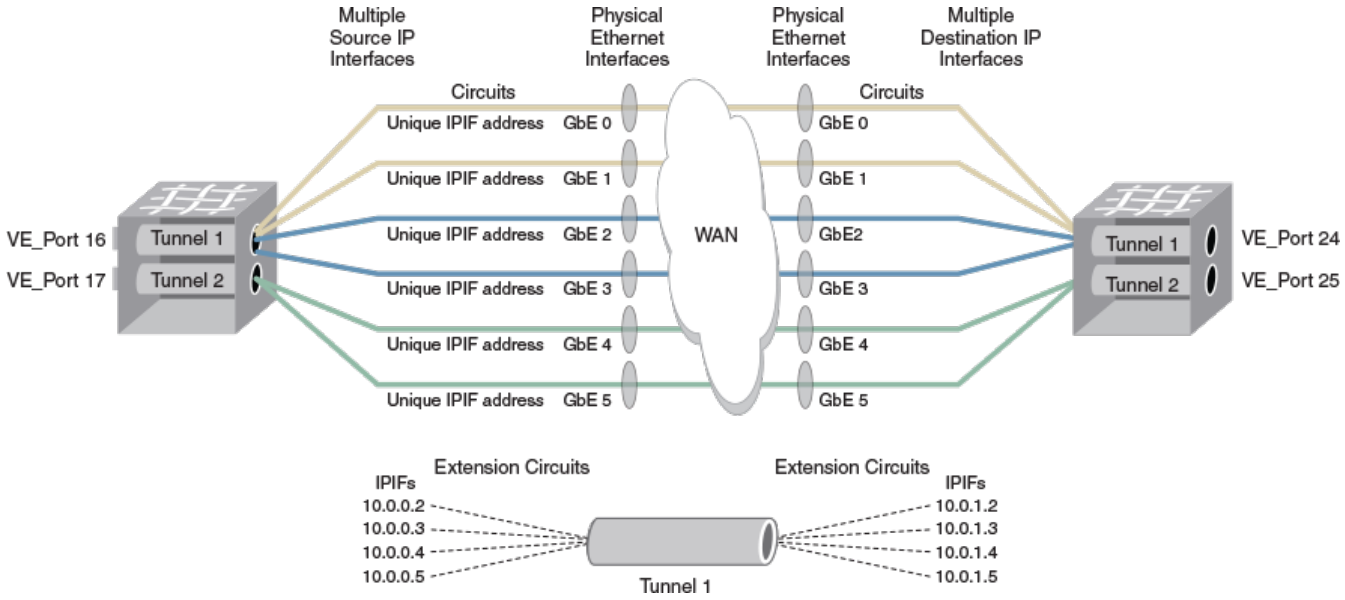

ARL and CIR

ARL utilizes Brocade WAN Optimized TCP to dynamically determine and adjust circuit rate limits. An ARL dynamically increases the rate limit up to the maximum until either it reaches the maximum or detects that no more bandwidth is available. If it detects that no more bandwidth is available, and ARL is not at the maximum, it continues to periodically test for more bandwidth. If ARL determines that more bandwidth is available, it continues to increase toward the maximum. On the other hand, if congestion events are encountered, ARL reduces the rate based on the selected back-off algorithm. ARL never attempts to exceed the maximum configured value and reserves at least the minimum configured value; everything in between is adaptive. If the min and max are configured to be equal; this is referred to as CIR.

An ARL or CIR configuration is required on every circuit to maintain, adjust, and limit the rate at which data is transmitted into the WAN. ARL is configured on a per-circuit basis, not per tunnel, because each circuit can take a unique path and can have a different amount of bandwidth. The ARL minimum and maximum bandwidths are configured when a circuit is created or modified. Circuit bandwidths might be limited based on the licensing or capacity of the platform. Circuits within a tunnel can have the same or different amounts of bandwidth.

The following table describes the maximum bandwidth that a single circuit can have on each platform:

Table 21: Maximum Circuit Bandwidth

Extension Platform	Max Bandwidth Base	Max Bandwidth Upgraded/Full
Brocade 7810 Extension Switch	1Gb/s	2.5Gb/s (GE interface permitting)
Brocade 7850 Extension Switch	25Gb/s (GE interface permitting)	
Brocade SX6 Extension Blade	10Gb/s (GE interface permitting)	

ARL Considerations

The following considerations should be taken into account when configuring ARL:

- As a best practice, the aggregate of maximum-rate bandwidth settings through a tunnel should not exceed the available bandwidth of the wide area network. When you have two circuits using the same backbone, you might oversubscribe the bandwidth so that if one circuit fails, the other can attempt to consume the long haul bandwidth. In 6VE mode, the aggregate of the maximum-rate bandwidth setting through a tunnel should not exceed 50Gb/s and 48Gb/s in 8VE mode. For ingress rates, there is no limitation because FC flow control rate-limits incoming data.
- The aggregate of the minimum configured bandwidth values on an Ethernet interface cannot exceed the speed of the Ethernet interface.
- Configure minimum rates of all tunnels so that the combined rate does not exceed the specifications listed for the extension product.
- 5:1 Min/Max Committed Rate Rule: The ratio between the minimum committed rate and the maximum committed rate for a circuit cannot exceed five times the minimum. For example, if the minimum is set to 2Gb/s, the maximum for that circuit cannot exceed 10Gb/s. This limit is enforced in software.
- 4:1 Lowest to Highest Circuit Rule: The ratio between any two circuits in the same tunnel, the highest bandwidth should not exceed four times the lowest bandwidth. For example, if one circuit is configured to 2Gb/s, any other circuit in that same tunnel should not exceed 8Gb/s. This rule is not enforced in software, but is recommended for a proper operation.
- For any circuit, the minimum and maximum bandwidth values must match on the local and remote sides.

ARL is enabled only when a circuit's minimum < maximum bandwidth value. If you configure the minimum to equal the maximum bandwidth, for example, 1,000,000 kb/s (1b/s), ARL is not used. Such configuration is referred to as CIR.

ARL FSPF Link Cost Calculations

Fabric Shortest Path First (FSPF) is a link-state path selection protocol that directs the traffic along the shortest path between the source and the destination. FSPF calculations are based on the link cost; every link has a cost. The link cost is calculated by summing the maximum rates of all active circuits in the tunnel.

The following formulas are used:

- If the bandwidth is greater than or equal to 2Gb/s, the link cost is 500.
- If the bandwidth is less than 2Gb/s but greater than or equal to 1Gb/s, the link cost is 1,000,000 divided by the bandwidth in Mb/s.
- If the bandwidth is less than 1Gb/s, the link cost is 2000 minus the bandwidth in Mb/s.

If multiple parallel tunnels (VE_Ports) are used, set identical static link costs for the VE_Ports and configure lossless DLS. This action helps to minimize the FC frame loss during bandwidth updates caused by circuit bounces in VEs.

NOTE
Multiple parallel tunnels are supported, but not recommended.

CIR

ARL can be disabled and a single CIR used. By setting min=max, ARL is disabled and the value that is used for min and max is the CIR. Most commonly, CIR is used when line-rate is needed. For example, there is a 100Gb/s WAN with no contention and the circuit is 20Gb/s; there is no need for the circuit to operate at less than 20Gb/s. Such a circuit should not use ARL, it should use CIR.

NOTE
ARL min and max values are not optional. Circuits do not come up until the min and max values have been configured.

For additional information on ARL, refer to the Brocade Adaptive Rate Limiting whitepaper.

Configuring ARL and CIR

Perform the following steps to configure ARL:

1. Connect to the switch and log on as admin.
2. Use the `portcfg fcipcircuit <ve> modify <cir_num> --min <min_bw_kbps> --max <max_bw_kbps>` command to set the min and max bandwidth values to a circuit.
   The following example sets ARL min and max bandwidth values for circuit 0 and 1. The minimum bandwidth value is 3Gb/s and the maximum bandwidth value is 5Gb/s.

```
Switch:admin> portcfg fcipcircuit 7/24 modify 0 --min 3000000 --max 5000000
Operation Succeeded
Switch:admin> portcfg fcipcircuit 7/24 modify 1 --min 3000000 --max 5000000
Operation Succeeded
```

3. Use the `portshow fcipcircuit <ve>` command to display the detailed circuit information.

The following example shows the settings for the circuits on slot 7 tunnel 24. The configured bandwidth values are min = 3Gb/s and max = 5Gb/s. The following is an example of the ARL backoff algorithm:

(Output is truncated)

```
switch:admin> portshow fcipcircuit 7/24

Circuit 7/24.0 (DP0)
==================================================
Admin/Oper State     : Enabled / Online
Warning Flags        : 0x00000000
IP Addr (L/R)        : 192.168.5.2 ge2 <-> 192.168.1.2
HA IP Addr (L/R)     : 192.168.5.12 ge2 <-> 192.168.1.12
Configured Comm Rates: 3000000 / 5000000 kbps
Peer Comm Rates      : 3000000 / 5000000 kbps
Actual Comm Rates    : 3000000 / 5000000 kbps
Keepalive (Cfg/Peer) : 6000 (6000 / 6000) ms
Metric               : 0
Connection Type      : Default
ARL-Type             : Auto
[...]

Circuit 7/24.1 (DP0)
==================================================
Admin/Oper State     : Enabled / Configuration Incomplete
Flags                : 0x00000000
IP Addr (L/R)        : 0.0.0.0 ge0 <-> 0.0.0.0
HA IP Addr (L/R)     : 0.0.0.0 ge0 <-> 0.0.0.0
Configured Comm Rates: 3000000 / 5000000 kbps
Peer Comm Rates      : 0 / 0 kbps
Actual Comm Rates    : 0 / 0 kbps
Keepalive (Cfg/Peer) : 0 (6000 / 0) ms
Metric               : 0
Connection Type      : Default
ARL-Type             : Step Down
[...]
```

ARL Backoff Algorithm

Although ARL provides shared WAN bandwidth, the amount of replicated storage data for extension connections continues to grow, which needs larger and faster links. On all supported extension platforms, the enhanced response time of ARL provides faster rate limiting adaptation, which permits optimized throughput of not only the extension traffic but also the competing flows.

The back-off mechanism that is implemented by ARL is optimized to increase the overall throughput. ARL dynamically preserves bandwidth and evaluates network conditions to see whether additional back-offs are required. ARL maintains the round-trip-time (RTT) state information to better predict network conditions and to allow intelligent and granular decisions about proper adaptive rate limiting. ARL examines the prior stateful information when it encounters a network error. Rate-limit decisions are made using the ARL algorithm. When configured for automatic, ARL dynamically determines which algorithm to use based on changing network conditions.

On all extension platforms, you can configure the ARL algorithm that is used for backing off traffic. The default is automatic, and the ARL logic determines which approach is best. The ARL algorithm choices are as follows:

- Automatic (default)
- Static reset
- Modified multiplicative decrease (MMD) or step-down
- Time-based decrease or timed step-down

NOTE
Do not change the ARL backoff algorithm unless instructed by Brocade support organization.

Compression

Brocade Extension platforms provide an advanced compression architecture that supports multiple algorithms to optimize compression ratios at various throughput requirements. The available algorithms include the hardware-based field-programmable gate array (FPGA) compression and software-based (HW assisted processor) compression. The available compression modes depend on the platform and the protocol (FCIP or IP Extension).

NOTE

- Throughput for all compression modes depends on the compression ratio achievable for the data. Broadcom makes no promises, guarantees, or assumptions about any compression ratio that any user application might achieve.
- The Brocade 7810 Extension Switch does not support Fast-Deflate compression.
- Each end of a tunnel must have the same compression setting.
- Fast-Deflate compression mode is not supported for IP Extension on any platform. Deflate and Aggressive Deflate function for both FCIP and IP Extension protocols.
- If a selected algorithm has insufficient throughput to accommodate all the tunnels using that algorithm, the compression is rate limiting. In each DP, Fast-Deflate is processed on a different engine than Deflate and Aggr-Deflate. In each DP, Deflate and Aggr-Deflate are processed on the same engine.

Compression Options

The following compression options are available in the Brocade Extension platforms:

- **None** – Compression is disabled.
- **Fast-Deflate** – Hardware-based compression. Fast-Deflate compresses data before entering the DP and decompresses it after leaving the DP. It provides the highest throughput with the least amount of compression ratio.

Because Fast-Deflate is hardware-based, it has the least amount of propagation delay and is the most appropriate for synchronous applications. The Brocade 7810 Extension Switch does not support Fast-Deflate.

- **Deflate** – Hardware-assisted compression. Deflate mode can process FCIP or IP Extension data. Ingress speed (pre-compression) of Deflate operates at a lower throughput than Fast-Deflate, and at a faster throughput than Aggressive-Deflate. Deflate compression provides more compression than Fast-Deflate, and not as much compression as Aggressive-Deflate.
- **Aggressive-Deflate** – Hardware-assisted compression. Aggressive deflate mode can process FCIP or IP Extension data. Ingress speed (pre-compression) of Aggressive-Deflate operates at the lowest throughput of the algorithms. Aggressive deflate compression provides the most compression of the algorithms.

  Hybrid mode allows configuration of a tunnel compression at a protocol level, per FCIP or IP Extension. The per protocol compression configuration overrides the compression setting at the tunnel level. Set the desired compression algorithm per protocol.

Table 22: Extension Hybrid Mode Protocol Compression Options

Compression Algorithm	FCIP	IP Extension
Fast Deflate	Yes (Brocade 7850 Extension Switch and Brocade SX6 Extension Blade only)	No
Deflate	Yes	Yes
Aggressive Deflate	Yes	Yes

Compression is set at the tunnel level and is not set per circuit. Compression can be set for the entire tunnel (regardless of protocol) or per protocol (FCIP and IP Extension). To set the compression per protocol, the platform must be in hybrid mode. The Brocade 7810 and Brocade 7850 Extension Switches only have hybrid mode. The Brocade SX6 Extension Blade must be configured in hybrid mode. Setting compression is a disruptive process.

Follow the following guidelines when assigning compression:

- Hybrid mode allows the compression algorithm selection at the protocol level.
- Protocol-level compression overrides the tunnel-level compression setting.
- The available compression algorithms depend on the platform (Brocade 7810 Extension Switch, Brocade 7850 Extension Switch, and Brocade SX6 Extension Blade).
- The available compression algorithms depend on the protocol (FCIP and IP Extension).
- The overall algorithm throughput depends on the platform.
- The compression ratio depends on the compressibility of the user data.

Follow the guidelines described in the following table to assign a compression algorithm for a tunnel:

Table 23: Assigning Compression on a Brocade 7810 Extension Switch

Total (WAN side) Bandwidth	Compression Algorithm	Protocol
1.5Gb/s and higher	Deflate	FCIP and IP Extension
1.5Gb/s or less	Aggressive deflate	FCIP and IP Extension

Table 24: Assigning Compression on a Brocade 7850 Extension Switch

Total (WAN side) Bandwidth	Compression Algorithm	Protocol
10Gb/s and higher	Fast Deflate	FCIP only
5Gb/s to 10Gb/s	Deflate	FCIP and IP Extension

Total (WAN side) Bandwidth	Compression Algorithm	Protocol
5Gb/s or less	Aggressive deflate	FCIP and IP Extension

Table 25: Assigning Compression on Brocade SX6 Extension Blade

Total (WAN side) Bandwidth	Compression Algorithm	Protocol
4Gb/s and higher	Fast Deflate	FCIP only
2Gb/s to 4Gb/s	Deflate	FCIP and IP Extension
2Gb/s or less	Aggressive deflate	FCIP and IP Extension

Configuring Compression

Perform the following steps to configure the compression:

1. Connect to the local switch and log on as admin.
2. Use the `portcfg fciptunnel <ve> modify --compression` command to set the general compression mode of a tunnel. If FC and IP compression settings are to be configured, there is no need to configure the general compression of the tunnel.
   The following example uses the `portcfg fciptunnel modify` command to set the compression of tunnel for tunnel-24 (VE24).

```
switch:admin> portcfg fciptunnel 24 modify --compression fast-deflate

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fciptunnel specified for a brief period of time. This
 operation will bring the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n] y
```

3. Use the `portcfg fciptunnel <ve> modify --ip-compression` command to set the IP Extension compression algorithm. Brocade SX6 Extension Blade must be in the hybrid mode.
   The following example uses the `portcfg fciptunnel <ve> modify` command to change the tunnel compression for IP Extension. `modify` is used in the staging configuration method.

```
switch:admin> portcfg fciptunnel 24 modify --ip-compression deflate

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fciptunnel specified for a brief period of time. This
 operation will bring the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n] y
```

4. Use the `--compression none` or `--fc-compression none` option to disable compression if it had been previously enabled.

```
switch:admin> portcfg fciptunnel 24 modify --compression none

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fciptunnel specified for a brief period of time. This
 operation will bring the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n] y
```

eHCL

On the Brocade 7850 Extension Switch and the Brocade SX6 Extension Blade, the eHCL allows the non-disruptive Fabric OS firmware updates on circuits that have been configured for HA; FCIP and IP Extension traffic continue to flow during a firmware update.

NOTE
eHCL is not supported on the Brocade 7810 Extension Switch; however, the Brocade 7810 Extension Switch does support the configuration of remote HA addresses for the peer eHCL support when connected to a Brocade 7850 Extension Switch or a Brocade SX6 Extension Blade.

Mainframe environments benefit from eHCL by supporting error-free, nonstop connectivity for FICON applications such as replication and tape, as well as FCIP and IP Extension. eHCL maintains the extension connectivity across the WAN during the firmware update without disrupting active I/O, data loss, and out of order delivery of data.

The Brocade 7850 Extension Switch and the Brocade SX6 Extension Blade have two data processor (DP) complexes, referred to as DP0 and DP1. During the eHCL process, each DP reloads sequentially, while the other DP remains operational. eHCL does not apply to the Brocade 7810 Extension Switch because it has only a single DP. A firmware update occurs on one DP at a time. When initiated, the update always starts on DP0. Before DP0 is updated to the new firmware, all eHCL enabled circuits failover to DP1 to maintain the communication between the local and remote switches. The failover process guarantees lossless and in-order delivery of data.

eHCL uses three tunnel groups to perform non-disruptive updates. Consider the main data center as *`local`* and the DR site as *`remote`*. The perspective in this section stays location consistent. The local backup tunnels (LBT) and remote backup tunnels (RBT) are automatically created when you configure an eHCL circuit by adding the HA IP addresses. A tunnel does not require that all circuits within it be enabled with eHCL; some circuits might be eHCL enabled while others might not.

To create HA tunnel circuits, in addition to the normally configured local and remote IPIF IP addresses for DP0, you must also configure the IPIF HA IP addresses for DP1. The main tunnel (MT) uses IPIF IP addresses on the local DP0 to the remote DP0. The LBT uses IPIF IP addresses on the local DP1 to the remote DP0. The RBT uses IPIF IP addresses on the local DP0 to the remote DP1. You only configure the production IP addresses, the HA IP addresses, and the production tunnel and circuits; do not create the backup circuits (LBT and RBT) as they are created automatically. When you configure production tunnel and circuits, you designate IP addresses specific to eHCL through the HA arguments `--local-ha-ip` and `--remote-ha-ip` in the `portcfg` command. This collection of circuits and tunnels form the *HA tunnel group*.

eHCL requires two additional IP addresses per circuit, one HA IP address at each end. Each circuit has four IP addresses, which are the production local and remote IP addresses and the eHCL local and remote HA IP addresses used during firmware updates. Each HA IP address has its own IPIF. A local pair of IP addresses is typically part of the same subnet, and a remote pair of IP addresses is part of the same subnet, although this is not a requirement. All IP addresses must be able to communicate across the IP infrastructure. If the production IPIF is on DP0, then the eHCL IPIF is on DP1, and conversely. An IP route must be configured for each IPIF because they are on different DPs.

The Brocade 7810 Extension Switch has only one DP and does not support eHCL by itself; however, when connected to a Brocade 7850 Extension Switch or a Brocade SX6 Extension Blade, it does support eHCL. Therefore, the `--local-ha-ip` option for the `portcfg fciptunnel|fcipcircuit` command is not configurable on a Brocade 7810 Extension Switch. The `--remote-ha-ip` option is configurable on a Brocade 7810 Extension Switch to facilitate the Brocade 7850 Extension Switch and the Brocade SX6 Extension Blade eHCL operation. The remote HA IP is required on the Brocade 7810 Extension Switch to allow eHCL to function on a connected Brocade 7850 Extension Switch or Brocade SX6 Extension Blade.

When the Brocade 7850 Extension Switch and the Brocade SX6 Extension Blade are configured for eHCL with a Brocade 7810 Extension Switch at the other end, only the `--local-ha-ip` option should be used. If you use the `--remote-ha-ip` option on the Brocade 7850 Extension Switch or the Brocade SX6 Extension Blade, the tunnel will not reach an ONLINE state and remain in a DEGRADED state. Only the MT group and RBT group appear on the Brocade

7810 Extension Switch, and only the MT group and the LBT group are found on a connected Brocade 7850 Extension Switch or Brocade SX6 Extension Blade.

The `portshow fciptunnel --hcl-status` command functions on a Brocade 7810 Extension Switch, even though it does not support eHCL. This command on the Brocade 7810 Extension Switch provides DP connectivity information for the MT and RBT groups.

eHCL performs coordinated upgrades, meaning the firmware update can be initiated on both ends of a tunnel simultaneously. Fabric OS coordinates the update process to maintain online, lossless, in-order delivery of data.

eHCL Operation

eHCL uses the dual DP complexes, DP0 and DP1, on the extension platforms to provide an uninterrupted traffic flow when the switch firmware is updated. In this example, we assume that the tunnel goes between local DP0 and remote DP0; however, in practice a tunnel can be configured between any DPs. The primary circuit, referred to as the MT, is the circuit that is created between the DP0 on the local and remote switches. The local backup tunnel is a circuit that is created between the local DP1 to the remote DP0. During a firmware update, traffic processing from local DP0 is transferred to local DP1, and DP0 is gracefully terminated. The local DP1 uses the backup tunnel to the remote DP0 to continue the uninterrupted traffic flow. When local DP0 comes back up, traffic processing is returned from local DP1 back to the main tunnel.

When configuring eHCL IPIFs, the LBT IPIF cannot be on the same DP as the MT IPIF. IPIFs are assigned to an Ethernet interface and DP when you configure the IPIF using the `portcfg ipif` command with the `ge_port.dp_num` parameter. Typically, the same GE interface as the normal operation circuits is used for eHCL.

In the following figure, the Brocade SX6 Extension Blade A and Brocade SX6 Extension Blade B both use DP0 as the main tunnel endpoints. Each blade uses DP1 for the LBT. LBT connects to DP0 on the remote switch. From the perspective of the remote switch, the LBT is the RBT. In a simplest eHCL configuration, a circuit is configured between the local and remote switches, and the production circuit consists of the main tunnel, during a firmware update traffic moves to the other DP and uses the LBT. From the perspective of the local switch, the RBT is not used until the remote side performs its firmware update.

Figure 10: eHCL MT, LBT, and RBT

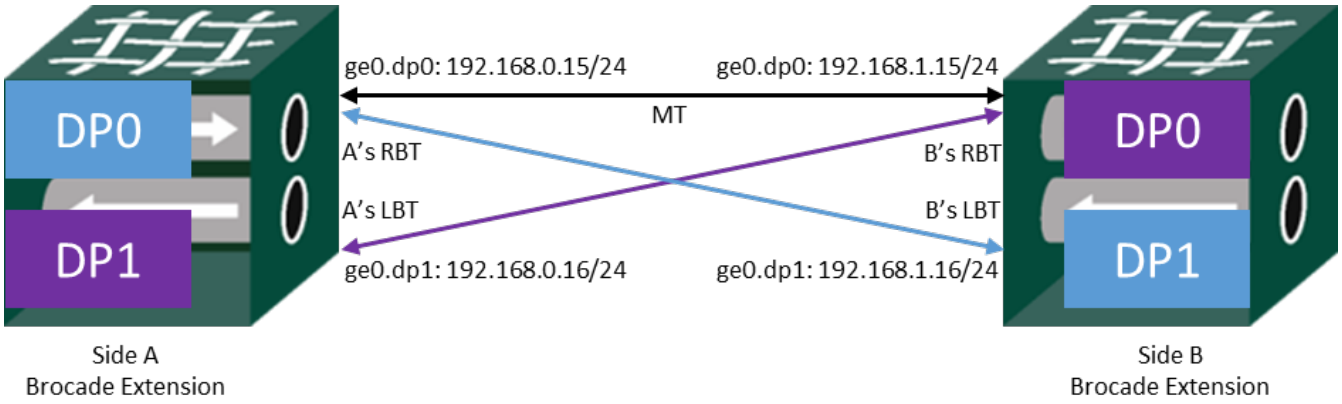

The DP complex creates the following tunnels between Side A and Side B:

- 192.168.0.15 to 192.168.1.15: The MT between DPs
- 192.168.0.16 to 192.168.1.15: Side A's RBT
- 192.168.0.15 to 192.168.1.16: Side A's LBT

  The following modes are the possible options for configuring MTs and LBTs on the Brocade 7850 Extension Switch and Brocade SX6 Extension Blade. The options show that ports are available in the 10VE mode and the 20VE mode.

  - In 10VE mode, the limit is five eHCL tunnels per DP.
  - In 20VE mode, the limit is 10 eHCL tunnels per DP.
- VE_Ports used for eHCL do not need to be on the same local and remote DP. That is for the main tunnel, VE_Ports from one end can be 16-25 and the VE_Ports from other end can be 26-35.

  When eHCL is configured and a main tunnel exists between the local and remote DPs, the following events occur during a firmware update. In the following example, the main tunnel is between DP0 on both ends; however, in practice the MT can be between either DP.
- The MT group provides connectivity during normal production operations.
- All FCIP and IP Extension traffic is maintained without any interruption during the firmware upgrade process.
- An LBT maintains connectivity from the local switch's DP1 to the DP0 of the remote switch while local DP0 is upgrading.
  – LBT is dormant when no firmware update is active.
  – LBT is created automatically when circuit HA IP addresses are added to a circuit.
  – The IPIF and associated HA IP address must exist for DP1 before the HA IP addresses are assigned to the circuit.
- An RBT maintains connectivity from remote switch DP1 to the local switch DP0 while the remote DP0 is upgrading.
  – RBT is dormant when a firmware update is not active.
  – RBT is created automatically when circuit HA IP addresses are added to a circuit.
  – The IPIF and associated HA IP address must exist for DP1 before the HA IP addresses are assigned to the circuit.
- QoS high, medium, and low priorities collapse into a single priority during an eHCL operation. QoS priorities are restored when the eHCL operation completes.
- During eHCL, if a circuit is configured with Differentiated Services Code Point (DSCP) marking, traffic is tagged with medium upon egress.

Creating the `circuit-0` circuit with the `portcfg fciptunnel` or `portcfg fcipcircuit` option creates the MT, which carries traffic through to the remote switch. The LBT is created upon specifying the `--local-ha-ip` option for the circuit, whereas the RBT is created upon specifying the `--remote-ha-ip` option for the circuit. All three tunnel groups (MT, LBT, and RBT) are associated with the same VE_Port on the DP. When a circuit is configured to be eHCL capable, the LBT, and RBT groups are present. These connections are established upon creation of enabled circuits or when the switch boots if eHCL is already configured.

These eHCL tunnels are utilized in the following upgrade process:

1. The firmware writes to the backup partition of the control processor.
2. The control processor reboots from the backup partition with the new firmware.
3. The local DP0 is updated with the new firmware using the following process:
   – Perform feature disable processing on the MT on DP0.
   – Traffic from the MT on DP0 is rerouted to DP1 through the LBT to the remote switch.
   – In-order data delivery is maintained.
   – DP0 reboots with the new firmware.
   – The configuration is reloaded.
   – Traffic from the LBT is failed-back to DP0 through the MT.
4. The local DP1 is updated with new firmware using the following process:
   – Perform feature disable processing on the MT on DP1.
   – Traffic from the MT on DP1 is rerouted to DP0 through the LBT to the remote switch.
   – In-order data delivery is maintained.
   – DP1 reboots with the new firmware.
   – The configuration is reloaded.
   – Traffic from the LBT is failed-back to DP1 through the MT.
5. After the firmware is updated on DP1, all MTs, LBTs, and RBTs are back online, the firmware update is complete.

During a firmware update process, Extension tunnels change their state (up and down). The MT provides connectivity during normal operations; therefore, it is normally up during production. The MT is down during the eHCL process. RBT and LBT are normally up during the normal production operation but do not carry traffic. They only carry the production traffic during the eHCL firmware update process. LBT handles traffic when the local switch DP0 undergoes the eHCL process. RBT handles traffic when the remote switch DP0 undergoes the eHCL process. The RBT is visible as a backup tunnel on the local DP0.

eHCL Limitations and Considerations

eHCL has the following limitations and considerations:

- eHCL has a circuit scope and must be configured per circuit.
- Only circuits that specify a local and remote HA IP address are protected.
- If a circuit is not configured for eHCL, it goes offline during the firmware update process.
- The production circuit (normal operating circuit) is referred to as the MT.
- The local backup circuit to the remote switch is the LBT.
- The tunnel from the remote switch back to the local switch is the RBT, which is used when the remote side performs a firmware update.
- The IPIFs for an MT and LBT must be on opposing DPs. For example, if the MT IPIF is on DP1, then the LBT IPIF is on DP0.
- MT tunnel and circuit parameters are cloned on the LBT and RBT, including the circuit properties such as QoS markings, ARL, FastWrite, OSTP, and FICON Acceleration.

  The management of ARL during eHCL is handled on a VE_Port basis. ARL is temporarily disabled when a VE_Port begins the failover process and no other VE_Ports on a DP are affected. When failover/failback is complete, ARL is enabled for that VE_Port.

  No configuration changes, including changing tunnel or circuit parameters, are allowed during the eHCL process. New device connections requiring zone checking can experience a timeout during the CP reboot phase of the firmware download. A CP performs zone checks and must be active to process new SID/DID connection requests, such as Port Logins (PLOGI).
- eHCL supports all features such as Virtual Fabrics (VF) and FC Routing (FCR).
- eHCL was designed for FICON environments such as IBM XRC, FICON tape, and TS7700 Grid.
- eHCL was designed for the FC replication such as Dell Technologies SRDF, Hitachi Universal Replicator, IBM Global Mirror, HP Remote Copy, and others.

eHCL supports asynchronous and synchronous environments, and IP Extension.

Planning and utilizing the switch for eHCL requires consideration of the following:

- The total switch capacity temporarily decreases by 50% as only one DP complex remains operational during eHCL.
- Available bandwidth during eHCL depends on licensing and the data compressibility.
- Implementation of eHCL requires proper planning.
- Redundant replication fabrics have A and B pathways, which also provide a bandwidth during a firmware update.
- eHCL does not cause a FICON interface control check (IFCC).
- No data inflight is lost.
- In-order delivery is guaranteed.
- eHCL issues the RASLog warnings and error messages (WARN or ERROR).
- If parallel tunnels (VE_Ports) are configured between local and remote sites, and if each tunnel has multiple circuits, you must manually set a static VE link cost. Otherwise, FC traffic might be disrupted during the eHCL activity.
- When Teradata Emulation is enabled on an Extension tunnel, eHCL is not supported. You must perform a disruptive firmware update.

Bandwidth planning might include dedicating one of the two DP complexes to HA or limiting each DP to a maximum of 50% of the licensed capacity to reserve the adequate bandwidth for high-availability operations. Most firmware updates for Fabric OS v9.2.0 and later will support eHCL; however, not every Fabric OS release guarantees a firmware update capable of eHCL. Refer to the Fabric OS Release Notes for potential caveats. The firmware at each end of the tunnel must be compatible, start with the same version of Fabric OS on each end and must upgrade incrementally until arriving at the same updated version on both ends. Normal production operation of extension with disparate Fabric OS versions might introduce instability and aberrant behavior or might prevent the successful tunnel formation or impaired feature functionality.

eHCL does not require any additional communication paths. Production data communications use the traditional extension tunnel that is connected across the WAN. In addition, the two eHCL backup tunnels (LBT and RBT) that exist alongside.

Do not configure certain pairs of VE_Ports in different logical switches with different Fabric OS routing policies. For example, on a Brocade SX6 Extension Blade, one logical switch has Exchange-Based Routing (EBR) and the other has Port-Based Routing (PBR), which is common in mainframe environments. VE_Ports 16 (FID 1) and 26 (FID 2) are used in the different logical switches. During eHCL, these VEs share backend virtual channels (VC) when on the HA path. This configuration can cause unexpected results. The following table shows the VE_Port pairs to avoid. On the Brocade SX6 Extension Blade, the pairing restriction is per blade.

Table 26: Brocade 7850 Extension Switch and Brocade SX6 Extension Blade Conflicting VE_Port Pairs with Differing LS Traffic Policies

Brocade 7850 Extension Switch		Brocade SX6 Extension Blade	
DP0	DP1	DP0	DP1
24	33	16	26
25	34	17	27
26	35	18	28
27	36	19	29
28	37	20	30
29	38	21	31
30	39	22	32
31	40	23	33
32	41	24	34
—	—	25	35

eHCL Firmware Updates

Brocade 7810 Extension Switch does not support eHCL, and a firmware update on a Brocade 7810 Extension Switch is disruptive to FCIP and IP Extension traffic. FC switching is not disrupted during firmware updates on Brocade 7810 Extension Switch.

Brocade 7850 Extension Switch and Brocade SX6 Extension Blade have two DPs that support the eHCL feature. During a firmware update, FC, FICON, FCIP, and IP Extension traffic is not disrupted. To perform non-disruptive firmware updates, eHCL must be configured on one or more circuits; otherwise, FCIP and IP Extension traffic is disrupted. eHCL should be configured on all circuits of a tunnel as a best practice.

NOTE
eHCL is not supported when Teradata Emulation is enabled on an extension tunnel. Even when eHCL is configured correctly, a firmware update disrupts any Teradata connections.

Update both ends of a tunnel with the same Fabric OS release. Both ends of the tunnel can initiate firmware updates simultaneously, and Fabric OS coordinates the update process across the platforms.

For more information about updating the firmware, refer to the "Installing and Maintaining Firmware" section in the *Brocade Fabric OS Administration Guide*.

Perform the following steps to configure eHCL:

1. Connect to the local switch and log on as an admin.
2. Use the `portshow ipif` command to confirm the IPIF addresses for the MT and LBT are configured on DP0 and DP1. The following example shows the IPIFs on the local switch. The following example shows the IPIFs on the local switch:

   The address 192.168.0.15 is on DP0 and used for the MT.

   The address 192.168.0.16 is on DP1 and used for the LBT.

```
local_switch:admin> portshow ipif

Port         IP Address                          / Pfx MTU   VLAN  Flags
-------------------------------------------------------------------------------
ge2.dp0      192.168.0.15                        / 24  1500  0     U R M I
ge2.dp1      192.168.0.16                        / 24  1500  0     U R M I
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

   The following example shows the IPIFs on the remote switch:

   The address 192.168.1.15 is on DP0 and used for the MT.

   The address 192.168.1.16 is on DP1 and used for the LBT.

```
remote_switch:admin> portshow ipif

Port         IP Address                          / Pfx MTU   VLAN  Flags
-------------------------------------------------------------------------------
ge2.dp0      192.168.1.15                        / 24  1500  0     U R M I
ge2.dp1      192.168.1.16                        / 24  1500  0     U R M I
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

3. Use the `portcfg fcipcircuit <ve> modify <cir_num>` command to add the local and remote IP addresses and the local and remote HA IP addresses to an existing circuit.

   The following example configures the tunnel-24 circuit-0 of the local switch with its local and remote IP and local and remote HA IP addresses:

```
local_switch:admin> portcfg fcipcircuit 24 modify 0 \
--local-ip 192.168.0.15 \
--remote-ip 192.168.1.15

local_switch:admin> portcfg fcipcircuit 24 modify 0 \
--local-ha-ip 192.168.0.16 \
--remote-ha-ip 192.168.1.16
```

   The following example configures the tunnel-24 circuit-0 of the remote switch with its local and remote IP and local and remote HA IP addresses:

```
remote_switch:admin> portcfg fciptunnel 24 modify 0 \
--local-ip 192.168.1.15 \
--remote-ip 192.168.0.15

remote_switch:admin> portcfg fciptunnel 24 modify 0 \
```

```
--local-ha-ip 192.168.1.16 \
--remote-ha-ip 192.168.0.16
```

4. Use the `portshow fcipcircuit` command to display the configured values for eHCL. The following example shows the MT endpoints and backup tunnels that are configured on the local switch. The command uses the long option notations.

```
local_switch:admin> portshow fciptunnel --circuits --ha --config

Tunnel Circuit   AdminSt Flags     Local IP                Remote Ip
--------------------------------------------------------------------------------
24     -         Enabled -M-------
24     0 ge2     Enabled ----ah-i4 192.168.0.15            192.168.1.15
24     -         Enabled -R-------
24     0 ge2     Enabled ----ah-i4 192.168.0.15            192.168.1.16
24     -         Enabled -L-------
24     0 ge2     Enabled ----ah-i4 192.168.0.16            192.168.1.15
--------------------------------------------------------------------------------
Flags (tunnel): l=Legacy QOS Mode
                M=MainTunnel L=LocalBackup R=RemoteBackup
                i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                I=IP-Ext
     (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                ARL a=Auto r=Reset s=StepDown t=TimedStepDown S=SLA
```

The following example shows the MT endpoints and backup tunnels that are configured on the remote switch. The command uses short option notations.

```
remote_switch:admin> portshow fciptunnel -chC

Tunnel Circuit  AdminSt Flags     Local IP        Remote Ip
------------------------------------------------------------
24     -        Enabled -M-------
24     0 ge2    Enabled ----ah-i4 192.168.1.15    192.168.0.15
24     -        Enabled -R-------
24     0 ge2    Enabled ----ah-i4 192.168.1.15    192.168.0.16
24     -        Enabled -L-------
24     0 ge2    Enabled ----ah-i4 192.168.1.16    192.168.0.15
------------------------------------------------------------
Flags (tunnel): l=Legacy QOS Mode
                M=MainTunnel L=LocalBackup R=RemoteBackup
                i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                I=IP-Ext
     (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                ARL a=Auto r=Reset s=StepDown t=TimedStepDown S=SLA
```

The M flag identifies the MT between DP0s.

The R flag identifies the RBT, which is from local DP0 to remote DP1 because the remote DP0 is undergoing eHCL.

The L flag identifies the LBT, which is from local DP1 to remote DP0 because the local DP0 is undergoing eHCL.

5. Use the `portshow fciptunnel` command to monitor the eHCL status. The following example shows the status of a correctly configured eHCL when no firmware is undergoing an update:

```
Local_:admin> portshow fciptunnel --hcl-status

Checking FCIP Tunnel HA Status.
Current Status            : Ready
```

```
CP Version               : v9.2.0
DP0 Status:
   State                 : Online - Inactive
   Version               : v9.2.0
   Current FC HA Stage   : IDLE
   Current IP HA Stage   : IDLE
   IP SVI Swapped        : NO
   DP COMM Status        : UP
DP1 Status:
   State                 : Online - Inactive
   Version               : v9.2.0
   Current FC HA Stage   : IDLE
   Current IP HA Stage   : IDLE
   IP SVI Swapped        : NO
   DP COMM Status        : UP

Tunnel 24 (FID:128) FC:HA Online IP:HA Online - Traffic will not be disrupted.
```

FastWrite and OSTP

Brocade FastWrite is an FCIP algorithm that reduces the number of round trips required to complete a SCSI write operation. FastWrite can maintain throughput levels over links that have a significant latency. Although remote data replication (RDR) applications still experience latency, the reduction in throughput due to that latency is minimized for asynchronous applications, and the reduction in response time for synchronous applications is up to 50%.

Protocol optimization features include the following:

- FCIP-FastWrite
- FCIP-Open Systems Tape Pipelining (OSTP)
- FICON Acceleration (The Brocade 7850 Extension Switch and Brocade SX6 Extension Blade only)

Brocade OSTP is an FCIP algorithm that enhances the open systems SCSI tape read and write I/O performance. The WAN network has the longest latency, and OSTP provides the accelerated tape read and write I/Os across the tunnel. OSTP accelerates tape read and write I/Os based on its sequential nature by reducing the number of round-trips that are required to complete the exchange.

As FastWrite and OSTP are in the tunnel scope, they are not configured per circuit. At both ends of a tunnel, Fabric OS should be the same version. Either FastWrite or FastWrite and OSTP are enabled at both ends. To use OSTP, you must enable FastWrite. OSTP and FastWrite are included in a tunnel's scope and enabled during configuration.

A Virtual Fabric (VF) or Logical Fabric (LF) must be implemented if multiple tunnels with protocol optimization are required within a single platform. FastWrite and OSTP require a deterministic FC path between initiator and target. When multiple protocol optimized tunnels exist on the same platform, two logical fabrics can be implemented with one tunnel in each, to provide a deterministic path for each tunnel. Non-controlled, parallel, equal-cost multipath tunnels are not supported between two domains when protocol optimization is enabled on one or more tunnels. Protocol optimization requires the source ID and destination ID (SID/DID) pairs to remain on a specific tunnel during the exchange of I/Os.

NOTE
The Brocade Extension recommends that both ends of an extension tunnel use the same version of Fabric OS when planning Fabric OS upgrade or downgrade.

FastWrite and OSTP Network Requirements

FastWrite and OSTP require a single deterministic path between the initiator and target. If multiple Equal-Cost Multi-Path (ECMP) tunnels exist between the same SID/DID pairs, protocol optimization fails. An outbound command that is routed

over VE_Port and the subsequent response returned over a different VE_Port results in an unknown state and breaks the algorithm. Most often, VF Logical Fabrics are used to create a deterministic path.

The following figure shows network configurations that are supported by FastWrite and OSTP. Neither configuration contains multiple ECMP paths.

> **NOTE**
> Only one emulating tunnel is supported between an initiator port and a target port. A tunnel can have multiple circuits.

Brocade Extension devices can distinguish between storage flows that use protocol optimization and those that do not use protocol optimization. For example, IBM Corporation products do not use FastWrite, but Asynchronous Symmetrix Remote Data Facility (SRDF/A) from Dell Technologies does use FastWrite. It is possible to run both applications over the same tunnel with FastWrite enabled since FastWrite does not engage with IBM flows, but engages with Dell Technologies SRDF/A flows. The same is true for OSTP and IBM products. All IBM flows can utilize the same VE_Port while FastWrite and OSTP are enabled.

Figure 11: FastWrite and OSTP Architecture

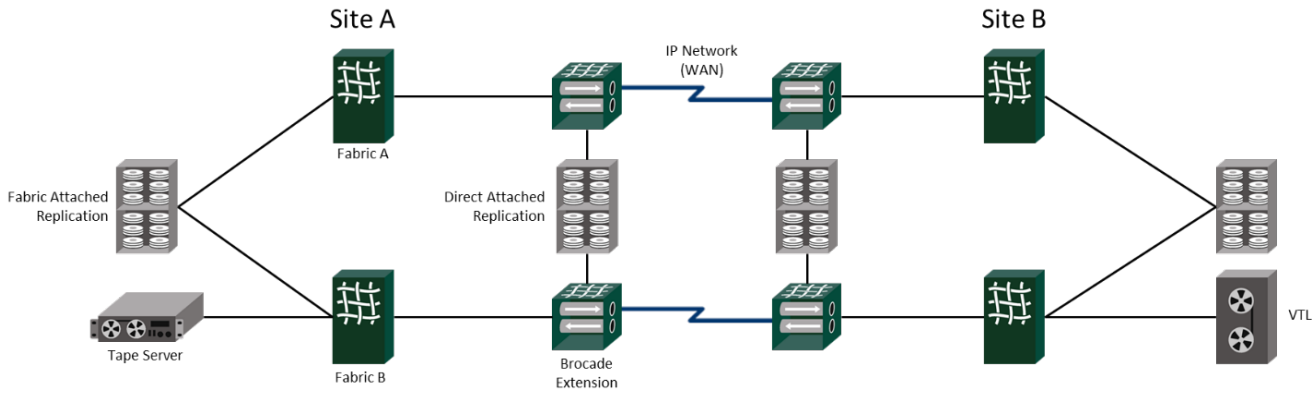

The following steps show how to enable and disable FastWrite and OSTP on a tunnel:

1. Connect to the local switch and log on as admin.
2. Enable FastWrite on tunnel-12 using the staging configuration method (modify an existing tunnel).

```
switch:admin> portcfg fciptunnel 12 modify --fastwrite enable

!!!! WARNING !!!!
The fastwrite feature is incompatible with multiple equal cost paths.
Please ensure that there are no multiple equal cost paths in your fabric before continuing.

Continue with operation (Y,y,N,n): [ n] y

!!!! WARNING !!!!
Delayed modify operation will disrupt traffic on the fcip tunnel specified. This operation will bring the
 existing tunnel down (if tunnel is up) for about 10 seconds before applying the new configuration.

Continue with delayed modification (Y,y,N,n): [ n]      y
Operation Succeeded.
```

3. Disable FastWrite on tunnel-12 using the staging configuration method (modify an existing tunnel).

```
switch:admin> portcfg fciptunnel 12 modify --fastwrite disable

!!!! WARNING !!!!
```

```
Delayed modify operation will disrupt traffic on the fcip tunnel specified. This operation will bring the
 existing tunnel down (if tunnel is up) for about 10 seconds before applying the new configuration.

Continue with delayed modification (Y,y,N,n): [ n]      y
Operation Succeeded.
```

4. Enable OSTP on tunnel-12 using the staging configuration method (modify an existing tunnel).

```
switch:admin> portcfg fciptunnel 12 modify --tape-pipelining enable

!!!! WARNING !!!!
Delayed modify operation will disrupt traffic on the fcip tunnel specified. This operation will bring the
 existing tunnel down (if tunnel is up) for about 10 seconds before applying the new configuration.

Continue with delayed modification (Y,y,N,n): [ n]      y
Operation Succeeded.
```

5. Disable OSTP on tunnel-12 using the staging configuration method (modify an existing tunnel).

```
switch:admin> portcfg fciptunnel 12 modify --tape-pipelining disable

!!!! WARNING !!!!
Delayed modify operation will disrupt traffic on the fcip tunnel specified. This operation will bring the
 existing tunnel down (if tunnel is up) for about 10 seconds before applying the new configuration.

Continue with delayed modification (Y,y,N,n): [ n]      y
Operation Succeeded.
```

Advanced FICON Acceleration

You can accelerate the read and write of FICON tapes and the replication of IBM Global Mirror for System z (XRC) by using FICON Acceleration, while maintaining the integrity of command and acknowledgment sequences. The Brocade SX6 Extension Blade configuration and FICON Emulation support Exchange Based Routing policy and this support comprises the following two major changes:

- VT and Egress VC assignments occur at the start of each FICON exchange.
- FICON Emulation processing creates a new internal object that is called an exchange. The exchange object is created when a new sequence (OXID) is received from either the channel or the control unit, and exist during the active sequences. The exchange control block is similar to the FCP/SCSI TWB structure, but for FICON flows.

As a result of this change, FICON over FCIP performance should be improved, since FICON flows will not all be mapped to the same VT or VC, thereby preventing head-of-line blocking on both the WAN and FC egress paths.

> **NOTE**
> EBR support is limited to a subset of FICON channels and devices. Only enable EBR in the configurations where IBM (or other FICON device providers) indicates the support for Exchange Based Routing.

1. Connect to the local switch and log on as admin.
2. Enable FastWrite on tunnel-12 using the staging configuration method (modify and existing tunnel).

Memory Use Limitations for Large-Device Tunnel Configurations

If you are enabling protocol optimization in tunnels supporting a large number of devices, consider the Brocade Extension platform's memory limitations. In situations where a large number of devices are present or active at the same time, protocol optimizations, such as FICON Acceleration, might be adversely affected.

The data processor (DP) on the Brocade Extension platforms has access to the reserved memory used for control block structure allocations. The following table describes the VE_Ports and DRAM pool sizes specifications for the Brocade Extension platforms:

Table 27: VE_Ports and DRAM Pool Sizes for Brocade Extension Products

Product	DP VE_Ports	DP DRAM Pool Size
Brocade 7850 Extension Switch	DP0: 24-32 DP1: 33-41	2.4 GB 2.4 GB
Brocade 7810 Extension Switch	DP0: 12–15	256 MB
Brocade SX6 Extension Blade	DP0: 16–25 DP1: 26–35	2.4 GB 2.4 GB

Tunnel processing creates more control blocks when any type of emulation feature is enabled, such as FastWrite, OSTP, or FICON Acceleration. When an excessive number of active devices are communicating across the tunnel at the same time, emulation can be adversely affected.

NOTE
Configurations that work with emulation disabled might not work when emulation is enabled.

Displaying the Control Block Memory Pool

Use the `portshow xtun` *`slot/ve`* `-dram2` command to display the current consumption of the tunnel DP complex control block memory pool. The following example shows the command output and displays the total DRAM2 pool size and current consumption for a Brocade SX6 Extension Blade, VE_Port 3/16:

```
switch:admin> portshow xtun 3/16 -dram2

Dram2 Pool Info:

Dram2 Cacheable: 0x8000000a8fce5b80 - 0x8000000b2fce5b7f (2684354560)
pMaster=0x00153029d0 heap_free=2675413120 heap_alloc=8941440 defrags=0 failedAllocs=0

Fast Free Usage=3401 Fast Free Allocates=2372625302
 Merged & Defrag'd Queues:
               User       MAX
 Q SegSize    Alloc     UserA     Alloc'd     Returnd       Max   Current  Bytes
== =======    =====     =====     =======     =======       ===   =======  =====
 1 0x0080      3200      3200        3200           0        70        70  0x00002300
 2 0x0100      3755      3755        3755           0        22        16  0x00001000
 3 0x0180        62        62          62           0        28        21  0x00001f80
 4 0x0200       298       298         298           0       616        28  0x00003800
 5 0x0280       568       568         568           0         3         0  0x00000000
 6 0x0300       258       258         258           0         2         0  0x00000000
 7 0x0380         3         3           3           0         1         0  0x00000000
 8 0x0400       524       524         524           0        31         0  0x00000000
10 0x0500      3584      3584        3584           0         1         0  0x00000000
11 0x0580        10        10          10           0        19         0  0x00000000
12 0x0600       285       285         285           0         1         1  0x00000600
13 0x0680        40        40          40           0         1         0  0x00000000
15 0x0780        59        59          59           0         1         0  0x00000000
16 0x0800         8         8           8           0        30         0  0x00000000
18 0x0900         2         2           2           0         1         0  0x00000000
19 0x0980         1         1           1           0         1         0  0x00000000
22 0x0b00         2         2           2           0         2         0  0x00000000
24 0x0c00        10        10          10           0        30         0  0x00000000
```

```
 25 0x0c80         1         1          1          0         1        0  0x00000000
 28 0x0e00         1         1          1          0         1        0  0x00000000
 29 0x0e80         2         2          2          0         1        0  0x00000000
 31 0x0f80        24        24         24          0         2        0  0x00000000
 32 0x1000         2         2          2          0         1        0  0x00000000
 33 0x1080         3         3          3          0         2        0  0x00000000
 34 0x1100         1         1          1          0         1        0  0x00000000
 40 0x1400        30        30         30          0         1        0  0x00000000
 42 0x1500         3         3          3          0         1        0  0x00000000
 43 0x1580         2         2          2          0         1        0  0x00000000
 44 0x1600        30        30         30          0         1        0  0x00000000
 48 0x1800        35        35         35          0         1        0  0x00000000
 51 0x1980         3         3          3          0         1        0  0x00000000
 53 0x1a80         8         8          8          0         1        0  0x00000000
 56 0x1c00         4         4          4          0         2        0  0x00000000
 59 0x1d80         1         1          1          0         1        0  0x00000000
 64 0x2000        25        32        754        729    327680   326584  0x9f770000

  Fast Free Queues:
   Q  SegSz/dec.     Allocs      Frees      Max  Current      Bytes
=== ======/==== ========== ========== ======== ======== ==========
  1 0x0080/ 128  540738249  540740811     3045     2562 0x00050100
  2 0x0100/ 256   48007719   48007720     3614        1 0x00000100
  3 0x0180/ 384       2109       2135       62       26 0x00002700
  4 0x0200/ 512  540644454  540644707      256      253 0x0001fa00
  5 0x0280/ 640        485        485       16        0 0x00000000
  6 0x0300/ 768       2663       2664        8        1 0x00000300
  7 0x0380/ 896         92         92        3        0 0x00000000
  8 0x0400/1024     340803     341326      524      523 0x00082c00
 10 0x0500/1280     179068     179068     3584        0 0x00000000
 11 0x0580/1408  478277990  478278000       10       10 0x00003700
 13 0x0680/1664        394        399       18        5 0x00002080
 15 0x0780/1920       1853       1853       58        0 0x00000000
 16 0x0800/2048        243        243        8        0 0x00000000
 24 0x0c00/3072  478277990  478278000       10       10 0x00007800
 25 0x0c80/3200   47827799   47827800        1        1 0x00000c80
 31 0x0f80/3968  190495335  190495343       16        8 0x00007c00
 33 0x1080/4224   47827799   47827800        1        1 0x00001080
 48 0x1800/6144        159        159        5        0 0x00000000
 51 0x1980/6528         98         98        3        0 0x00000000
               ---------- ---------- -------- -------- ----------
                2372625302                          3401 0x0010BA80
 Total Bytes in DRAM2 Pool: 2676509440 (free)    1096320 (fastfreed)
Total DRAM Bytes Allocated:    8941440 (in use)
```

Control Blocks Created During FCP Traffic Flows

For FCP/SCSI traffic flows, tunnel processing creates control block structures based on the SID/DID pairs used between devices. Upon enabling FastWrite or OSTP (read or write), additional structures, and control blocks are created for each logical unit number (LUN) based on SID/DID pairs. If multiple SID/DID pairs can access the same LUN, FCP processing in an emulated tunnel creates multiple control blocks for each LUN. The following structures and control blocks are created:

- Initiator-Target Nexus (ITN) structure: Each FCP-identified SID/DID flow is recorded in an ITN data structure.
- Initiator-Target-LUN (ITL) control block: Each specific LUN on a SID/DID flow has an ITL control block created for the flow.
- Turbo-Write Block (TWB) structure: FCP emulation processing creates a TWB structure for each outstanding FC exchange.

Control Blocks Created During FICON Traffic Flows

For FICON traffic flows, tunnel processing creates a control block structure based on the SID/DID pairs that are called a FICON device path block (FDPB). When the FICON emulation is enabled, additional control blocks are created for each SID/DID pair, logical partition (LPAR) number (FICON channel block structure), and LCU number (FICON control unit block structure), and each individual FICON device address on those LCUs. FICON Exchange control blocks are also created if the switch is configured to operate in EBR mode.

The total number of FICON device control blocks (FDCBs) that is created over a FICON emulating tunnel is represented by the following equation:

FDCBs = Host Ports × Device Ports × LPARs × LCUs × FICON Devices per LCU

In extended direct-attached storage device (DASD) configurations, such as those used for IBM z/OS Global Mirror, also known as Extended Remote Copy (XRC), this number increases rapidly.

FDCB Example

In the following example, the tunnel is used to extend two channel path IDs (CHPIDs) from a System Data Mover (SDM) site to a production site. There are also two SDM-extended LPARs, the IBM DS8000 production controllers have 32 LCUs per chassis, and each LCU has 256 device addresses. Based on the previous equation, the number of extended FICON device control block images created is as follows:

2 Host Ports × 2 Device Ports × 2 LPARs × 32 LCUs × 256 Devices per LCU = 56,536 FDCBs

NOTE
The Brocade 7810 Extension Switch does not support FICON emulation.

Considerations for Tunnel Control Block Memory and Device Configuration

The `portshow xtun` *`slot/ve_port`* `-fcp -port -stats` command displays the current tunnel usage and the size of the control blocks once control blocks have been allocated. Use the output from this command to determine the unique characteristics for a specific tunnel configuration. The highlighted text in the following example shows statistics for the control block structures that are created for FCP and FICON traffic flows during tunnel processing.

The following example shows the FICON VE_Port output:

```
switch:admin> portshow xtun 36 -fcp -port -stats

Slot(0.dp1) VePort(36) Port Stats:
 Global Queue Stats:
  Name,cnt,max,usage,size,total size
  Data,0,2,4182,8192,0
  Message,0,1,2091,7448,0
  Stat,0,0,0,1152,0
  Stat Cache,0,0,0,0,0
  Global stats,0,0,0,0,23832
 Port Queue Stats:
  Name,cnt,max,usage,size,total size
  Image,9,9,9,0,0
  SRB,0,0,0,0,0
```

```
  TWB,1,1,8,0,0
 Port Struct Allocation Stats:
  Name,cnt,max,usage,size
  IMAGE,9,9,9,6104
  FDPB,9,0,0,6528
  OXTBL,21,0,0,8192
  FCHB,15,0,0,1544
  FCUB,54,0,0,1816
  FDCB,352,0,0,1128
Global Buffer Stats:
  Name,current,min,max
  Write Data Storage,0,0,0
  Read Data Storage,0,0,0
  XBAR % avail,100,100,100
  WIRE % avail,99,99,100
  SWCOMP % avail,100,100,100
```

The following example shows the FCP/SCSI VE_Port output:

```
switch:maintenance> portshow xtun 24 -fcp -port -stats

Slot(0.dp0) VePort(26) Port Stats:
Global Queue Stats:
Name,cnt,max,usage,size,total size
Data,5,7,407,8192,40960
Message,0,1,175,7448,0
Stat,0,0,0,1152,0
Stat Cache,0,0,0,0,0
Global stats,0,0,0,0,64792
Port Queue Stats:
Name,cnt,max,usage,size,total size
Image,4,4,4,0,0
SRB,0,0,0,0,0
TWB,31,35,783,0,0
Port Struct Allocation Stats:
Name,cnt,max,usage,size
IMAGE,4,4,4,6104 ITN,4,0,0,5624
OXTBL,9,0,0,8192
ITL,3,0,0,5048
TWB,30,35,783,1008
Global Buffer Stats:
Name,current,min,max
Write Data Storage,0,0,0
Read Data Storage,0,0,0
XBAR % avail,100,100,100
WIRE % avail,99,99,100
SWCOMP % avail,100,100,100
Port Buffer Stats:
Name,current,min,max
Write Data Storage,0,0,0
Read Data Storage,0,0,0
PT XFRDY'd Data,0,0,0
Port-based Flow Control Stats:outstanding bytes,accel I/Os,accel bytes,passthr I/Os,passthr
 bytes,pipesize,window closed,max port data
```

```
0,0,0,0,0,192000000,0,0
TCB sharemem alloc: 199776
```

Use output from the `portshow xtun slot/ve_port -fcp -port -stats` command with output from the `portshow xtun slot/ve_port -dram2` command to determine how a tunnel configuration is affecting tunnel control block memory. As a guideline, no more than 80% of the tunnel DP control block memory pool (dram2) should be allocated for SID/DID pair-related control blocks (ITNs, ITLs, FDPBs, FCHBs, FCUBs, and FDCBs). When more than 80% of the pool is allocated, consider redesigning the tunnel configuration to ensure a continuous operation. Consider the existing number of SID/DID pairs in the tunnel configuration when redesigning it and determine whether new switches, chassis, or blades are necessary to reduce the impact on DRAM2.

DP complexes such as the Brocade 7850 Extension Switch and Brocade SX6 Extension Blade generate the XTUN-1008 RASLog message when the following percentages of DRAM memory remain:

- 50%
- 25%
- 12.5%
- 6.25%
- 0.05%

RASLog messages include information regarding the amount of allocated memory, the amount of free memory, and the total size of the pool. Refer to the RASLog messages to determine if you must reduce the size of the extended configuration or to plan for additional switch resources.

The Brocade Extension DPs are limited to the number of FICON device control blocks (FDCB) and extended LUNs (ITL) listed in the following table:

Table 28: FDCB and ITL per DP

Platform	FDCB	ITL
Brocade 7810 Extension Switch	0 (FICON is not supported on the Brocade 7810 Extension Switch)	30,000
Brocade 7850 Extension Switch	512,000	200,000
Brocade SX6 Extension Blade	512,000	200,000

For the Brocade 7850 Extension Switch and Brocade SX6 Extension Blade, the following applies:

During eHCL, emulation, and non-emulation control blocks are duplicated on the remote DP for the RBT. During the eHCL process, twice as much memory is consumed as is normally consumed. This temporary duplication process occurs on the remote DP when the local DP is undergoing an eHCL firmware update.

The amount of DRAM2 memory on the Brocade 7850 Extension Switch or Brocade SX6 Extension Blade should be able to support eHCL operations with approximately 512,000 FDCB active through the VE_Ports on that DP.

Due to the unique configurations of each customer, the number and type of devices supported might vary. For large configurations, you should periodically review the memory usage to ensure reliable and continued operation of the emulated tunnel.

Circuit Failover and Failback

Multiple extension tunnels can be defined between pairs of extension switches or blades, but this defeats the purpose of defining a single tunnel for multiple circuits. Two tunnels between switches or blades are not as redundant or fault-tolerant as multiple circuits within one tunnel.

By providing a lossless link loss (LLL), BET ensures that all data that is lost in flight is retransmitted and reordered before being delivered to the upper-layer protocols. LLL is an essential feature to prevent interface control checks (IFCCs) on mainframes that use FICON and SCSI timeouts for the open-system-based replication.

NOTE
To create multiple parallel tunnels between two switch domains, you must set static link costs on the VEs and must enable lossless dynamic load sharing (DLS). The DLS minimizes FC frame loss due to bandwidth updates that are caused by tunnel circuit bounces or offline tunnels. Bandwidth updates might change a link cost resulting in fabric route changes. Avoid link cost and route changes by setting static link costs. Bandwidth updates also result in DPS rebalancing workloads across parallel tunnels, which might result in an FC frame loss as exchanges move. Enabling lossless DLS helps minimize these events, but frame loss might still occur.

Failover Metrics

Failover circuits and groups are a feature supported on all Brocade Extension platforms. A group defines a set of metric-0 and metric-1 circuits that provide a failover protection for each other. All metric-0 circuits in a group must fail before a metric-1 circuit becomes active. When all metric-0 circuits within the group fail, metric-1 circuits are immediately active, even if metric-0 circuits are active in other failover groups. Typically, one metric-0 and one metric-1 circuit is configured in a group, which provides the one-to-one circuit protection. Failover and failback circuits of a group must be connected to the same VE_Port and have LLL, meaning no data is lost during failovers and failbacks. The Keepalive Timeout Value (KATOV) of the circuit is the amount of time to detect that a circuit has gone offline.

Failover and failback of circuits is a method of preventing interruptions by seamlessly switching circuits that have become offline to a designated backup. To manage failover, each circuit is assigned a metric value, either 0 (default) or 1, which is used to manage the failover from one circuit to another.

BET has a LLL, and no data in-flight is lost during a failover or failback event. If a circuit goes offline, LLL retransmits pending data over another metric 0 circuit. In the following figure, circuit 1 and circuit 2 are both metric 0 circuits. Circuit 1 has gone offline and transmission fails over to circuit 2, which has the same metric. Pending traffic at the time of failure is retransmitted over circuit 2. LLL ensures in-order delivery.

Figure 12: LLL with Metric 0 Circuits

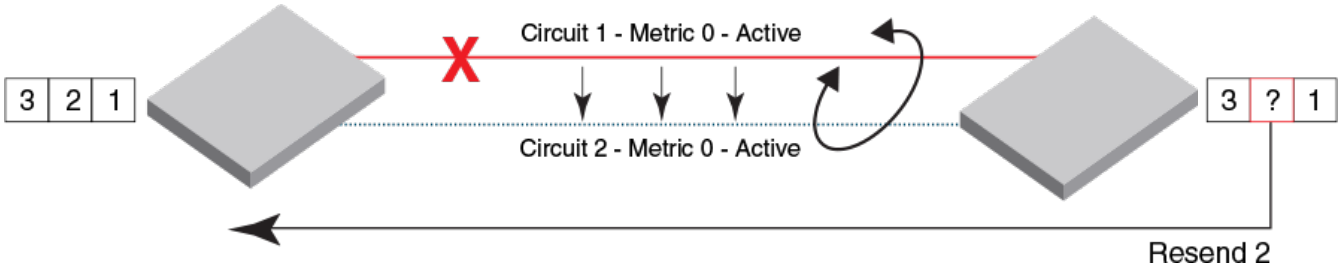

NOTE
Modifying a metric of a circuit briefly disrupts traffic.

Failover Groups

Circuit failover groups control which metric 1 circuits are activated when metric 0 circuits fail. A set of metric 0 and metric 1 circuits are grouped. When all metric 0 circuits within the group have gone down, metric 1 circuits in that group become active. As each failover group is autonomous, circuits with metric 0 in other failover groups are not affected.

The purpose of failover groups is to group circuits so that they can backup each other if one fails. For example, two circuits (metric 0 and metric 1) in one group would form a one-to-one backup. Failover groups can be numbered; group ID is a value from 0 through 9 and the default is 0.

Typically, one metric 0 circuit and one metric 1 circuit are grouped together. The purpose of this is to replace the degraded metric 0 circuit with a corresponding metric 1 circuit as soon as possible. As a result, the problem of a degraded tunnel is avoided and the remaining metric 0 circuits remain operational.

The following two types of configuration are supported:

- Active-Active: Data is weighted balanced across multiple circuits. All circuits are metric 0 (default).
- Active-Passive: If all metric 0 circuits go down within a failover group, data fails over to the metric 1 circuits. If there is a failover/failback, LLL guarantees that all data in delivered in order.

In the following figure, circuit 1 is assigned a metric of 0, and circuit 2 is assigned a metric of 1. Both circuits are in the same tunnel. Circuit 2 does not become active until no all the metric 0 circuits within its failover group have gone offline. If all metric 0 circuits go offline within the failover group, the traffic is retransmitted and sent over available metric 1 circuits within the failover group. Failover between like metric circuits or between different metric circuits is lossless.

Figure 13: Failover to Metric 1 Circuits

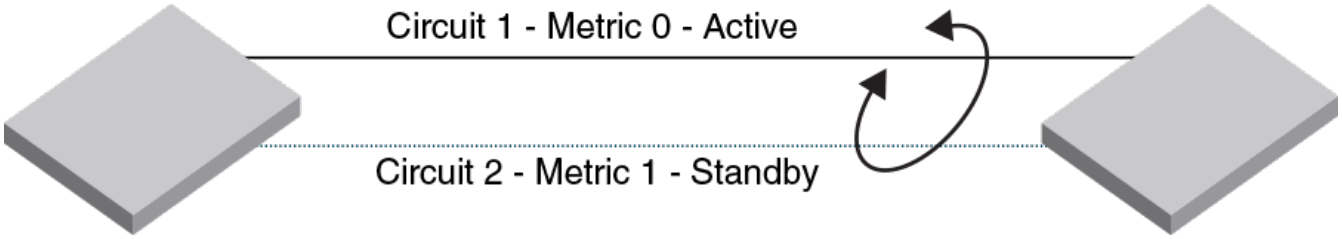

There might be differences in behavior between the metric 1 circuits and the production circuits if the configuration of the metric 1 circuits is not identical to the configuration of the metric 0 circuits. Also, if the WAN path for the metric 1 circuits has different characteristics, the failover behavior might be different from the production behavior.

With circuit failover groups, you can control which metric 1 circuits are activated if metric 0 circuits fails. To create circuit failover groups, define a set of metric 0 and metric 1 circuits within a common failover group. When all metric 0 circuits within the group fail, metric 1 circuits are activated. Each failover group is autonomous.

Typically, there is only one metric 0 circuit in a group, so when the metric 0 circuit fails, a specific metric 1 circuit takes over; this is known as one-to-one failover grouping. A one-to-one configuration prevents a tunnel from operating in a degraded mode because a subset of the overall metric 0 circuits failed.

Circuit Failover Considerations and Limitations

Circuit failover groups operate under the following conditions:

- Each failover group is independent and operates autonomously.
- All metric 0 circuits in a failover group must fail before the metric 1 circuits in that group are used.
- All metric 1 circuits in a failover group are used if there is no metric 0 circuits in the group.
- Circuits can participate in only one failover group.
- Both ends of the tunnel must have the same circuit failover groups defined.
- Tunnel and circuit states indicate a misconfiguration error if circuit failover group configurations are not valid.
- Modifying the metric is a disruptive operation.
- Modifying the failover group ID is a disruptive operation.
- Circuit failover groups do not define load balancing over circuits.

When no circuit failover groups are defined, failover reverts to the default operation: all metric 0 circuit must fail before failing over to metric 1 circuit.

A valid failover group requires at least one metric 0 circuit and at least one metric 1 circuit; otherwise, a warning is displayed. The number of failover groups per tunnel is limited by the number of circuits that can be created on the extension platform as follows:

- The Brocade 7810 Extension Switch can configure up to three valid groups on a 6 circuit tunnel.
- The Brocade 7850 Extension Switch can configure up to five valid groups on a 10 circuit tunnel.
- The Brocade SX6 Extension Blade can configure up to five valid groups on a 10 circuit tunnel.

Consider available WAN bandwidth requirements when configuring failover circuit groups. For more information, see Bandwidth Calculation during Circuit Failover.

Bandwidth Calculation during Circuit Failover

When all metric 0 circuits within the failover group fail, the bandwidth of metric 1 circuits is not calculated as the available bandwidth in a tunnel.

Consider the following circuit configurations for circuits 0 through 3:

- Circuits 0 and 1 are created with a metric of 0.
- Circuit 0 is created with a maximum transmission rate of 1Gb/s.
- Circuit 1 is created with a maximum transmission rate of 500Mb/s.
- The combined bandwidth of circuits 0 and 1 is 1.5Gb/s.
- Circuits 2 and 3 are created with a metric of 1.
- Circuit 2 is created with a maximum transmission rate of 1Gb/s.
- Circuit 3 is created with a maximum transmission rate of 1Gb/s.
- The combined bandwidth of circuits 2 and 3 is 2Gb/s.
- The metric 1 bandwidth is held in reserve. The following actions occur during circuit failures:
- If either circuit 0 or 1 goes offline, traffic continues to flow over the remaining metric 0 circuit. The available bandwidth is still considered to be 1.5Gb/s.
- If both circuits 0 and 1 go offline, failover to circuits 2 and 3; the available bandwidth is updated to 2Gb/s.
- If a metric 0 circuit comes online, the metric 1 circuits return to a standby status, and the available bandwidth is updated as each circuit comes online. For example, if circuit 0 is recovered, the available bandwidth is updated to 1Gb/s. When circuit 1 recovers, the available bandwidth is updated to 1.5Gb/s.

During failover and failback, a metric paradigm is used on a tunnel circuit. Metric 0 is used for the normal production, and metric 1 is used for passive standby failover circuits. When configuring a circuit, metric 0 is the default. All circuits are active if no metrics are configured, which is a common configuration. When all circuits are active, and a circuit goes down, data still fails over to the other remaining circuits using LLL. However, if the amount of licensed bandwidth is limited, the amount of bandwidth consumed is only the aggregate of the circuits that are active at the time. The metric 1 circuits are not counted toward the aggregate if the metric 0 circuits are online.

Examples of Circuit Failover in Groups

During a circuit failover, the following events occur:

- If circuit 1 fails, circuit 3 becomes active and data is load-balanced over circuit 2 and circuit 3.
- If circuit 2 fails, circuit 4 becomes active and data is load-balanced over circuit 1 and circuit 4.
- If both circuit 1 and circuit 2 fail, circuit 3 and circuit 4 become active and data is load-balanced over circuit 3 and circuit 4.

For circuit failover in a tunnel, two failover groups, each with two circuits, are shown in the following table:

Table 29: Tunnel with Two Failover Groups with Two Circuits

Circuits in Tunnel	Failover Group ID	Tunnel Bandwidth	FSPF Link Cost if Circuit Goes Offline	In Use for Tunnel Data?
Circuit 1 Metric 0	1	500Mb/s	1500	If active, yes
Circuit 2 Metric 0	2	1000Mb/s	1000	If active, yes
Circuit 3 Metric 1	1	500Mb/s	1500	Only when circuit 1 fails
Circuit 4 Metric 1	2	1000Mb/s	1000	Only when circuit 2 fails

The following table shows a tunnel with five circuits, two circuits have a failover group defined and three circuits are left in the default failover group. In this example, all data is initially load-balanced across circuits 1, circuit 2, and circuit 3 because they all have metric 0. The following events occur when various circuits go offline:

- If circuit 1 goes offline, circuit 4 becomes active and data is load-balanced over circuit 2, circuit 3, and circuit 4. As the circuit 1 fails over to circuit 4 because both are in the same failover group, and circuit 2 and circuit 3 are online with an equal cost (500Mb/s) bandwidth.
- If circuit 2 goes offline, data is load-balanced over circuit 1 and circuit 3, and no other circuits are active. As the circuit 1 and circuit 3 are the only online circuits because circuit 4 becomes active only when circuit 1 goes offline, and circuit 5 becomes active only when both circuit 2 and circuit 3 go offline.
- If circuit 2 and circuit 3 go offline, circuit 5 becomes active and data is load-balanced over circuit 1 and circuit 5. As the circuit 1 has a defined group and is still active, and undefined group circuit 2 and circuit 3 (metric 0) failover to undefined group circuit 5 (metric 1).
- If circuit 1, circuit 2, and circuit 3 go offline, circuit 4 and circuit 5 become active and data is load-balanced over circuit 4 and circuit 5. As the circuit 1 fails over to circuit 4, which is the failover circuit for group 1. Default group circuit 5 is the failover for default group offline circuit 2 and circuit 3.

The following table describes a tunnel with five circuits. Two circuits have a failover group defined and three circuits are left in the default failover group:

Table 30: Tunnel with Failover Groups and Non-Grouped Circuits

Circuits in Tunnel	Failover Group ID	Tunnel Bandwidth	FSPF Link Cost if Circuit Goes Offline	In Use for Tunnel Data?
Circuit 1 Metric 0	1	500Mb/s	1500	If active, yes
Circuit 2 Metric 0	0	500Mb/s	1500	If active, yes
Circuit 3 Metric 0	0	500Mb/s	1500	If active, yes
Circuit 4 Metric 1	1	500Mb/s	1500	Only when circuit 1 fails
Circuit 5 Metric 1	0	1000Mb/s	1000	Only when circuits 2 and 3 fail

Active-Active Circuits

An active-active configuration has two circuits, and both are configured with the same metric, metric 0. One circuit uses ge0 and the other circuit uses ge1. Both circuits are configured for 1Gb/s and are sending data. The load is balanced across two circuits. The effective tunnel bandwidth is 2Gb/s.

Perform the following steps to configure an active-active circuit:

1. Configure IPIFs.

```
portcfg ipif ge0.dp0 create 192.168.10.10/24
portcfg ipif ge1.dp0 create 192.168.10.11/24
```

2. Create a tunnel.
```
portcfg fciptunnel 12 create --ipsec MyIPsecPolicy --compression deflate
```
3. Create two circuits.
```
portcfg fcipcircuit 12 create 0
portcfg fcipcircuit 12 create 1
```
4. Add IP addresses and 1Gb/s bandwidth to each circuit.
```
portcfg fcipcircuit 12 modify 0 --local-ip 192.168.10.10 --remote-ip 192.168.10.20
portcfg fcipcircuit 12 modify 1 --local-ip 192.168.10.11 --remote-ip 192.168.10.21

portcfg fcipcircuit 12 modify 0 --min 1000000 --max 1000000
portcfg fcipcircuit 12 modify 1 --min 1000000 --max 1000000
```
5. Display tunnels, circuits, bandwidth, failover metrics, and group settings.
```
portshow fciptunnel -c
```

NOTE
To reach the gateway to the destination, you must configure IP routes if the source and destination addresses are located on different subnets.

Active-Passive Circuits

The following example shows an active-passive configuration with two circuits. Circuit 0 has the default metric of 0 and circuit 1 has a metric of 1. Both circuits are grouped into failover group 1. One circuit uses ge0 and the other circuit uses ge1. In this example, circuit 1 is the failover circuit because it has metric 1. When circuit 0 goes down, the traffic fails over to circuit 1. The effective bandwidth of the tunnel in this example is 1Gb/s.

Perform the following steps to configure an active-passive circuit:

1. Create IPIFs.
```
portcfg ipif ge0.dp0 create 192.168.10.10/24
portcfg ipif ge1.dp0 create 192.168.10.11/24
```
2. Create a tunnel.
```
portcfg fciptunnel 12 create --ipsec MyIPsecPolicy --compression deflate
```
3. Create two circuits.
```
portcfg fcipcircuit 12 create 0
portcfg fcipcircuit 12 create 1
```
4. Add IP addresses and 1Gb/s bandwidth to each circuit.
```
portcfg fcipcircuit 12 modify 0 --local-ip 192.168.10.10 --remote-ip 192.168.10.20
portcfg fcipcircuit 12 modify 1 --local-ip 192.168.10.11 --remote-ip 192.168.10.21

portcfg fcipcircuit 12 modify 0 --min 1000000 --max 1000000
portcfg fcipcircuit 12 modify 1 --min 1000000 --max 1000000
```
5. Add metrics and groups to each circuit.
```
portcfg fcipcircuit 12 modify 0 --metric 0 --failover-group 0
portcfg fcipcircuit 12 modify 1 --metric 1 --failover-group 0
```
6. Display circuit and interface details for ge0.
```
portshow fcipcircuit -c
```

NOTE
If the source and destination addresses are on different subnets, you must configure IP routes to the destination addresses.

Circuit Spillover

The spillover feature is supported on all Brocade Extension platforms. The spillover feature enables you to configure primary circuits that are regularly used, and secondary circuits that are only used during times of a high utilization. When circuits are configured for spillover, the metric-0 circuits are utilized until the configured bandwidth is reached, after which, the remaining traffic is sent over the higher metric-1 circuits. For example, consider three 1Gb/s circuits that are configured on a tunnel as follows:

- Two circuits are metric-0 circuits
- One circuit is a metric-1 circuit
- Each circuit supports a maximum of 1Gb/s

In this configuration, if a host traffic is started at 2Gb/s, it runs over the metric-0 circuits. If the host traffic increases to 2.5Gb/s, the metric-0 circuits continue to support the 2Gb/s traffic and the additional 500Mb/s of traffic runs over the metric-1 circuit.

NOTE
Failover groups are not used when the tunnel is configured for spillover, and any failover groups that are defined are ignored. Spillover behavior is similar to failover behavior, in that if metric 0 circuits are not available, metric 1 circuits are used.

Circuit spillover is a load-balancing method and is defined at the tunnel level, not per circuit. This feature allows circuits to be used only when there is congestion or during periods of a high-bandwidth utilization.

Primary circuits (metric 0) are specified and these circuits are the normal operation production circuits. Secondary circuits (metric 1) are specified and these circuits are the spillover circuits. Spillover circuits are used during high-utilization or congestion periods.

Primary and secondary circuits are controlled using the metric field of the circuit. For example, when a tunnel is configured for spillover, traffic uses the metric 0 circuits until their maximum bandwidth is reached. When the bandwidth in the metric 0 circuits is exceeded, both the metric 0 and metric 1 circuits are used. Conversely, when the utilization drops below the maximum on the metric 0 circuits, the metric 1 circuits are no longer used.

Tunnel failover remains intact. If a tunnel has spillover configured and the metric 0 circuits within a failover group go offline, the metric 1 circuits within the group are activated. Whether configured for failover or spillover, metric 1 circuits behave as backup circuits. In comparison to spillover, failover requires all metric 0 circuits to go offline before the metric 1 circuits are used within a failover group.

Circuit failover groups can be configured along with spillover and behave the same as without spillover configured; however, only the metric value determines whether a circuit is considered a spillover circuit or a primary circuit. Failover groups do not affect which circuits are used for spillover. For example, if a tunnel has four circuits, group 0 has a metric 0 and a metric 1 and group 1 has a metric 0 and a metric 1. When the metric 0 circuits become saturated, the metric 1 circuits start to be used. Also, if the metric 0 circuit in group 0 goes offline, the metric 1 circuit becomes active, and the same would be true in group 1 if its metric 0 went offline.

Due to the way spillover is implemented, the observed behavior might not be as expected. As an example, consider traffic flows with QoS of high, medium, and low, each with a bandwidth percentage (50/30/20). QoS traffic is carried by the active circuits according to the percentage allocated to each priority. If the QoS low-priority traffic reaches saturation before hitting its bandwidth allocation limit, it spills over from a metric 0 circuit to a metric 1 circuit.

For example, consider QoS traffic flows designated as high, medium, and low. The high QoS traffic flow can be assigned to metric 1 circuits, while the medium and low QoS traffic flow can be assigned to metric 0 circuits. In this instance, the spillover circuits (metric 1) are used even though the metric 0 circuits are not saturated. When the metric 0 circuits are saturated, additional traffic will spillover to the metric 1 circuits.

Using the `portshow fciptunnel` and `portshow fcipcircuit` commands with the `--qos` and `--perf` options display more details.

NOTE
Spillover has a tunnel scope. In the flag list for both tunnel and circuit output, a spillover flag (s = Spillover) is listed. The spillover flag applies only to output for the `portshow fciptunnel` command.

Understanding Circuit Spillover Utilization

To verify the spillover usage of metric 1 circuits, you can view the protocol data unit (PDU) and byte count information that is maintained for each circuit. The following `portshow` commands can be used to verify the spillover usage of metric 1 circuits. Using the `--qos` option provides more detail about which QoS tunnel that the spillover is occurring on. Some commands provide PDU/byte counts, and others provide the throughput information:

```
portshow fciptunnel --circuits
portshow fciptunnel --circuits --qos
portshow fciptunnel 24 -c
portshow fciptunnel 24 -scq --perf
```

Use the `portshow fciptunnel` and `portshow fcipcircuit` commands with and `--perf` option to display the main utilization values used by the spillover algorithm. An example header and flag table for these commands are as follows:

```
Tunnel Circuit St Flg TxMBps RxMBps CmpRtio RTTms ReTx TxWAN% TxQ%/BW Met/G

<display output values not shown>

Flg (tunnel): c=Control h=HighPri m=MedPri l=LowPri, I=IP-Ext, s=Spillover St: High level state, Up or Dn
TxWAN (tunnel): Tx WAN utilization high of primary circuits (--qos for range) (circuit): Tx WAN utilization
 high (--qos for range)
TxQ (tunnel): Tx data buffering utilization high (--qos for range)
```

The utilization numbers can be better understood by understanding how resources are allocated. A VE has multiple logical QoS tunnels. Each of the tunnels is separated internally into independent resource groups, each with its own instantiation of the configured circuits on the VE, the TCP connections, and the prorated bandwidth allocations. The number of resource allocations is based primarily on the configured bandwidth.

The TxWAN and TxQ utilizations numbers shown in the output are calculated over a 10-second interval for each resource group. TxWAN represents the average utilization of all primary TCP connections within the group. TxQ represents the average use of the buffering mechanism provided by the group. TxQ provides a close estimate of the amount of data the host application is sending and data that is readily available (buffered) when bandwidth becomes available on the WAN. The presence of utilization suggests that there is a push back from an end device.

The values that are displayed for the tunnel and circuit represent the highest level that is measured within that object. The values that are displayed for the QoS tunnel level represent the range of these measurements, lowest, and highest measured.

In general, the WAN utilization of 96% or greater must be sustained for approximately 15 to 20 seconds before spillover is activated and the metric 1 circuits become spillover active. When the metric 1 circuits are activated, they are used only if all the data cannot be sent over the metric 0 circuits.

The utilization numbers might not be an exact value because they depend on the timing of how the data is collected. For example, a circuit might show 94%, whereas the tunnel might show 93% high utilization. The circuit utilization is measured independently across all resource groups and the tunnel utilization is measured as a single resource group.

Use the `--qos` option to determine whether throughput spillover is active. You can determine if the resources are being used evenly, or which Quality of Service is being used, by examining the displayed utilization range. When the low-high utilization range is narrow, it indicates that the resource groups are being used equally for the QoS.

Spillover Examples

Using the following series of `--perf` options, we can see that utilization reaches spillover thresholds for one QoS but not for others that have traffic flowing over them.

```
switch:admin> portshow fciptunnel 24 -s --perf

Tunnel Circuit St   Flg TxMBps  RxMBps CmpRtio RTTms ReTx TxWAN% TxQ%/BW Met/G
-------------------------------------------------------------------------------
24     -            Up   --s    1138.8  6.2    -     -     0     99     0    -
-------------------------------------------------------------------------------
switch:admin> portshow fciptunnel 24 -sc --perf

Tunnel Circuit St  Flg TxMBps RxMBps CmpRtio RTTms ReTx TxWAN% TxQ%/BW Met/G
-------------------------------------------------------------------------------
24     -           Up --s 1138.2    6.2    -    -    0    99    0    -
24     0 ge0       Up --- 990.4     6.2    -    1    0    99    9500/9500    0/-
24     1 ge1       Up --- 147.9     0.0    -    1    0    13    9500/9500    1/-
-------------------------------------------------------------------------------
switch:admin> portshow fciptunnel 24 -scq --perf

Tunnel Circuit St Flg TxMBps RxMBps CmpRtio RTTms ReTx TxWAN% TxQ%/BW Met/G
-------------------------------------------------------------------------------
24     -          Up    c-s    0.0    0.0    -    -    0    0-0    0-0       -
24     0    ge0   Up    ---    0.0    0.0    -    1    0    0-0    0/9500    0/-
24     1    ge1   Up    ---    0.0    0.0    -    1    0    0-0    0/9500    1/-
24     -          Up    h-s    712.6  0.7    -    -    0    99-99  0-0       -
24     0    ge0   Up    ---    564.4  0.7    -    1    0    99-99  2375/9500 0/-
24     1    ge1   Up    ---    148.2  0.0    -    1    0    12-13  2375/9500 1/-
24     -          Up    m-s    0.0    0.0    -    -    0    0-0    0-0       -
24     0    ge0   Up    ---    0.0    0.0    -    1    0    0-0    1425/9500 0/-
24     1    ge1   Up    ---    0.0    0.0    -    1    0    0-0    1425/9500 0/-
24     -          Up    l-s    0.0    0.0    -    -    0    0-0    0-0       -
24     0    ge0   Up    ---    0.0    0.0    -    1    0    0-0    950/9500  0/-
24     1    ge1   Up    ---    0.0    0.0    -    1    0    0-0    950/9500  1/-
24     -          Up    hIs    426.7  5.5    -    -    0    72-82  0-0       -
24     0    ge0   Up    ---    426.7  5.5    -    1    0    72-82  2375/9500 0/-
24     1    ge1   Up    ---    0.0    0.0    -    1    0    0-0    2375/9500 1/-
24     -          Up    mIs    0.0    0.0    -    -    0    0-0    0-0       -
24     0    ge0   Up    ---    0.0    0.0    -    1    0    0-0    1425/9500 0/-
24     1    ge1   Up    ---    0.0    0.0    -    1    0    0-0    1425/9500 1/-
24     -          Up    lIs    0.0    0.0    -    -    0    0-0    0-0       -
24     0    ge0   Up    ---    0.0    0.0    -    1    0    0-0    950/9500  0/-
24     1    ge1   Up    ---    0.0    0.0    -    1    0    0-0    950/9500  1/-
-------------------------------------------------------------------------------
```

In the following series, the utilization does not quite reach the spillover thresholds, so spillover does not occur.

```
switch:admin> portshow fciptunnel 25 -s --perf
Tunnel Circuit St Flg TxMBps RxMBps CmpRtio RTTms    ReTx    TxWAN%    TxQ%/BW    Met/G
-------------------------------------------------------------------------------
25     -       Up --s 505.4    1.8     -      -       0       84        0          -
-------------------------------------------------------------------------------
```

```
switch:admin> portshow fciptunnel 25 -sc --perf
Tunnel Circuit St Flg    TxMBps   RxMBps CmpRtio RTTms    ReTx    TxWAN%    TxQ%/BW     Met/G
-------------------------------------------------------------------------------------------
25    -        Up --s    505.2    1.8    -       -        0       85        0            -
25    1 ge5    Up ---    505.2    1.8    -       1        0       85        2000/5000    0/-
25    2 ge2    Up ---    0.0      0.0    -       1        0       0         2000/5000    1/-
-------------------------------------------------------------------------------------------
switch:admin> portshow fciptunnel 25 -sc --qos --perf

Tunnel Circuit  St      Flg    TxMBps   RxMBps CmpRtio RTTms   ReTx    TxWAN%     TxQ%/BW            Met/G
-----------------------------------------------------------------------------------------------------------
25    -          Up    c-s    0.0      0.0    -       -       0       0-0        0-0                -
25    1    ge5   Up    ---    0.0      0.0    -       1       0       0-0        0/5000             0/-
25    2    ge2   Up    ---    0.0      0.0    -       1       0       0-0        0/5000             1/-
25    -          Up    h-s    0.0      0.0    -       -       0       0-0        0-0                -
25    1    ge5   Up    ---    0.0      0.0    -       1       0       0-0        500/5000           0/-
25    2    ge2   Up    ---    0.0      0.0    -       1       0       0-0        500/5000           1/-
25    -          Up    m-s    506.7    1.8    -       -       0       85-85      0-0                -
25    1    ge5   Up    ---    506.7    1.8    -       1       0       85-85      300/5000           0/-
25    2    ge2   Up    ---    0.0      0.0    -       1       0       0-0        300/5000           1/-
25    -          Up    l-s    0.0      0.0    -       -       0       0-0        0-0                -
25    1    ge5   Up    ---    0.0      0.0    -       1       0       0-0        200/5000           0/-
25    2    ge2   Up    ---    0.0      0.0    -       1       0       0-0        200/5000           1/-
-----------------------------------------------------------------------------------------------------------
```

Circuit Spillover Considerations

In the event of a circuit spillover, the following factors should be considered:

- 5:1 Circuit Min/Max Rule: In the `portcfg fciptunnel` command, the value for the `--max-comm-rate` option cannot exceed 5 times the value for the `--min-comm-rate` option.
- The minimum supported circuit rate for the Brocade 7810 Extension Switch and Brocade SX6 Extension Blade is a minimum of 20Mb/s and for the Brocade 7850 Extension Switch is a minimum of 50Mb/s.
- When a tunnel is configured for spillover, all metric 1 circuits are treated as metric 0 circuits when calculating aggregate bandwidth restrictions.

To configure a spillover, use the `portcfg fciptunnel <ve> modify` command. If you do not configure a spillover, the default action is to use failover.

1. Connect to the switch and log on as admin.
2. Use the `portcfg fciptunnel` command as shown in the following example to modify an existing tunnel for spillover.

   ```
   switch:admin> portcfg fciptunnel 12 modify --load-leveling spillover
   Operation Succeeded
   ```

3. Use the `portshow fciptunnel <ve>` command to display the current load leveling algorithm. The value is checked on both ends; therefore, the configured, peer, and actual value (the value used) are displayed.

   ```
   switch:admin> portshow fciptunnel 12

   Tunnel: VE-Port:12 (idx:0, DP0)
   ================================================
   Oper State    : Online
   ...
   Load-Level (Cfg/Peer): Spillover (Spillover / Spillover)
   ```

4. The following example shows the configuration warnings when one side of the tunnel is configured with spillover and the other side is configured with failover:

```
switch:admin > portshow fciptunnel 12

Tunnel: VE-Port:12 (idx:0, DP0)
==================================================
Oper State    : Online Warning
...
Load-Level (Cfg/Peer): Failover (Failover / Spillover)
...
Configuration Warnings: Load Leveling (failover|spillover)

Output is truncated…
```

5. Send and receive packets, bytes, uptime, Bps, RTT, and other statistics are maintained for each circuit and can be used to verify the spillover usage on metric 1 circuits. The `--qos` option provides a more detailed view as to which QoS priority that the spillover is occurring on.

   Following are some examples of commands commonly used to display the tunnel and circuit status and statistics. Due to the variable nature of output based on configuration and traffic conditions, only the commands are displayed.

```
switch:admin> portshow fciptunnel 12
switch:admin> portshow fciptunnel 12 -c
switch:admin> portshow fciptunnel 12 -c -q -s
switch:admin> portshow fcipcircuit 12
switch:admin> portshow fcipcircuit 12 -q
```

Example of Spillover Circuits at Various Rates

The following example shows the expected behavior of a tunnel that is configured with spillover, and the circuits at various rates. The maximum output of the application is 309Mb/s. The primary circuit-0 is initially configured at 3Gb/s. The spillover circuit-1 is configured for 3Gb/s. Because circuit-0 can support the 309MB/s max output, the spillover circuit-1 is not used. To facilitate a spillover, the bandwidth for circuit-0 is stepped down, the extra bandwidth is spilled over into circuit-1, and the overall throughput is maintained.

```
switch:admin> portshow fciptunnel -c

Tunnel Circuit OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt   CommRt   Met/G
-------------------------------------------------------------------------------
12     -       Up      --------I  17m52s  309.05  0.00   2          -          -
12     0 ge2   Up      ----a---4  17m51s  309.05  0.00   2      3000/3000     0/-
12     1 ge3   Up      ----a---4  17m52s  0.00    0.00   2      3000/3000     1/-
-------------------------------------------------------------------------------


Tunnel Circuit OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt   CommRt   Met/G
-------------------------------------------------------------------------------
12     -       Up      --------I  20m40s  306.29  0.00   2          -          -
12     0 ge2   Up      ----a---4  20m40s  285.33  0.00   2       2400/2400    0/-
12     1 ge3   Up      ----a---4  20m40s   20.96  0.00   2       3000/3000    1/-
-------------------------------------------------------------------------------

Tunnel Circuit OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt   CommRt   Met/G
-------------------------------------------------------------------------------
12     -       Up      --------I  23m30s  305.17  0.00   2          -          -
```

```
12     0 ge2   Up      ----a---4  23m30s  249.70  0.00   2       2100/2100    0/-
12     1 ge3   Up      ----a---4  23m31s   55.47  0.00   2       3000/3000    1/-
-------------------------------------------------------------------------------


Tunnel Circuit OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt   CommRt   Met/G
-------------------------------------------------------------------------------
12     -       Up      --------I  24m52s  304.78  0.00   2           -         -
12     0 ge2   Up      ----a---4  24m51s  178.11  0.00   2       1500/1500    0/-
12     1 ge3   Up      ----a---4  24m52s  126.67  0.00   2       3000/3000    1/-
-------------------------------------------------------------------------------
```

Service-Level Agreement (SLA)

The service-level agreement (SLA) feature provides automated testing of a circuit before it is put into service. An SLA works with the existing WAN Tool, which is supported on all the Brocade Extension platforms. When using an SLA, you can define a set of network service-level parameters, such as packet loss, and can make certain that the network meets those parameters before bringing the circuit into service.

When configured, SLA automatically invokes the WAN Tool and runs a test on the circuit. The WAN Tool uses the same circuit configuration parameters to verify the network path condition, so no additional configuration is needed. After the network parameters are met, the WAN Tool session stops, and the circuit is placed into service. If the circuit bounces, it reverts to test mode, and the SLA parameters must be met before the circuit is enabled again and placed in service.

The SLA checks circuits for a packet loss; if other circuit verifications are required such as throughput, congestion, and out-of-order delivery, use WAN Tool to manually run tests.

After configuring an SLA, you must assign the SLA to a specific circuit with the `portcfg fcipcircuit` command. An SLA session must be configured at each end of the circuit being tested. If the circuit configurations specify different transmission rates, the SLA negotiates and uses the lower configured rate. This allows the SLA to start even when circuit configurations have a mismatch. If the circuit configurations specify different packet loss values, the SLA negotiates and uses the lower (lower the packet loss value, better the circuit quality) of the two values. This ensures that the circuit adheres to better quality standards. When the session is established, traffic starts automatically. During the test, the traffic must remain under the specified packet loss percentage before the circuit can be placed into service.

NOTE

- During eHCL, SLA is disabled and no SLA sessions are created until after all eHCL operations have completed. After all eHCL operations are complete, SLA is re-enabled.
- On the Brocade 7810 Extension Switch, when you have SLA enabled on a circuit configuration that maxes out the primary and secondary bandwidth limits (2.5Gb/s of metric 0 with 2.5Gb/s of metric 1), the total running throughput might exceed the 2.5Gb/s, which is the platform maximum. In this configuration, SLA use should be avoided or used with caution.

SLA interaction while it is running is limited. Statistics can be viewed and an active session can be aborted. However, an active session cannot be modified.

After a circuit has gone down, there are two ways a circuit can be placed back into service. One, for the packet loss to remain under the specified loss percentage for the amount of time set by the SLA runtime; the circuit is placed backed into service. Two, the packet loss did not remain under the specified loss percentage during runtime; however, the SLA timeout was set and expired; the circuit is placed back into service. The SLA timeout value must either be greater than the runtime or set to zero. A timeout entry of zero equals no timeout, and the default is no timeout. If a timeout value has not been set, SLA testing continues until the packet loss remains under the threshold for the specified runtime. The default runtime is 5 minutes.

- Brocade 7850 Extension Switch: Up to 20 SLA session can be defined per DP
- Brocade SX6 Extension Blade: Up to 20 SLA sessions can be defined per DP
- Brocade 7810 Extension Switch: Up to 12 SLA sessions can be defined

Attempts to modify an active session are blocked. WAN Tool commands cannot be used while an SLA session is running. Active SLA sessions can be aborted and SLA statistics can be viewed.

Configured SLA sessions are persistent across reboots because they are part of the circuit configurations; unlike, user configured WAN Tool sessions, which are not persistent.

The following steps show how to configure, display, and abort SLA sessions. You can apply an SLA to multiple circuits, and you can configure multiple SLAs.

1. Use the `portcfg sla` command to create an SLA session. SLA sessions are created at each end of a circuit; the session names do not need to match.
   - The SLA named `networkA` is created with a packet-loss limit of 0.5 percent. It runs for the default 5 minutes and the default never times out.
   - The SLA named `networkB` is created with a packet-loss limit of 1 percent. It will run for 10 minutes and timeout after 30 minutes.

   ```
   switch:admin> portcfg sla networkA create --loss 0.5
   switch:admin> portcfg sla networkB create --loss 1 --runtime 10 --timeout 30
   ```
2. Use the `portcfg fcipcircuit` command to assign an SLA to a circuit. The following command modifies a circuit by assigning the SLA networkA to the circuit. Remember to configure the other end of the circuit with a matching SLA.
   ```
   switch:admin> portcfg fcipcircuit 24 modify 0 --sla networkA
   ```
3. Use the `portshow fciptunnel -c` command to display the status of tunnel and circuits and see whether the SLA session is actively testing the circuit.
   ```
   switch:admin> portshow fciptunnel -c

   Tunnel Circuit OpStatus Flags    Uptime TxMBps RxMBps ConnCnt CommRt    Met/G
   ------------------------------------------------------------------------------
   24     -       Degrad   -------I 1m36s    0.00   0.00    1      -         -/-
   24     0 ge2   Test     -S-ah--4 0s       0.00   0.00    0  5000/10000    0/-
   24     1 ge3   Up       ---ah--4 1m36s    0.00   0.00    1  5000/10000    0/-
   ------------------------------------------------------------------------------
   Flags (tunnel): l=Legacy QOS Mode
                   i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                   a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                   I=IP-Ext
        (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                   ARL a=Auto r=Reset s=StepDown t=TimedStepDown  S=SLA
   ```
   The sample output assumes that both ends of the circuit are configured with matching SLAs: packet loss, runtime, and timeout values are the same. Any difference in the transmit speed between the endpoints is chosen to be the lower value. The output shows tunnel 24, circuit 0 is under active test and the flags show that an SLA has been assigned.
4. Use the `portcmd --wtool show --detail` command to display details about active WAN Tool sessions. SLA invokes WAN Tool to run its tests, so the active SLA session are displayed.
   ```
   switch:admin> portcmd --wtool show --detail

   WTool Session: 24.0 (DP0)
   =========================================
   Admin / Oper State    : Enabled / Running
   Up Time               : 10s
   Run Time              : 9s
   ```

```
Time Out              : 3m50s
Time Remaining        : 1m51s
IP Addr (L/R)         : 213.70.2.10 ge2 <-> 213.70.2.20
IP-Sec Policy         : (none)
PMTU Discovery (MTU)  : disabled (1500)
Bi-Directional        : disabled
L2CoS / DSCP          : (none) / (none)
Configured Comm Rate  : 1000000 kbps
Peer Comm Rate        : 1000000 kbps
Actual Comm Rate      : 1000000 kbps
Tx rate               : 999624.45 Kbps ( 124.95 MB/s)
Rx rate               : 1000000.00 Kbps ( 125.00 MB/s)
Tx Utilization        : 99.96%
Rx Utilization        : 100.00%
RTT (Min/Max)         : 1 ms/1 ms
RTT VAR (Min/Max)     : 1 ms/1 ms Local
Session Statistics
Tx pkts               : 810024
Peer Session Statistics
Rx pkts               : 792029
Ooo pkts              : 0
Drop pkts             : 0 (0.00%)
```

When an SLA session completes, information is displayed in the `portcmd --wtool show --detail` command that displays when the session either completed or was stopped, and the completion reason. The following partial command output provides an example of a session that was aborted.

```
switch:admin> portcmd --wtool show --detail

WTool Session: 24.0 (DP0)
=================================================
Admin / Oper State     : Disabled / Disabled
Last Session End       : Thu Feb 23 07:12:31 2017
Last Session Completion: <Pass:Fail(reason):Aborted>
(output truncated)
```

5. Use the `portcmd --wtool show --sla` command to display the summary information about active SLA sessions. The main advantage of the summary display is to see the amount of time remaining for each session.

```
switch:admin> portcmd --wtool show --sla

Session      OperSt           TxMBps   RxMBps     Drop%    RunTime TimeOut   TimeRemaining
24.0         Running          6.22     6.25       0.00     5s      3m55s     1m56s
24.1         Running          6.26     6.25       0.00     5s      19m55s    1m56s
24.2         Running          6.22     6.25       0.00     4s      19m56s    1m57s
24.3         Running          6.20     6.23       0.00     3s      19m57s    1m58s
25.0         Running          12.52    12.45      0.00     3s      19m57s    3m58s
25.1         Running          12.50    12.50      0.00     4s      19m56s    3m57s
25.2         Running          12.47    12.51      0.00     4s      19m56s    3m57s
```

6. Use the `portcmd --wtool <session> stop` command to abort an SLA session. The session ID information is obtained from the `portcmd --wtool show` command. The `portcmd --wtool show --detail` command to display session details and the completion reason.

```
switch:admin> portcmd --wtool show

Session OperSt  Flags   LocalIP             RemoteIp            TxMBps RxMBps Drop%
--------------------------------------------------------------------------------
```

```
24.0    Running ----4S- 176.196.2.2          177.195.2.2          12.52  12.50  0.00
24.1    Running ----6S- 2002:176:196:dead::2 2002:177:195:dead::2 12.48  12.50  0.00
24.2    Running ----4S- 176.196.3.2          177.195.3.2          12.51  12.50  0.00
24.3    Running ----6S- 2003:176:196:dead::2 2003:177:195:dead::2 12.52  12.50  0.00
24.4    Running ----4S- 176.196.4.3          177.195.4.3          12.47  12.50  0.00
24.5    Running ----6S- 2004:176:196:dead::3 2004:177:195:dead::3 12.52  12.49  0.00
24.6    Running ----4S- 176.196.5.3          177.195.5.3          12.51  12.50  0.00
24.7    Running ----6S- 2005:176:196:dead::3 2005:177:195:dead::3 12.50  12.49  0.00
24.8    Running ----4S- 176.196.5.4          177.195.5.4          12.52  12.49  0.00
24.9    Running ----6S- 2005:176:196:dead::4 2005:177:195:dead::4 12.51  12.50  0.00
-------------------------------------------------------------------------------
Flags (wtool): S=SLA v=VLAN i=IPsec 4=IPv4 6=IPv6 L=Listener I=Initiator
```

Individual TCP sessions can be displayed for each session with the following command:

VE.Circuit 24.0 is 1Gb/s, and WAN Tool session 15 consists of multiple connections making up 2.5Gb/s.

```
switch:admin> portcmd --wtool show -c

Session OperSt  LocalIP/sPort      RemoteIP/dPort     TxMBps  RxMBps  Drop%
-------------------------------------------------------------------------------
15      Running 10.10.10.76/3225   10.10.10.77/62722  62.50   62.50   0.00
15      Running 10.10.10.76/3226   10.10.10.77/64673  62.47   62.50   0.00
15      Running 10.10.10.76/3226   10.10.10.77/64676  62.50   62.50   0.00
15      Running 10.10.10.76/3225   10.10.10.77/62716  62.46   62.50   0.00
24.0    Running 176.196.3.2/3226   177.195.3.2/61986  62.43   62.48   0.00
24.0    Running 176.196.3.2/3225   177.195.3.2/63957  62.50   62.49   0.00

switch:admin> portcmd --wtool 24.0 stop

The WAN Tool session 24.0 is aborted.

switch:admin> portcmd --wtool show --detail

WTool Session: 24.0 (DP0)
=================================================
Admin / Oper State    : Disabled / Disabled
Last Session End    : Thu Feb 23 07:12:31 2023
Last Session Completion: <Pass:Fail(reason):Aborted>
(output truncated)
```

7. Use the `portcmd --wtool stop-all` command to abort all running SLA sessions.

```
switch:admin> portcmd --wtool stop -all
```

Extension QoS

QoS refers to policies for expediting particular traffic flows over other flows. These policies are based on application delivery requirements. For example, email data is tolerant of delays and dropped packets, but real-time voice and video are not. Some storage replication applications are more sensitive to delay than others. QoS settings provide a framework for accommodating these differences in data as it passes through the extension tunnel or IP network.

The Brocade Extension platform has several Quality of Service (QoS) features. The Brocade Extension platform has two general categories of QoS functionality. First, there is a two tier traffic classification into FCIP/IP Extension and second each of those into high, medium, and low priorities. The second category is traffic marking with DSCP and/or L2CoS, which are used by the IP network.

This section describes these features and configuration.

There are two QoS tiers for WAN bandwidth allocation, protocol distribution, and priority ratio. Each extension protocol, FCIP, and IP Extension, have a QoS distribution. By default each protocol receives 50% of the bandwidth. The distribution for a protocol can be set as low as 10%, and the total must equal 100%. If a protocol path is utilized, the configured percentage of bandwidth is reserved. The other protocol cannot use the reserved bandwidth. If a protocol path is not utilized, for example, FCIP is used and IP Extension is not, then the IP Extension bandwidth is not reserved and FCIP can consume all the bandwidth regardless of the configured settings. The same is true if IP Extension is used and FCIP is not.

The following figure shows the internal architecture of TCP connections that handle Priority TCP QoS (PTQ). The figure illustrates a tunnel containing a single circuit. Hybrid mode provides three QoS priorities for each FCIP and IP Extension, a total of six QoS priorities, plus one for Class-F.

Figure 14: TCP Sessions for QoS

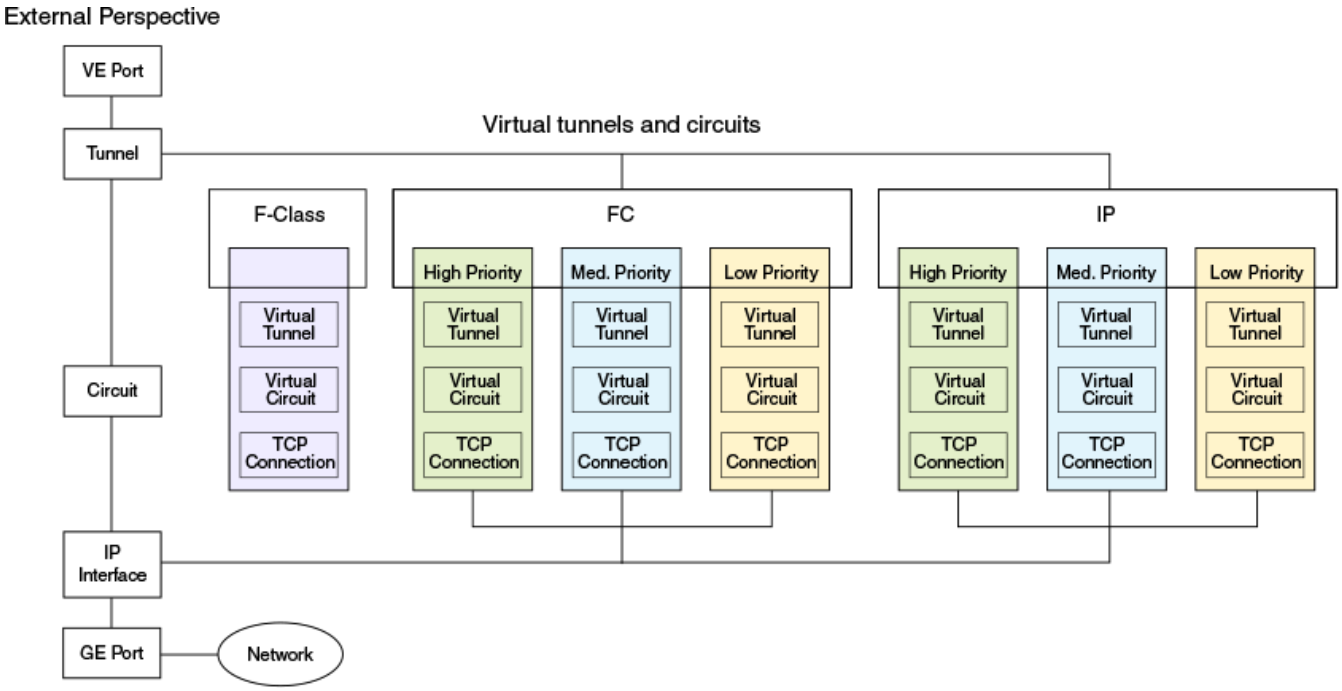

NOTE
Modifying QoS distribution and priority ratios of a tunnel is not disruptive to the circuits or tunnel, this action only changes the QoS priority distribution within the tunnel; it does not reconfigure the fabric.

Protocol Bandwidth Distribution

The first tier of Extension QoS, each protocol, FCIP, and IP Extension, receives a proportion of the overall configured bandwidth; this is referred to as protocol distribution. Distribution allows the apportioning bandwidth between FCIP and IP Extension. FCIP and IP Extension protocol distribution is rate-limiting to bandwidth. If there is utilization of bandwidth within a distribution, the entire percentage of bandwidth becomes reserved for that distribution and is not available to the other distribution. Distribution occurs at the tunnel level, not per circuit. For example, you can specify a QoS distribution ratio of FCIP 60% and IP Extension 40%. The default distribution is 50/50%. Allocation for either FCIP or IP Extension cannot exceed 90%. QoS allocations must total 100%.

Consider the following when you configure QoS protocol bandwidth distribution percentages:

- The distribution (FCIP:IP Extension) must equal 100%
- 10% is the minimum value
- 90% is the maximum value
- Distribution ratios must be the same on both ends of a tunnel

NOTE

- Protocol distribution, and FCIP and IP Extension circuit QoS percentages are rate-limiting to bandwidth. If there is utilization of bandwidth in a distribution or QoS path, the entire percentage of bandwidth is reserved and not available to other paths.
- If your storage device supports Fibre Channel CS_CTL prioritization, you can use the CS_CTL values in the FC header to prioritize the QoS traffic. For more information, refer to the *Brocade Fabric OS Administration Guide*.
- Changing the protocol distribution, resets the high, medium, low ratios to their default percentages (50%, 30%, and 20%).

1. Connect to the switch and log on as admin.
2. Use the `portcfg fciptunnel <ve> modify --distribution` command to change the protocol bandwidth distribution. The default is 50:50.

The following example modifies the protocol bandwidth distribution to FCIP 60% and IP Extension to 40%. Using the `modify` keyword facilitates staging configuration of the tunnel, as well, the distribution can be specified when the tunnel is first created.

```
switch:admin> portcfg fciptunnel 24 modify --distribution 60,40

Warning: Modification of the distribution ratio will reset the QoS ratio values to default.

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fciptunnel specified for a brief period of time. This
 operation will bring the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n] y
```

Priority Bandwidth Ratios

The second tier of Extension QoS, Brocade Extension platforms support FCIP and IP Extension and each has their own high, medium, and low QoS within their distribution. The Brocade Fabric OS platforms provide the traffic classification into high, medium, and low priorities. QoS priority ratios have a tunnel scope, and are configured at the tunnel level, not per circuit. A set of internal virtual tunnels is created for both protocols and each QoS priority. Created within each circuit are seven virtual tunnels, one F-Class, three for IP Extension (H/M/L), and three for FCIP (H/M/L). Medium priority is the default for both protocols. Brocade Extension QoS uses PerPriority-TCP-QoS (PTQ), which opens a dedicated WAN Optimized TCP session for the QoS priority of each protocol. Flows are automatically mapped to the TCP session of the corresponding priority. No TCP session carries more than one priority. The designated QoS priority is based on the virtual circuit (VC) that carries the data into the DP. For FC, zoning with QoS prefixes assigns the zoned flows to a specific set of VCs, in turn, VCs are mapped to extension priority TCP sessions. For example, if data enters through a high VC, it is placed in a high TCP session; if it enters through a low VC, it is placed in a low TCP session. For IP, the Traffic Control List (TCL) assigns matched flows to the correlating extension TCP session.

High, medium, and low are each assigned a percentage of bandwidth, the default is 50%, 30%, and 20%. The ratios must add up to 100%. If there are traffic flows on a priority, its bandwidth becomes reserved and is not available to other priorities. If only the medium priority (default) is being used, for example, no flows have been designated as high or low, all bandwidth is made available to medium, even if the medium queue remains configured at 30% (default). An egress queue is serviced once every millisecond, and the amount of traffic that is sent on that queue is based on the priorities being used and the configured ratios.

An Extension tunnel consists of the following virtual tunnels:

- Control virtual-tunnel
- High-priority data virtual-tunnel
- Medium-priority data virtual-tunnel
- Low-priority data virtual-tunnel

Each priority is permitted a percentage of the bandwidth as shown in the following table:

Table 31: Traffic Classes and Default Settings

Traffic Class	Default
Class-F	Control traffic and has a strict priority, the highest priority is sent first There is relatively little class-F traffic
Distribution FCIP	50% of bandwidth
Distribution IP Extension	50% of bandwidth
FCIP QoS High	50% of protocol's bandwidth
FCIP QoS Medium	30% of protocol's bandwidth
FCIP QoS Low	20% of protocol's bandwidth
IP Extension QoS High	50% of protocol's bandwidth
IP Extension QoS Medium	30% of protocol's bandwidth
IP Extension QoS Low	20% of protocol's bandwidth

The circuit adheres to the minimum bandwidth configured, and it receives that amount of bandwidth. For example, a 10Gb/s circuit used for both FCIP and IP Extension and experiencing WAN contention, by default, receives 1.5Gb/s for a medium QoS FCIP circuit (10Gb/s × 50% × 30% = 1.5Gb/s). Minimum bandwidth allocations are not oversubscribed. Circuits are limited to consuming the maximum bandwidth rate.

On the Brocade Extension platforms supporting eHCL (Brocade 7850 Extension Switch and Brocade SX6 Extension Blade), an additional tunnel group is created for eHCL. When eHCL is active during a firmware update, the high, medium, and low QoS data tunnels collapse into a single QoS priority. For more information, refer to Extension Hot Code Load (eHCL).

When configuring QoS protocol priority bandwidth percentages, note the following:

- The three priorities must total 100%
- 10% is the minimum value
- 90% is the maximum value
- Priorities must be the same on both ends of a tunnel

NOTE

- FCIP and IP Extension circuit QoS percentages are rate-limiting to bandwidth. If there is utilization of bandwidth in a QoS path, the entire percentage of bandwidth becomes reserved and not available to another QoS path.
- Modifying priority percentages is not disruptive to tunnel traffic.
- If your storage device supports Fibre Channel CS_CTL prioritization, you can use the CS_CTL values in the FC header to prioritize QoS traffic. For more information, refer to the *Brocade Fabric OS Administration Guide* .

1. Connect to the switch and log on as admin.

2. Use the `portcfg fciptunnel modify` command to change the QoS priority bandwidth allocations for FCIP and IP Extension traffic.
   a. Use the `portcfg fciptunnel modify --fc-qos-ratio` command to configure the FCIP QoS priorities.
      This command modifies the FCIP QoS priority bandwidth values in this order:
      - high 30%
      - medium 50%
      - low 20%

      ```
      switch:admin> portcfg fciptunnel 12 modify --fc-qos-ratio 30,50,20
      Operation Succeeded
      ```
   b. Use the `portcfg fciptunnel modify --ip-qos-ratio` command to configure the IP Extension QoS priorities.
      This command modifies IP Extension QoS priority bandwidth values in this order:
      1. high – 60%
      2. medium – 30%
      3. low – 10%

      ```
      switch:admin> portcfg fciptunnel 12 modify --ip-qos-ratio 60,30,10
      Operation Succeeded
      ```
3. Use the `portcfg fciptunnel --qos-bw-ratio` command to change the FCIP and IP Extension priority bandwidths.
   The following example modifies QoS bandwidth ratios for FCIP (high, med, low): 60%, 25%, 15% and for IP Extension (high, med, low): 60%, 20%, 20%. In order for each protocol to be successful, its three values must add up to 100%.
   ```
   switch:admin> portcfg fciptunnel 24 modify --qos-bw-ratio 60,25,15,60,20,20
   Operation Succeeded
   ```
   The following example modifies the priority QoS bandwidth ratios back to the original default values.
   ```
   switch:admin> portcfg fciptunnel 24 modify --qos-bw-ratio default
   Operation Succeeded
   ```
4. The following example displays the QoS protocol distribution and the FCIP and IP Extension QoS priority bandwidth values. Use the `portshow fciptunnel <ve>` command to display the configured distribution and priority QoS values.
   In the following example, the QoS values for tunnel-12 are configured as follows:
   - QoS Distribution: FC:50% / IP:50%
   - FC QoS BW Ratio: 50% / 30% / 20%
   - IP QoS BW Ratio: 50% / 30% / 20%

   ```
   switch:admin> portshow fciptunnel 12

    Tunnel: VE-Port:12 (idx:0, DP0)
    ==================================================
     Oper State           : Online
     TID                  : 12
     Flags                : 0x00000000
     IP-Extension         : Enabled
     Compression          : Aggressive Deflate
     FC-Compression       : Deflate (Override)
     IP-Compression       : Aggressive Deflate (Override)
     QoS Distribution     : Protocol (FC:50% / IP:50%)
     FC QoS BW Ratio      : 50% / 30% / 20%
     IP QoS BW Ratio      : 50% / 30% / 20%
   ```

```
Output is truncated…
```

5. Use the `portshow fciptunnel` command to display the QoS virtual circuits. A legend is shown at the end of output. The following example shows the QoS virtual circuit information. Over each circuit there is a virtual circuit for high, medium, and low, and a VC for the F-class control traffic. In this example, there are two circuits configured, circuit 0 and circuit 1.

```
switch:admin> portshow fciptunnel --qos -c

Tunnel Circuit   OpStatus  Flags       Uptime      TxMBps  RxMBps ConnCnt CommRt    Met/G
-------------------------------------------------------------------------------------------
 24    -         Up        c-i-----I    72d14h     0.00    0.00   9       -         0
 24    0 ge2     Up        ----ah-i4    72d14h     0.00    0.00   12      0/100     0/-
 24    1 ge3     Up        ----ah-i4    72d14h     0.00    0.00   9       0/100     0/-
 24    -         Up        h-i------    72d14h     0.00    0.00   9       -         1
 24    0 ge2     Up        ----ah-i4    3d 10h39m  0.00    0.00   12      12/100    0/-
 24    1 ge3     Up        ----ah-i4    72d14h     0.00    0.00   9       12/100    0/-
 24    -         Up        m-i------    72d14h     0.00    0.00   9       -         1
 24    0 ge2     Up        ----ah-i4    72d14h     0.00    0.00   12      7/100     0/-
 24    1 ge3     Up        ----ah-i4    72d14h     0.00    0.00   9       7/100     0/-
 24    -         Up        l-i------    72d14h     0.00    0.00   9       -         1
 24    0 ge2     Up        ----ah-i4    3d10h39m   0.00    0.00   12      5/100     0/-
 24    1 ge3     Up        ----ah-i4    72d14h     0.00    0.00   9       5/100     0/-
 24    -         Up        h-i----DI    72d14h     0.00    0.00   9       -         2
 24    0 ge2     Up        ----ah-i4    3d10h39m   0.00    0.00   12      12/100    0/-
 24    1 ge3     Up        ----ah-i4    72d14h     0.00    0.00   9       12/100    0/-
 24    -         Up        m-i----DI    72d14h     0.00    0.00   9       -         2
 24    0 ge2     Up        ----ah-i4    72d14h     0.00    0.00   12      7/100     0/-
 24    1 ge3     Up        ----ah-i4    72d14h     0.00    0.00   9       7/100     0/-
 24    -         Up        l-i----DI   72d14h      0.00    0.00   9       -         2
 24    0 ge2     Up        ----ah-i4   3d10h39m    0.00    0.00   12      5/100     0/-
 24    1 ge3     Up        ----ah-i4   72d14h      0.00    0.00   9       5/100     0/-
-------------------------------------------------------------------------------------------
 Flags (tunnel): c=Control h=HighPri m=MedPri l=LowPri
                 i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                 a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                 I=IP-Ext
      (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                 ARL a=Auto r=Reset s=StepDown t=TimedStepDown  S=SLA
```

QoS Marking

Whenever IPsec is enabled on a tunnel, all QoS priorities inside a circuit must use the same markings for QoS.

Traffic in the different priority TCP sessions can be marked with DSCP or 802.1P (L2CoS) based on the FCIP or IP Extension priority (high, medium, and low). IP Extension assigns a priority by matching a TCL rule. Data flows that have been assigned to high, medium, or low are marked accordingly with a particular DSCP and/or L2CoS value and each protocol (FCIP and IP Extension) can be marked independently.

QoS marking is configured per circuit, not per tunnel. Each circuit takes its own IP network route, which could have its own QoS paradigm. The default is no marking. Marking is not enforcement, it is up to the IP network to perform a QoS enforcement based on the markings. Therefore, it is essential to work with the network administrators to establish an enforcement.

The following two options are available for QoS marking:

- DSCP
- L2CoS (Layer 2 Class of Service, 802.1P)

IP marking (layer 3) is done using DSCP. Ethernet marking (layer 2) is done using 802.1P in the 802.1Q VLAN tag header.

DSCP Marking

Differentiated Services Code Point (DSCP) operates across IP (layer 3). Its implementation for establishing QoS policies is defined in RFC 2475. DSCP uses six bits of the Type of Service (TOS) field in the IP header to establish up to 64 different classes of service. The network administrators in an enterprise define and assign service classes and configure the network to handle the classed traffic accordingly. DSCP settings are useful only if routers and switches are configured to enforce QoS policies uniformly across the network. Control and high, med, low-priority flows can be configured with different DSCP values. Make sure you determine if the IP network implements Per Hop Behavior (PHB) and the appropriate DSCP values before configuring DSCP marking.

Perform the following steps to configure DSCP marking:

1. Connect to the switch and log on as admin.
2. Use the `portcfg fcipcircuit <ve> modify` command to set the DSCP marking.
   The following command modifies the FC High DSCP value to 46 for tunnel-24 circuit-0. DSCP values must be coordinated with networking to be effective.

```
switch:admin> portcfg fcipcircuit 24 modify 0 --dscp-high 46
Operation Succeeded
```

3. Use the `portshow fcipcircuit <ve>` command to display the configured values for DSCP.
   The following example shows the FC High DSCP value of 46 on tunnel-24 circuit-0.

```
switch:admin> portshow fcipcircuit 24

 Circuit 24.0 (DP0)
 ==================================================
  Admin/Oper State     : Enabled / Online
  Flags                : 0x00000000
  IP Addr (L/R)        : 10.0.0.16 ge2 <-> 10.0.1.16
  HA IP Addr (L/R)     : 10.0.0.17 ge2 <-> 10.0.1.17
  Configured Comm Rates: 100000 / 100000 kbps
  Peer Comm Rates      : 100000 / 100000 kbps
  Actual Comm Rates    : 100000 / 100000 kbps
  Keepalive (Cfg/Peer) : 5000 (5000 / 5000) ms
  Metric               : 0
  Connection Type      : Default
  ARL-Type             : Reset
  PMTU                 : Disabled
  HA PMTU              : Disabled
  SLA                  : (none)
  Failover Group       : 0
  VLAN-ID              : NONE
  L2Cos (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
  L2Cos (IP:h/m/l)     : 0 / 0 / 0
  DSCP (FC:h/m/l)      : 46/ 0 / 0 (Ctrl:0)
  DSCP (IP:h/m/l)      : 0 / 0 / 0
  cfgmask              : 0x40000000 0x00010caf
  Flow Status          : 0
  ConCount/Duration    : 33 / 138d20h
  Uptime               : 11h
```

```
 Stats Duration       : 11h
 Receiver Stats       : 632401248 bytes / 3124879 pkts /    4.87 KBps Avg
 Sender Stats         : 880290585 bytes / 3917984 pkts /    7.16 KBps Avg
 TCP Bytes In/Out     : 300185609897 / 352630995311
 ReTx/OOO/SloSt/DupAck: 25507 / 102368 / 13160 / 12836
 RTT (min/avg/max)    : 21 / 21 / 41 ms
 Wan Util             : 0.6%
```

L2CoS (Layer 2 Class of Service, 802.1P)

VLAN traffic uses 802.1Q tags within an Ethernet header. The tag includes a VLAN ID and Class of Service (CoS) priority bits. To set L2CoS, the Brocade Extension IPIF must be configured with a VLAN ID. The 802.1P L2CoS priority scheme uses three priority bits, which allows eight priorities. Consult with your WAN administrator to determine if setting L2CoS is useful and which L2CoS value can be used.

1. Connect to the switch and log on as admin.
2. Use the `portcfg fcipcircuit <ve> modify <cir_num>` command to configure the L2CoS values. The following example modifies tunnel-24 circuit-0 to configure L2CoS values.

   NOTE
   To configure L2CoS, the IPIFs of the circuit must have a VLAN ID configured.

```
switch:admin> portcfg fcipcircuit 24 modify 0 \
--l2cos-f-class 32 \
--l2cos-high 16 \
--l2cos-medium 8 \
--l2cos-low 4
Operation Succeeded
```

3. Use the `portshow fcipcircuit <slot/ve>` command to display the L2CoS values.
   The following example shows the FC L2CoS values for tunnel-24 circuit-0:

```
switch:admin> portshow fcipcircuit 24

 Circuit 24.0 (DP0)
 ==================================================
  Admin/Oper State     : Enabled / Online
  Flags                : 0x00000000
  IP Addr (L/R)        : 10.0.0.16 ge2 <-> 10.0.1.16
  HA IP Addr (L/R)     : 10.0.0.17 ge2 <-> 10.0.1.17
  Configured Comm Rates: 100000 / 100000 kbps
  Peer Comm Rates      : 100000 / 100000 kbps
  Actual Comm Rates    : 100000 / 100000 kbps
  Keepalive (Cfg/Peer) : 5000 (5000 / 5000) ms
  Metric               : 0
  Connection Type      : Default
  ARL-Type             : Reset
  PMTU                 : Disabled
  HA PMTU              : Disabled
  SLA                  : (none)
  Failover Group       : 0
  VLAN-ID              : NONE
  L2Cos (FC:h/m/l)     : 16/ 8 / 4 (Ctrl:32)
  L2Cos (IP:h/m/l)     : 0 / 0 / 0
  DSCP (FC:h/m/l)      : 46/ 0 / 0 (Ctrl:0)
  DSCP (IP:h/m/l)      : 0 / 0 / 0
```

```
cfgmask              : 0x40000000 0x00010caf
Flow Status          : 0
ConCount/Duration    : 33 / 138d20h
Uptime               : 11h
Stats Duration       : 11h
Receiver Stats       : 632401248 bytes / 3124879 pkts /    4.87 KBps Avg
Sender Stats         : 880290585 bytes / 3917984 pkts /    7.16 KBps Avg
TCP Bytes In/Out     : 300185609897 / 352630995311
ReTx/OOO/SloSt/DupAck: 25507 / 102368 / 13160 / 12836
RTT (min/avg/max)    : 21 / 21 / 41 ms
Wan Util             : 0.6%
```

WAN Side IP Network Considerations

Brocade Extension uses TCP connections over the WAN. Consult the WAN carrier and IP network administrator to ensure that TCP connections are supported by the network equipment on the data path.

Ensure that routers and firewalls in the data path are configured to pass extension traffic through specific ports and to provide the throughput that replication through extension requires. Consider the following when you configure the routers and firewalls in the data path:

- Brocade IPsec uses Encapsulating Security Payload (ESP) and UDP port 500 (IKE).
- Brocade WO-TCP randomly selects an ephemeral (source or initiating) port between 49152 and 65535.
- Brocade WO-TCP uses well-known TCP ports 3225 and 3226.
- Brocade WO-TCP requires TCP URG flags to not be dropped or modified.
- Brocade Extension sets the IP DF bit (Do Not Fragment). Fragmentation in the IP network is not supported.
- Brocade Extension PMTU requires end-to-end ICMP echo and replies.

To enable the resiliency against a WAN failure or device outage, be sure that redundant network paths are available from end to end. Ensure that the underlying WAN infrastructure supports the redundancy and performance that is expected in your implementation. If you configure jumbo frames on the extension platforms, ensure that the IP network supports them.

NOTE
Brocade Extension sets the IP DF bit (Do Not Fragment). Fragmentation in the IP network is not supported. Use AUTO MTU if the MTU is unknown to avoid an unexpected circuit behavior that might occur when the IP network cannot support a large MTU value.

KATOV

KATOV is essential for error recovery and should be based on application requirements. Check with your IP storage application provider to determine the I/O timeout. When the KATOV sum from all circuits in a tunnel is slightly less than the I/O timeout, this is the appropriate KATOV. For example, a mirroring application has a 6-second I/O timeout. There are three circuits belonging to a VE_Port (three circuit members in the tunnel). Set the KATOV to 2 seconds on each circuit. This allows maximum retries overall available circuits before an I/O is timed out by the IP storage application. Often, a 2-second to 3-second KATOV is optimal for IP Extension and should be considered.

NOTE
KATOV supported RTT at 2 seconds and above is 250 ms with a maximum of 1% packet loss. KATOV supported RTT at 2 seconds or less is 200 ms with a maximum of 0.1% packet loss.

Consider the following when you configure the KATOV:

- KATOV is within the circuit scope, it is configured per circuit.
- The default non-FICON tunnel KATOV is 6 seconds.
- KATOV must be the same on both ends of a circuit. If the local and remote KATOV does not match, the tunnel uses the shorter of the configured values.
- FICON
  - A VE_Port used for FICON should be in a FICON LS.
  - A tunnel that is used for FICON should be created with the FICON argument.
  - FICON requires a KATOV of 1 second or less on each circuit.
  - Tunnels that are not originally created with the FICON argument must be modified to become FICON tunnels to correct the KATOV to 1 second or less. This will not happen automatically.
- The aggregate KATOV across all circuits must be less than the I/O timeout. If the I/O timeout is less than the aggregate KATOV, the I/O times out before the extension platform can recover. The Brocade Extension platform retries the I/O across each online circuit.

NOTE

- Changing the KATOV can be disruptive.
- The Brocade 7810 Extension Switch does not support FICON or FICON Acceleration.

For more information about options and value ranges available with the `portcgf fcipcircuit --keepalive-timeout <ms>` command, refer to the *Brocade Fabric OS Command Reference Manual*.

1. Use the `portcfg fcipcircuit <ve> modify <cir_num> --keepalive-timeout <ms>` command to change the KATOV value.
   The following example modifies tunnel-24 circuit-0 and changes the KATOV to 2 seconds (2000 milliseconds).

```
switch:admin> portcfg fcipcircuit 24 modify 0 -k 2000

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fciptunnel specified for a brief period of time. This
 operation will bring the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n]    y
Operation Succeeded
```

2. Use the `portshow fcipcircuit <ve>` command to display the circuit configuration.
   KATOV tunnel-24 circuit-0 setting in bold is shown in the following example:

```
Local_switch:admin> portcfgshow fcipcircuit 24

Circuit 24.0 (DP0)
==================================================
Admin/Oper State         : Enabled / --
Flags                    : 0x00000000
IP Addr (L/R)            : 192.168.5.2 <-> 192.168.1.2
HA IP Addr (L/R)         : 192.168.5.12 <-> 192.168.1.12
Configured Comm Rates    : 5000000 / 5000000 kbps
Peer Comm Rates          : 0 / 0 kbps
Actual Comm Rates        : 0 / 0 kbps
Keepalive                : 2000 ms
Metric                   : 0

Output is truncated…
```

Configuring the LAN Side (IP Extension)

Brocade IP Extension provides an IP storage with the benefits of protocol optimization, bandwidth management, encryption, and compression.

Brocade IP Extension is used primarily for IP storage applications, such as remote host-based or database-based replication, NAS replication, IP backups, and tape grids. IP Extension can use the same tunnel (VE_Port) and circuits that FCIP uses, or it can use its own. The LAN side connects to a data center LAN so that IP storage end devices can access IP Extension.

IP Extension uses VE_Ports as logical endpoints for a tunnel, even if no FC or FICON traffic is being transported. IP Extension forwarding and processing are done within the extension platform, then field-programmable gate arrays (FPGAs), and the data processors (DPs). IP Extension uses an IP Extension gateway on each DP as a communication port with the IP storage end-devices (L2 deployment) or the IP network (L3 deployment). The IP Extension Gateway becomes the gateway for data flows headed to the remote data center. At the IP Extension Gateway, end-device TCP sessions are locally terminated, termed TCP proxy, and reformed on the remote side. This process has the advantage of providing local acknowledgments and acceleration. WAN-Optimized TCP (WO-TCP) is used as the transport between data centers, not to the storage end-devices.

IP Extension and IP networking require an understanding of the following key points:

- LAN IPIFs serve as gateways for DPs, which is why they are referred to as IP Extension Gateways.
- IP Extension forms a tunnel between data centers allowing end-devices to communicate.
- The local IP Extension Gateway and remote IP Extension Gateway must be on different subnets.
- A single subnet cannot be extended from one data center to another.
- IP Extension optimizes the TCP-based traffic.
- Layer 2 Implementation
  - The LAN IPIF is local and in the same broadcast domain as the attached end-devices, typically using an Ethernet data center switch.
  - Brocade IP Extension acts as a next-hop gateway for the IP storage flows traversing two data centers. The next-hop gateway is managed by creating the IP Extension LAN IPIF plus configuring on the end-device a specific route for the remote IP address/subnet that forwards the traffic to that IP Extension Gateway (ipif).
- Layer 3 Implementation
  - IP Extension on the LAN side can forward traffic to a local router so the traffic can be forwarded to its final destination.
  - As end devices send traffic to their traditional network gateway, the router intercepts the data flow and redirects it to the IP Extension Gateway, most commonly Policy Based Routing (PBR).
  - IP Extension supports a route for destination subnets to be configured on the LAN side.
- LAN side GE interfaces support dynamic and static link aggregation groups (LAGs). A LAG is a group of physical Ethernet links that are treated as a single logical link.
- LAN and WAN side GE interfaces support Link Level Discovery Protocol (LLDP). In contrast, WAN side GE interfaces connect to the network that sends traffic to the other data center.
- IP Extension uses a traffic control list (TCL) to direct traffic to the appropriate tunnel. If no configured TCL rule matches the traffic, the default action is to deny all traffic.
- Double VLAN tagging, often referred to as QinQ or nested VLAN tagging, is not supported.

NOTE
Creating a TCL is required, even if a single tunnel (VE_Port) is implemented. IP Extension functions without a configured TCL at each end.

Brocade IP Extension is supported on the following platforms:

- Brocade 7810 Extension Switch
- Brocade 7850 Extension Switch
- Brocade SX6 Extension Blade in Brocade X6 or Brocade X7 Directors

IP Extension provides the Extension acceleration and encryption for the IP storage replication. The IP Extension gateway is on the LAN of the data center. The end devices on each end of an IP Extension network must be on different subnets. The WAN side connectivity between data centers can be on the same subnet [Layer 2, that is, Dense Wave Division Multiplexing (DWDM)] or across different subnets (Layer 3, that is, routed).

IP Extension receives the same benefits as traditional FCIP:

- High-speed encryption
- Bandwidth management and flow-control
- QoS
- Compression
- Load balancing
- Lossless failover/failback

IP Extension requires Brocade SX6 Extension Blade to be in hybrid mode. Enabling hybrid mode is disruptive because a blade reboot is required to load the hybrid mode image.

Brocade 7850 Extension Switch and Brocade 7810 Extension Switch only have hybrid mode; therefore, there is no hybrid mode setting to configure.

Consider the following before you configure IP Extension:

- The WAN side configuration should be completed before configuring IP Extension.
- IP Extension requires a TCL to be configured; the default is *deny all*.
- The limit of LAN side TCP connections for each platform must be considered.
  - Exceeding the number of supported TCP sessions prevents the occurrence of additional TCP sessions.
  - Brocade 7810 Extension Switch: 256 TCP and 32 UDP connections
  - Brocade 7850 Extension Switch: 512 TCP and 64 UDP connections per DP
  - Brocade SX6 Extension Blade: 512 TCP and 64 UDP connections per DP
- Balance the load on each DP and the number of TCP sessions. Estimate bandwidth and TCP usage when planning end-devices per DP, and determine which IP Extension Gateway will be used by the end-device.

Perform the following tasks to configure the IP Extension configuration:

- Create LAN side LAG (portchannel) on GE interfaces.
- Create LAN side IPIF (IP Extension Gateway). On LAN side IPIF, add VLAN ID if applicable.
- Create LAN side IP routes, if applicable (layer 3 deployment only).
- Verify the LAN side IP connectivity.
- Create TCL.

NOTE

- IP extension should be enabled for each individual tunnel, even if the platform is operating in hybrid mode. This can be done during the creation or modification of the tunnel.
- When tunnels and circuits have already been configured in FCIP mode, the configuration is carried over to hybrid mode in the supported VE mode. You are prompted to make the necessary changes if there is a configuration incompatibility.
- Ethernet interfaces on the LAN side of IP Extension do not participate in Layer 2 switching. All traffic is either sent over the tunnel or dropped based on the TCL. As a result, loops cannot be formed, and STP is not required or used.

Brocade SX6 Extension Blade Hybrid Mode

Configure the Brocade SX6 Extension Blade in hybrid mode to support FCIP and IP Extension, as FCIP is the default application mode. The Brocade 7810 Extension Switch and Brocade 7850 Extension Switch only have hybrid mode; therefore, no application mode configuration is required. If you plan to use IP Extension, configure hybrid mode during an initial deployment. Brocade SX6 Extension Blade must be in 10VE mode before enabling hybrid mode. Changing VE modes requires rebooting the platform and is disruptive. In the 10VE mode, it is normal for the unused VE_Ports to be disabled. Configuring hybrid mode requires a blade reboot and is disruptive to the blade.

Consider the following when you configure Brocade SX6 Extension Blade IP Extension:

- **App mode**: The Brocade SX6 Extension Blade must be in hybrid mode (not the default).
- **VE mode**: The Brocade SX6 Extension Blade must be in 10VE mode (the default).

LAN Side GE Interfaces

IP Extension cannot function without at least one LAN side GE interface, which requires one or more GE interfaces to be in LAN mode. IP storage replication ports communicate with IP Extension Gateways (LAN side IPIF) through the data center LAN. The IP Extension Gateways are accessible through the LAN side GE interfaces. GE interfaces default to the WAN side, which is designated as *FCIP* in a `switchshow` . Incoming LAN traffic is not recognized on a WAN side GE interface and is dropped.

All Brocade Extension platforms have speed settable GE interfaces. GE speeds are not auto-negotiated; extension GE interfaces operate at the configured speed and with the optic supporting that speed. GE optics might or might not be able to operate at more than one speed.

Table 32: Supported LAN side GE Interfaces

Extension Platform	Platform's GE Interfaces	LAN side supported?	Interface Speeds
Brocade 7810 Extension Switch	GE0-GE1	Yes	1GbE
	GE2-GE7	Yes	1GbE, 10GbE
Brocade 7850 Extension Switch	GE0-GE15	Yes	1GbE, 10GbE, 25GbE
	GE16-GE17	No	100GbE
Brocade SX6 Extension Blade	GE0-GE1	No	40GbE
	GE2-GE17	Yes	1GbE, 10GbE

Any existing interface configuration must be removed before changing a WAN or LAN mode of the GE interface. Once configured as a LAN side interface, it cannot be used for circuits as a WAN side interface.

Brocade 7850 Extension Switch has 16 1/10/25GE interfaces in which eight can be configured for LAN.

Brocade SX6 Extension Blade has 16 1/10GE interfaces in which eight can be configured for LAN.

Brocade 7810 Extension Switch has six GE interfaces available (6 x 1/10GE optical ports or 4 x 1/10GE optical ports and 2 x 1GE RJ-45 copper ports), and up to four interfaces can be configured for LAN. The number of available interfaces depends on whether the upgrade license has been installed.

Consider the following when you configure GE interfaces:

- Determine which GE interfaces will be used for the LAN side connectivity.
- Ensure that these GE interfaces are configured in LAN mode.
- Ensure that these GE interfaces are in the Default LS.

Configuring LAN side GE interfaces for IP Extension:

1. Connect to the switch and log on as admin.
2. Configure the necessary GE interfaces for LAN mode using the `portcfgge <ge#> --set -lan` command. Use the `portcfgge` command for more information.

```
switch:admin> portcfgge --help

Usage: portCfgGe [<slot>/][<port>] [<args>]

Port Format:
   ge#

Args:
   --enable -autoneg         - Enable the auto-negotiate mode of the 1 GE ports only.
   --disable -autoneg        - Disable the auto-negotiate mode of the 1 GE ports only.
   --set -speed <speed>      - Set the port speed of the GE ports.
                               Allowable speeds are [1G|10G|25G]
   --set -lan                - Set the port as lan.
   --set -wan                - Set the port as wan.
   --set -channel <channel>  - Set the port tunable SFP channel ID of the 10 GE ports only.
                               Allowable channel ID Range [1] - [102]
   --show                    - Show the current GE port configurations.
   --show -lmac              - Show the Local MAC addresses for GE/LAN ports.
```

1. Use the `portcfgge <ge#> --set -lan` command to configure a GE interface for the LAN operation. The following example puts port GE7 in LAN mode.

```
switch:admin> portcfgge ge7 --set -lan
Operation Succeeded.
```

2. Use the `portcfgge --show` command to verify that the GE interface is in LAN mode, show the speed, and to see if GE interfaces are part of a LAG.

```
switch:maintenance-service> portcfgge --show

Port  Speed  Flags  Channel  FEC  Lag Name
-------------------------------------------------------------------------------
ge0   1G     AC--   N/A      Off  -
ge1   1G     AC--   N/A      Off  -
ge2   10G    ----   N/A      Off  -
ge3   10G    ----   N/A      Off  -
ge4   1G     ----   N/A      Off  -
ge5   10G    ----   N/A      Off  -
ge6   10G    ----   N/A      Off  -
ge7   10G    -L     N/A      Off  -
-------------------------------------------------------------------------------
Flags: A:Auto-Negotiation Enabled C:Copper Media Type
L:LAN Port G: LAG Member
```

Alternatively, you can use the `switchshow` command to verify the GbE interface speed and mode.

```
switch:admin> switchshow
switchName:     switch
switchType:     178.0
switchState:    Online
switchMode:     Native
switchRole:     Principal
```

```
switchDomain:   20
switchId:       fffc14
switchWwn:      10:00:88:94:71:60:99:7e
zoning:         ON (EXT-Testing)
switchBeacon:   OFF
FC Router:      OFF
Fabric Name:    REPLICATION
HIF Mode:       OFF
Index Port Address  Media Speed   State       Proto
==================================================
   0   0   140000   id    N32     No_Light    FC
   1   1   140100   id    N32     No_Light    FC
   2   2   140200   id    N32     No_Light    FC
   3   3   140300   id    N32     No_Light    FC
   4   4   140400   --    N32     No_Module   FC
   5   5   140500   --    N32     No_Module   FC
   6   6   140600   --    N32     No_Module   FC
   7   7   140700   --    N32     No_Module   FC
   8   8   140800   --    N32     No_Module   FC
   9   9   140900   --    N32     No_Module   FC
  10  10   140a00   --    N32     No_Module   FC
  11  11   140b00   --    N32     No_Module   FC
  12  12   140c00   --    --      Offline     VE
  13  13   140d00   --    --      Offline     VE
  14  14   140e00   --    --      Offline     VE
  15  15   140f00   --    --      Offline     VE
      ge0           cu    1G      Offline     FCIP  Copper  Disabled (Unsupported blade mode)
      ge1           cu    1G      Offline     FCIP  Copper  Disabled (Unsupported blade mode)
      ge2           id    10G     No_Light    FCIP
      ge3           id    10G     No_Light    FCIP
      ge4           id    10G     No_Module   FCIP
      ge5           id    10G     No_Module   FCIP
      ge6           id    10G     No_Light    FCIP
      ge7           id    10G     No_Light    LAN
```

Portchannels (LAG)

IP Extension supports portchannels on the LAN side, also referred to as link aggregation groups (LAG). It depends on the platform whether a portchannel can have multiple links or how many portchannels it can contain. A LAG can be defined as static or dynamic. Dynamic LAGs use Link Aggregation Control Protocol (LACP).

Table 33: Maximum Portchannels per Platform

Extension Platform	Max. IP Extension Gateways per DP	Max. LAN Interfaces per Platform	Max. LAGs per Platform	Max. Links per LAG
Brocade 7810 Extension Switch	4	4	2	2
Brocade 7850 Extension Switch	8	8	4	4
Brocade SX6 Extension Blade	8	8	4	4

NOTE

An IP Extension LAN side LAG cannot span multiple Brocade SX6 Extension Blade platforms.

Consider the following when you configure Layer 2 for IP Extension LAN side ports:

- IP Extension LAN side interfaces support jumbo frames. A jumbo frame can be up to 9216 bytes
- IP Extension LAGs support VLAN tagging (802.1Q). Stacked tagging (QinQ) is not supported.

A portchannel is also referred to as a LAG, and applies only to the IP Extension LAN side. A LAG forms a single logical connection from multiple links. The GE interfaces must be set to LAN before a LAG can be configured. On the Brocade 7810 Extension Switch, up to four 1/10GE interfaces can be set to LAN. On the Brocade 7850 Extension Switch, up to eight 1/10/25GE interfaces can be set to LAN. On the Brocade SX6 Extension Blade, up to eight 1/10GE interfaces can be set to LAN. The Brocade Extension platforms support dynamic LAGs, which use LACP, and static LAGs. Multiple extension blades installed on a chassis cannot form a LAG across different blades.

Configuring a LAG is optional; however, without a LAG, you can only connect a single Ethernet link from the IP Extension platform to the data center LAN switch. A single link does not provide redundancy; a LAG provides redundant ports, cables, and optics. The recommended practice is to configure a LAG when connecting to a data center LAN switch or switches.

LAN side IP Extension links cannot connect to more than one data center LAN switch unless the switches are logically one switch. For example, VCC, vPC, and MLAG logically creates a single switch. When connecting to multiple data center LAN switches that are logically a single switch, the LAN side links must be in a portchannel.

A LAG can be configured with or without 802.1Q VLAN Tagging. In a Layer 2 type of deployment, a single LAG can accommodate multiple VLANs, if it has VLAN tagging enabled; otherwise, without VLAN tagging only the access VLAN is communicating across the LAG. End devices do not require a separate LAG for each VLAN.

All links in a LAG must be operating at the same speed. The Brocade 7850 Extension Switch and Brocade SX6 Extension Blade can support LAGs with up to four links. The Brocade 7850 Extension Switch can support a LAG with up to four links. The maximum number of LAGs is two, each having two links. The Brocade 7810 Extension Switch 1GbE RJ-45 copper ports can participate in forming a LAG; therefore, a LAG can have a mix of copper and optical 1GbE interfaces.

NOTE

- Ethernet must be configured for LAN before adding an interface to a LAG. LAG will not enhance the ability of a link to come online. Each link must be capable of coming online independently.
- There are no WAN side LAGs. The WAN side uses BET to achieve the same enhanced functionality with additional benefits.

A LAG is given a name when it is configured. GE interfaces are added or removed from a LAG. The interface speeds and auto-negotiation (1GbE only) settings must match for all interfaces that are added to the LAG. Individual LAG interfaces can enable or disable.

Perform the following steps to configure a LAG:

1. Connect to the switch and log on as admin.
2. Use the `lacp` command with the `--help` option to view the available options.

```
switch:admin> lacp --help
lacp cli usage:
lacp --help
lacp --show
lacp --config -sysprio <priority>
lacp --default
         Sets all the configurations to default
```

3. For dynamic LAGs, set the global LACP priority from 0 through 65535, the default is 32768. Use the `lacp --config -sysprio <priority>` command. Then run `lacp --show` to view the setting. This step is optional.
   The LACP system priority is set to 100.

```
switch:admin> lacp --config --sysprio 100 switch:admin> lacp --show
LACP system priority: 100
LACP System ID: 0x8000,00-05-33-74-85-42
```

4. Use the `portchannel` command with the `--help` option to view the available options.

```
switch:admin> portchannel --help
portchannel cli usage:
portchannel --help
portchannel --create <po_name> -type <static|dynamic> [-key <po-number>]
portchannel --delete <po_name>
portchannel --enable <po_name>
portchannel --disable <po_name>
portchannel --add <po_name> -port <port-num|port-range> [-timeout <l/s> Dynamic Port-Channel]
portchannel --remove <po_name> -port <port-num|port-range>
portchannel --show -detail | -static | -dynamic | -all | <po_name> | -stats [<po_name>]
portchannel --config  <po_name> -autoneg <on|off> [Not applicable on Empty port-channel]
portchannel --config  <po_name> -type <static|dynamic>
portchannel --config  <po_name> -rename <new-po-name>
portchannel --config  -port <port-num|port-range> -timeout <s|l> [Dynamic port-channel]
portchannel --reset -stats [<po_name>]

Specifying port ranges:
        port_num -- <port_num> is <[slot]/port>
        port_range -- Specifies a set of ports as a range:
                 (examples: '10/6' or '10/6-9' or '10/ge6-ge9' or '10/ge6-9')
Actions:
        --show  : Show the configured portchannel details

                -static : Show all configured static portchannels.
                -dynamic : Show all configured dynamic portchannels.
                -all : Show all configured portchannels.
                -detail : Show portchannels in details.
                -stats : Show all statistics on FCIP portchannels.
        --create  : Creates portchannel. Portchannel name is a string of max 31 character
                characters allowed:alphanumeric with special char _-

                -type : Specify the type of portchannel.
                -key : Specify the Key of portchannel. key range is <1-1000>
        --delete  : Delete the specified portchannel

        --config  : Configures portchannel and member port specific parameters
                -type : Change type of the portchannel from dynamic to static and vice versa.
                -rename : Change name of the portchannel.
                -autoneg : Configure Auto-neg on or off for portchannel of speed 1G.
                -port <port_num | port-range> : The port number or range of port
                -timeout : Configures member port timeout option. possible options are S and L
        --enable : Enable the specified portchannel
        --disable : Disable the specified portchannel
        --add : Add member port to portchannel
                -port <port_num | port-range> : The port number or range of port to be added
        --remove : Remove member port from portchannel
                -port <port_num | port-range> : The port number or range of port to be removed
        --reset  : reset portchannel parameters
```

```
                -stats : reset sttistics for FCIP portchannels
```

5. Create either a static or dynamic LAG with the `portchannel --create` command.

```
switch:admin> portchannel --create MyDynLag -type dynamic -key 555
switch:admin> portchannel --create MyStaticLag -type static -key 100

switch:admin> portchannel --show

Name         Type     Oper-State  Port-Count Member Ports
--------------------------------------------------------
MyDynLag     Dynamic  Offline     0
MyStaticLag  Static   Offline     0
```

GE interfaces are not members of a portchannel after configuration. Add ports individually or by specifying a range of ports.

6. Add GE ports to the groups with the `portchannel --add` command.

```
switch:admin> portchannel --add MyStaticLag -port ge16-17

WARNING: While making configuration changes the modified LAN GE ports will be disabled. Please manually
 enable the modified LAN GE ports after completing all the configuration changes.

switch:admin> portchannel --show -all

Name          Type      Oper-State  Port-Count   Member Ports
-------------------------------------------------------------
MyDynLag      Dynamic   Online      1            ge6
MyStaticLag   Static    Offline     3            ge15 ,ge16 ,ge17
```

In static LAGs, LAN side GE interfaces are disabled when added to the LAG and must be re-enabled after completing the LAG configuration. When adding an interface to a disabled dynamic LAG, the interface is disabled. When you remove an interface from a disabled dynamic LAG, the interface is enabled. If a disabled dynamic LAG is deleted using the `portchannel --delete` command, all the member interfaces are enabled.

Set auto-negotiation for all the 1GbE interfaces in the LAG using the `portchannel --config` command.

```
switch:admin> portchannel --config slag101 -autoneg on
```

NOTE

GE interfaces are disabled after changing the auto-negotiation configuration. Fabric OS automatically re-enables the ports.

The following RASLog example shows that Fabric OS has re-enabled the ports:

```
2023/02/16-22:37:18 (GMT), [ESM-2801], 94, FID 128 | PORT 0/GE7, INFO, Awing-2, Port ge7 speed set to
1G, auto negotiation Enabled, FEC clause Off [API].
2023/02/16-22:37:19 (GMT), [PORT-1008], 95, FID 128, INFO, Awing-2, GigE Port (0/GE7) has been enabled.
```

7. Enable or disable a LAG with the `portchannel --enable | --disable` command.

NOTE

By default, admin state of a portchannel is enabled.

```
switch:admin> portchannel --enable MyStaticLag

switch:admin> portchannel --show -all

Name          Type     Oper-State  Port-Count   Member Ports
-------------------------------------------------------------
MyDynLag      Dynamic  Online      1            ge6
MyStaticLag   Static   Online      3            ge15 ,ge16 ,ge17
```

8. View detailed information about the LAG groups with the `portchannel --show -detail` command.

```
switch:admin> portchannel --show -detail
```

```
Name :test
Type :Dynamic Key :1
Speed :1G
Admin-state: Enable Oper-state : Online
Admin Key: 0001 - Oper Key 0001
LACP System ID: 0x8000,c4-f5-7c-01-31-4a PART System ID: 0x0001,00-24-38-9b-03-00
Portchannel Member count = 2
    Port    Oper state         Sync    Timeout         Auto-Negotiation
    ------------------------------------------------------------------------------
    *ge5    Online             1       Long            Disabled
     ge6    Offline            0       Long            Disabled

Name :static
Type :Static Key :2
Speed :1G
Admin-state: Enable Oper-state : Offline
Portchannel Member count = 2

    Port    Oper state    Auto-Negotiation
   -------------------------------------------------
    ge0     offline       Enabled
    ge1     Offline       Enabled
```

9. Reset the statistics with the `portchannel --reset` command.

```
switch:admin> portchannel --reset portchannel --reset -stats [<po_name>]

switch:admin> portchannel --reset -stats switch:admin> portchannel --show -stats

Name:    static_lag1  LAG Name
Speed:           10G  LAG Speed
Admin-State:  Enable  LAG Admin State
Oper-State:   Online  LAG Operational  Status
InPkts:             0  LAG Received Frames
InOctets:           0  LAG Received Octets
InUcastPkts:        0  LAG Received Unicast Frames
InMcastPkts:        0  LAG Received Multicast Frames
InBcastPkts:        0  LAG Received Broadcast Frames
InVlanPkts:         0  LAG Received VLAN Frames
InPausePkts:        0  LAG Received Pause Frames
InDiscards:         0  LAG Received Frames Discarded
InErrors:           0  LAG Received Frames with Errors
OutPkts:            0  LAG Transmitted Frames
OutOctets:          0  LAG Transmitted Octets
OutUcastPkts:       0  LAG Transmitted Unicast Frames
OutMcastPkts:       0  LAG Transmitted Multicast Frames
OutBcastPkts:       0  LAG Transmitted Broadcast Frames
OutVlanPkts:        0  LAG Transmitted VLAN Frames
OutPausePkts:       0  LAG Transmitted Pause Frames
OutDiscards:        0  LAG Transmitted Frames Discarded
OutErrors:          0  LAG Transmitted Frames with Errors
CRCErrors:          0  LAG CRC Errors
CarrierErrors:      0  LAG lost carrier sense
```

```
JabberErrors:       0   LAG Jabbers
LAG Uptime:   3 Days 8 Hours 43 Mins 27 Seconds
```

IP Extension Deployment Types

IP Extension connectivity is supported in direct connections to the Brocade Extension platform, layer 2 (switched), and layer 3 (routed) networks.

Layer 2 Deployment

Layer 2 communicates at the Ethernet level. End devices on the same VLAN with IP addresses on the same subnet as the IP Extension Gateway can communicate at the Ethernet level. Alternatively, IP storage can be connected directly to the IP Extension LAN ports; however, data center LAN switches can scale the number of end device Ethernet connections that are needed for IP Extension. By utilizing a static route on the end device, the end device is capable of differentiating the gateway to which it sends traffic based on the destination IP address or subnet. The following figure shows an IP Extension deployment using direct connections, which is a Layer 2 method.

Figure 15: IP Extension Direct Connect to IP Storage

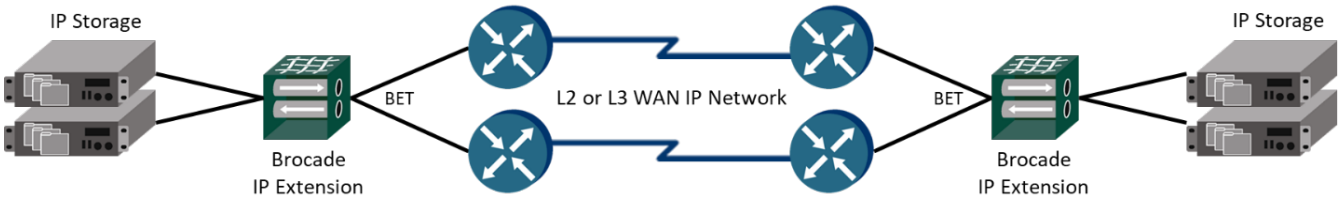

In this case, configure the IP storage with the IP Extension Gateway as the next-hop. The IP Extension Gateway is configured on the DP (DP0 or DP1) owning the VE_Port of the tunnel. The IP Extension Gateways are IPIF LAN IP addresses specific to a DP. The same IP Extension Gateway cannot be configured on more than one DP on an Extension platform. Also, the same IP Extension Gateway cannot be configured on more than one DP across platforms, which cause an IP address conflict on the data center network. Based on the next-hop of the end device, the end device learns the IP Extension Gateway's MAC address through an ARP or Neighbor Discovery (ND) request.

NOTE
The Brocade 7810 Extension Switch has only one DP, which is DP0.

When IP storage devices are connected to a Layer 2 switch, as shown in the following figure, you can use Link Aggregation 802.3ad (LAG) to connect IP Extension LAN interfaces to the switch. The Brocade Extension platforms support static and dynamic LAG to data center switches. Dynamic LAG uses LACP. The LAG can be 802.1Q tagged if various end devices live on more than one VLAN. Tagging enables multiple VLANs to communicate across the same physical LAG. The default is no tagging.

NOTE

- Only one path might exist between the Layer 2 broadcast domain and the Brocade IP Extension LAN interfaces. A LAG is considered a single path.
- NIC teaming is not supported.

Figure 16: P Extension Connectivity using a L2 Switch

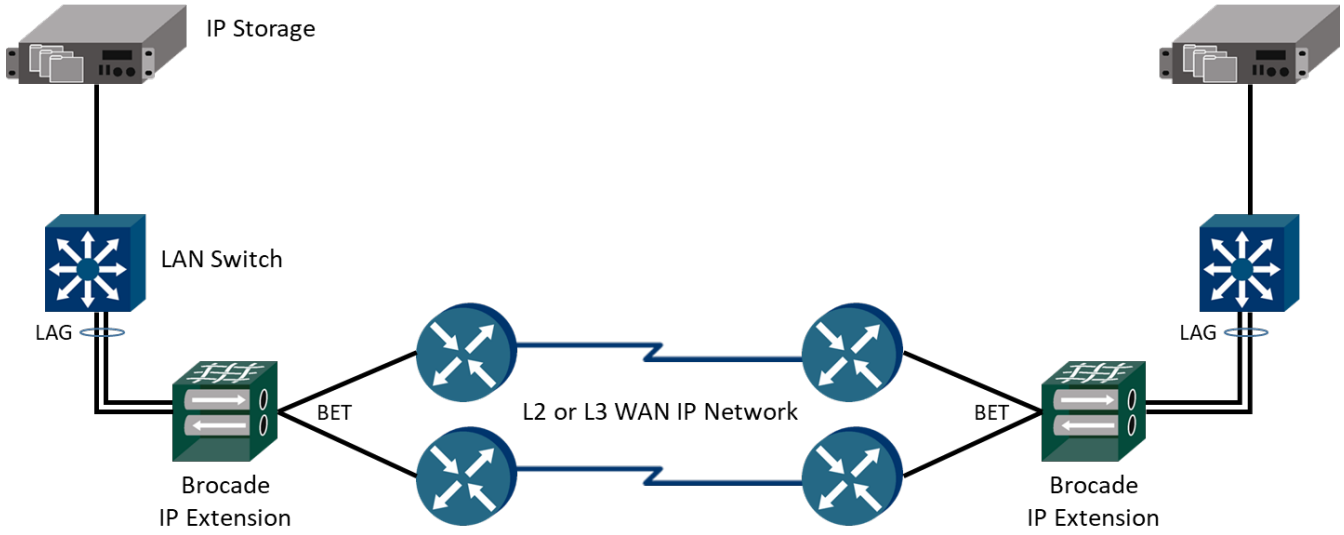

As shown in the figure, the Brocade IP Extension LAN interfaces are connected to the IP storage by the Layer 2 LAN switch. You must configure at least one IP Extension Gateway, which is accomplished by configuring an IPIF LAN IP address. For a Layer 2 architecture, each VLAN on each DP owning an IP Extension enabled tunnel requires a LAN side IPIF. The LAN side IPIF has an IP address on the same subnet as the IP addresses of end devices and acts as the IP Extension Gateway.

Layer 3 Deployment

End devices do not require special routes and send traffic normally instead of sending a specific traffic to the IP Extension Gateway. The changes are in the IP network routers, which are configured to intercept the traffic and forward it to the IP Extension Gateway. A policy-based routing (PBR) mechanism is the most common method of performing this task within an IP network.

The following figure shows IP Extension that is used in a routed network configured with PBR. The network intercepts traffic that is matched by an access control list (ACL) and redirects the traffic to an IP Extension Gateway. A Layer 3 deployment simplifies IP storage connectivity and can scale to a greater number of subnets than a Layer 2 deployment. IP storage ports might be located in different subnets from the IP Extension Gateway. The IP Extension Gateway only requires the routed network of the data center as its next-hop, and relies on the routed network to deliver traffic to the end devices.

Figure 17: IP Extension Layer 3 Connectivity

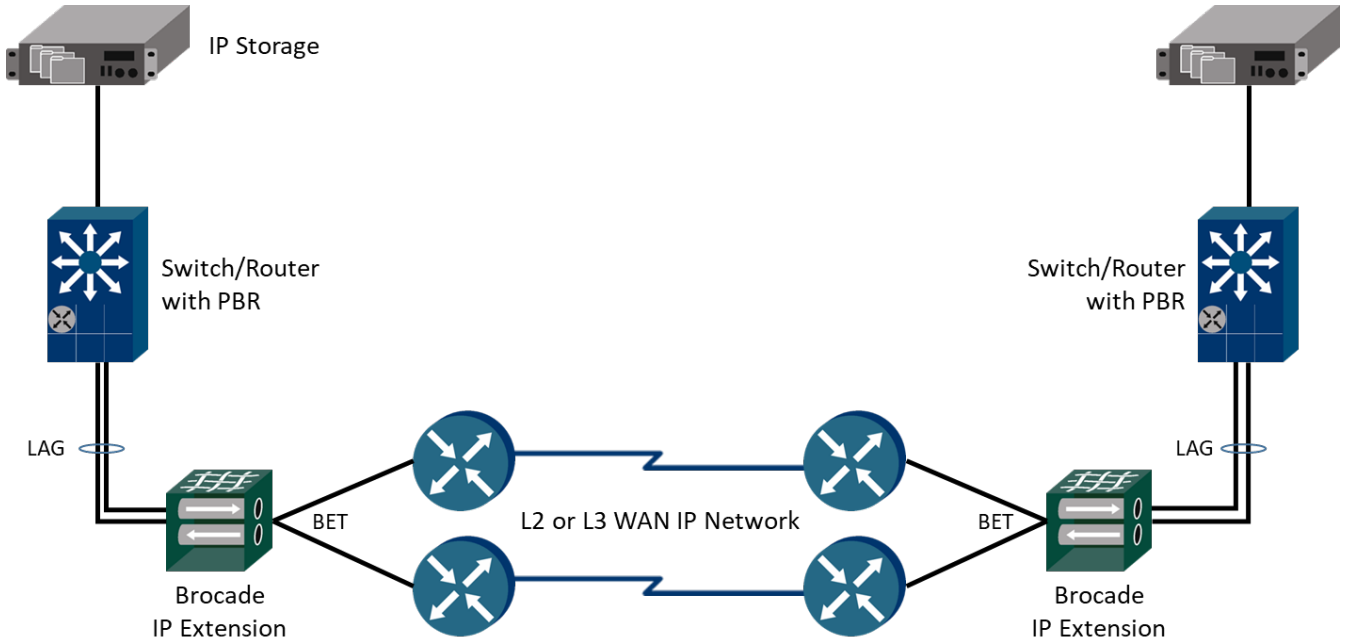

IP Extension Limitations and Considerations

For the TCP traffic, the following considerations apply:

- The LAN side uses the RX window advertisement for flow control data from the source for flow control preventing buffer overflows. However, the window closes when the WAN experiences congestion or the destination end device exhibits slow drain behavior.
- Up to 512 TCP open requests per second are allowed, additional open requests are dropped. Limiting open requests per second mitigates Denial of Service (DoS) attacks.
- Statistics are available for the number of dropped connections.

Table 34: Maximum Supported TCP Session and RASLog Warning Threshold

Brocade Extension Platform	TCP sessions per DP	RASLog Warning Threshold
Brocade 7810 Extension Switch	256 (1 DP)	244
Brocade 7850 Extension Switch	512 (2 DPs)	500
Brocade SX6 Extension Blade	512 (2 DPs)	500

IP Extension Gateway (LAN IPIF)

The LAN side IPIF acts as a gateway for IP storage devices replicating data across IP Extension. Each DP has its own set of IP Extension Gateways, one for each subnet in which IP storage devices are located. On the LAN side, both IPv4 and IPv6 are supported with IP Extension.

An IP Extension Gateway address cannot be used on both DPs of a Brocade 7850 Extension Switch or Brocade SX6 Extension Blade, or anywhere else in the network. An IP Extension Gateway can be located on the same subnet, however it must be unique.

LAN side GE interfaces can access all IP Extension Gateways; therefore, all untagged (no 802.1Q) Ethernet traffic can access IP Extension Gateways on either DP. If the LAN side traffic is tagged (802.1Q), then access depends on matching the incoming VLAN tag with the VLAN ID of the LAN side IPIF of the IP Extension Gateway.

The IP Extension LAN side IPIF is a gateway for IP storage devices to access IP Extension, referred to as IP Extension Gateway. End-devices send the IP storage traffic to the IP Extension Gateways. The IP Extension Gateways are used for both layer 2 and layer 3 deployments.

In a Layer 2 implementation, the IP Extension Gateway is the next-hop address for the storage end device. A layer 2 implementation requires IP routes be put on the end device to direct its traffic to the IP Extension Gateway. Changes have to be made to the end devices. On each end-device, for each subnet that is used on an interface, an IP route must be added to the end-device which points to the IP Extension Gateway. For each subnet on which an end device has an interface, an IP Extension Gateway must be configured. These end-device IP routes are required before the end device sends traffic to the IP Extension Gateway.

The default route to the data center router is the next-hop address; therefore, no changes must be made to the end devices in a layer 3 implementation. However, it is necessary to modify the data center routers so that IP storage flows are intercepted and forwarded to the IP Extension Gateway.

Locally, device ports can be on the same subnet or different subnets; however, the local and remote subnets cannot be the same, they must be different.

The following considerations apply to the IP Extension IPIF:

- An IPIF lan.dp# is referred to as an IP Extension Gateway.
- The same IP Extension Gateway IP address cannot be used on another DP.
- Identify the end-device replication ports that are extended through IP Extension.
- Identify the subnets that the IP storage end-devices use for replication.
  - With a Layer 2 deployment, create IP Extension Gateways (LAN IPIF) using IP addresses from each subnet.
  - Identify which VLANs are used.
  - With a Layer 2 deployment, create IP Extension Gateways (LAN IPIF) using IP addresses from each subnet.
- The LAN side IP MTU range is 1280 bytes to 9216 bytes.
- PMTU is not supported on LAN side IP Extension Gateway interfaces (LAN IPIF).
- IPIF IP addresses and subnet masks cannot be modified, the IPIF has to be deleted and recreated.
- Multiple IP Extension Gateways on the same subnet cannot be configured on the same DP.
- A separate LAN side IPIF (IP Extension Gateway) is needed for each VLAN ID on an 802.1Q tagged LAG.

  There is only one Software Virtual Interface (SVI) per DP, which means there is only one LAN side MAC per DP. A Brocade 7850 Extension Switch and Brocade SX6 Extension Blade can have a maximum of eight IP Extension Gateways per DP, and the Brocade 7810 Extension Switch can have a maximum of four IP Extension Gateways.

  In the following situations, distinct LAN gateways must be configured:

  If the LAG of the data center LAN switch is not configured to use VLAN tagging (802.1Q), do not associate a VLAN ID with the LAN side IPIF. Ethernet links cannot be made operational when tagging and non-tagging are not matching. Tagged traffic can only talk to interfaces that are configured for the tagged traffic.

  To accommodate multiple VLANs, the LAG must be a trunk, which means the traffic is tagged. In a trunk, multiple logical ISLs pass through the physical LAG and traffic of each VLAN is tagged with the corresponding VLAN ID. For the tagged traffic to be forwarded to the correct IP Extension Gateway, the LAN IPIF must be configured with the VLAN ID. If the LAG to the LAN switch is tagged, an IP Extension Gateway is created for each VLAN.

  Perform the following steps to configure an IP Extension Gateway:

  a. Connect to the switch and log on as admin.
  b. Use the `portcfg ipif` command to create a LAN side IP Extension Gateway. The VLAN and MTU arguments are optional. The VLAN default is none, which means there is no tagging. The MTU default is 1500 bytes. The following example shows how to enter the IP address and subnet using CIDR IP address notation (/subnet mask length) or by using the netmask argument.

The IP Extension Gateway is local to the platform, and the gateway is registered with DP0, which is important when associating the gateway with a tunnel (VE_Port) - both must be on the same DP.

Using the CIDR notation (that is, /24).

```
switch:admin> portcfg ipif lan.dp0 create 10.0.0.4/24
```

Using the netmask argument.

```
switch:admin> portcfg ipif lan.dp0 create 10.0.0.4 netmask 255.255.255.0
```

c. Delete an IPIF.

```
switch:admin> portcfg ipif lan.dp0 delete 10.0.0.4/24
```

While IP Extension Gateway IPIFs are in use, they can be deleted. Clean-up enforcement checks are not done on LAN side IPIFs.

NOTE
If a LAN side IPIF is accidentally deleted, all IP Extension flows that use that IP Extension Gateway will be disrupted.

d. Use the `portshow ipif` command to display the IPIF ports.

The following example shows two LAN side IP Extension Gateways that are configured on slot 7 DP0, and two WAN side circuit IPIFs configured on slot 7 DP0:

- 10.0.0.1 on VLAN 100 with MTU 1500
- 10.0.1.1 on VLAN 200 with MTU 9216
- 192.168.60.20 on GE4
- 192.168.10.107 on GE17

```
switch:admin> portshow ipif

Port        IP Address      / Pfx  MTU    VLAN  Flags
-----------------------------------------------------
7/ge4.dp0   192.168.60.20   / 24   1500   0     U R M
7/ge17.dp0  192.168.10.107  / 24   1500   0     U R M
7/lan.dp0   10.0.0.1        / 24   1500   100   U R M
7/lan.dp0   10.0.1.1        / 24   9216   200   U R M
-----------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

LAN side VLANs

LAN side VLAN IDs are configured on LAN side IPIFs. It is possible to configure the same VLAN ID on more than one IPIF, for example, on an IPIF for DP0 and an IPIF for DP1. IPIFs for DP0 and DP1 cannot share the same gateway address, but they can belong to the same VLAN.

A VLAN ID is needed only when the traffic from the data center network switches is tagged. The Ethernet port directly connected to a Brocade LAN side GE must be configured for VLAN tagging; otherwise, do not configure a VLAN ID on the IPIF. Usually, the switch port is part of a particular VLAN, but is not configured to perform VLAN tagging, so a VLAN ID on the IPIF is not necessary. The Ethernet link cannot be established if there is a mismatch between the Brocade Extension platform and the data center switch.

When VLANs are used, they are used to differentiate between traffic flows passing through the same physical GE interface. The data center switch can forward each flow based on its VLAN ID.

Perform the following steps to add VLAN ID when creating an IPIF:

1. Connect to the switch and log on as admin.
2. Add a VLAN ID when creating an IPIF.

```
switch:admin> portcfg ipif lan.dp0 create 10.0.0.4/24 vlan 100
```

3. Change a VLAN ID on an IPIF.
```
switch:admin> portcfg ipif lan.dp0 modify 10.0.0.4/24 vlan 101
```
4. Remove a VLAN ID from an IPIF.
```
switch:admin> portcfg ipif lan.dp0 create 10.0.0.4/24 vlan none
```
5. Enter the `portshow ipif` command to verify settings.
```
switch:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.10.1                  / 30   9216  0     U R M I
 ge2.dp0      10.0.0.5                      / 24   AUTO  0     U R M
 lan.dp0      10.10.10.123                  / 24   1500  100   U R M
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

LAN side MTUs

Determine the supported MTU from both the IP Storage provider and the Network Administrator. Set the MTU value on the LAN side IPIF. The default MTU is 1500 bytes. The larger the MTU, the less overhead and the more efficient the data transfer is.

1. Connect to the switch and log on as admin.
2. The following example shows an untagged (no VLAN ID) IP Extension Gateway on slot 7 DP0 configured with an MTU of 9216:
```
switch:admin> portcfg ipif 7/lan.dp0 create 10.0.1.5/24 mtu 9216
```
3. The following example shows modifying an existing LAN IPIF on DP1 to set the MTU to 1500 bytes:
```
switch:admin> portcfg ipif lan.dp1 modify 10.0.1.6/24 mtu 1500
```
4. Enter the `portshow ipif` command to verify settings.
```
switch:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.10.1                  / 30   9216  0     U R M I
 ge2.dp0      10.0.0.5                      / 24   AUTO  0     U R M
 lan.dp0      10.10.10.123                  / 24   1500  100   U R M
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

LAN side IP Routes

IP Extension supports Layer 2 and Layer 3 deployments. LAN side IP routes are only used with layer 3 deployments. An LAN side IP route in a layer 3 deployment is used to forward arriving IP Extension traffic from the tunnel to the next hop router. With layer 2 deployments, the end-devices are on the same subnet and there is no intermediate router. LAN side IP routes play no role in the forwarding traffic on or toward the WAN side. Layer 3 deployments employ IP Extension to deliver traffic that has already passed through the tunnel to a remote router for delivery to the end device's subnet. IP routes define a destination subnet where the end devices reside, and the router gateway address that can forward the data to that subnet.

In a Layer 3 deployment, when traffic arrives at the local router from an IP storage end-device, it is intercepted by the router and forwarded to the IP Extension Gateway. Typically, access control list (ACL) and policy-based routing (PBR) are used to match the incoming traffic and forward it to the IP Extension Gateway. The router must be configured to perform this function. End-devices do not require any special configuration with Layer 3.

Configuration steps for PBR are done on the router. The router is configured with a next-hop gateway address that identifies the extension platform's IP Extension Gateway. An ACL and PBR policy determine exactly what traffic is sent to the IP Extension Gateway.

On the extension platform, a LAN side IP route to the next-hop gateway is configured. The IP Extension Gateway and the next-hop gateway are on the same subnet. Steps for configuring a LAN side IP route are similar to configuring a WAN side IP route. Traffic from the end-device going towards the tunnel is forwarded by PBR, and traffic from the tunnel towards the end-device is forwarded by the IP Extension LAN side IP route.

The following illustration shows an example network topology where Brocade IP Extension is the next-hop gateway for traffic forwarded by PBR from the DC router. Traffic not diverted by PBR to Brocade IP Extension is routed based on the traditional routing table.

Figure 18: Layer 3 Deployment

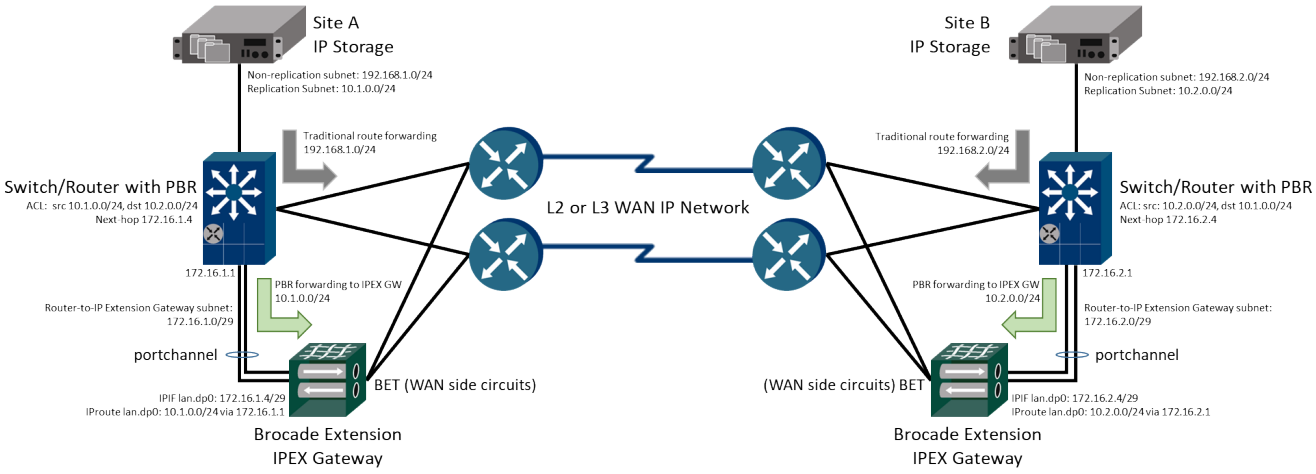

PBR is implemented in routers at each end. In the configuration steps, the router configurations and the IP storage configurations are generic. For the detailed steps, refer to the documentation of your network equipment.

Configuring LAN Side IP Routes for Layer 3 IP Extension Deployment

1. Connect to the switch and log on as admin.
2. Use the `portcfg ipif lan.dp#` command to configure a LAN side IPIF on each end. The `portcfg ipif` command creates a LAN side IPIF (IP Extension Gateway) on DP0. The following example shows the netmask keyword; however, the CIDR nomenclature without the keyword will also work (/29):

   ```
   admin:switch_siteA> portcfg ipif lan.dp0 create 172.16.1.4 netmask 255.255.255.248
   admin:switch_siteB> portcfg ipif lan.dp0 create 172.16.2.4 netmask 255.255.255.248
   ```

3. Use the `portcfg iproute lan.dp#` command to specify the next-hop gateway address. Perform this action on each end.

   ```
   admin:switch_siteA> portcfg iproute lan.dp0 create 10.1.0.0/24 172.16.1.1
   admin:switch_siteB> portcfg iproute lan.dp0 create 10.2.0.0/24 172.16.2.1
   ```

   These commands configure a LAN side IP route with the destination subnet of the local storage end device. The IP Extension Gateway and DC router's gateway are connected using the LAG, both gateways are on the same subnet.
4. On the routers in DC A and B, configure the portchannel (LAG) and router interface that is connected to Brocade Extension. Usually, there are two routers and each has a gateway interface. Virtual IP addresses (VIPs) are created

between them using protocols such as VRRP and HSRP. Using VRRP or HSRP protocols allows one router to be active and the other for failover. For example, router-1 has 172.16.1.2, router-2 has 172.16.1.3, and the shared VIP is 172.16.1.1; therefore, for this subnet, the Brocade IP Extension GW is 172.16.1.4 as shown below.

```
DC Router Site A:
IP: 10.1.0.1 mask 255.255.255.0
DC Router Site B:
IP: 10.2.0.1 mask 255.255.255.0
```

NOTE

For detailed instructions, refer to the documentation provided with your router equipment.

The following steps provide the generic information:

5. On the routers in DC A and B, configure an ACL and PBR to direct the specific IP traffic destined to the remote end devices. Designate the IP Extension Gateway as the next hop gateway.

```
DC Router Site A:
Destination subnet: 10.2.0.0/24
Source subnet: 10.1.0.0/24
Next-Hop Gateway: 172.16.1.4
DC Router Site B:
Destination subnet: 10.1.0.0/24
Source subnet: 10.2.0.0/24
Next-Hop Gateway: 172.16.2.4
```

6. On the IP storage end-devices in DC A and B, configure the replication ports with IP addresses from the subnet that has been enabled by the network for an IP extension. All IP addresses are examples.

```
Site A: IP Storage Replication Port:
IP: 10.1.0.10 mask 255.255.255.0
Site A: IP Storage Non-Replication Port:
IP: 192.168.1.10 mask 255.255.255.0
Site B: IP Storage Replication Port:
IP: 10.2.0.10 mask 255.255.255.0
Site B: IP Storage Non-Replication Port:
IP: 192.168.2.10 mask 255.255.255.0
```

7. Typically, this step is not required since the default route of the IP storage end device already sends traffic to the gateway of the DC router.

```
Site A: IP Storage Default Route
Subnet: 10.1.0.0
Mask: 255.255.255.0
Gateway 10.1.0.1
Site B: IP Storage Default Route
Subnet: 10.2.0.0
Mask: 255.255.255.0
Gateway 10.2.0.1
```

Verifying LAN Side IP Connectivity

Brocade Extension platforms can ping from the IP Extension Gateway out to devices in the data center LAN. It is not possible to ping from the LAN side to any destination in the direction of the WAN side. The IP Extension Gateway is a demarcation point between the LAN side and WAN side and pinging across the demarcation point is not supported. A tunnel does not have to be configured to ping from an IPIF. A LAN side IP route does not need to be configured for a layer 2 deployment because pinging is on the same subnet. In a layer 3 deployment, the router gateway on the same subnet as the IPIF can be pinged without an IP route; however, the end-device ports beyond the router gateway are on different subnets and can require a LAN side IP route to be configured.

1. Connect to the switch and log on as admin.

2. Enter the `portcmd --ping` command.

```
switch:admin> portcmd --ping
--ping: Usage

--ping [<slot>/]<port> <options>

Options:
  -s,--src <ipaddr>              - Specify source IP Address to send from.
  -d,--dst <ipaddr>              - Specify destination IP Address to send to.
  -n,--num-pings <num>           - Specify the number of pings to send.
  -t,--ttl <ttl>                 - Set the TTL value to use.
  -w,--wait-time <sec>           - Set the timeout value in seconds for each ping
                                   request.
  -z,--size <bytes>              - Set the payload size of each ping request in
                                   bytes.
  -i,--ip-tos <0x0-0xFF|0-255>   - Set the IP TOS value to use; accepts hex or
                                   decimal.
  --dscp <0x0-0x3F|0-63>         - Set the dscp value to use; accepts hex or
                                   decimal.
  -c,--l2cos <0-7>               - Set the L2CoS value to use.
  -y,--quiet                     - Quiet output. Only show summary lines at start
                                   and finish.
  -h,--help                      - Print this usage statement.
```

3. Use the `portcmd --ping` command to ping a known operational end-device port or router gateway on the LAN side. The source IP address is the LAN side IPIF that you want to test. The destination IP address should be on the same subnet. In the example below, the following IP addresses are used, or can be used:

   - Pining from the Brocade Extension platform out onto the data center LAN.
   - Source is the IP Extension Gateway, which is the LAN side IPIF: 10.0.0.4
   - Destination is a gateway on the router for this subnet: 10.0.0.1 (shown below)
   - Alternatively, the IP address of a replication port on the end-device can be pinged: 10.0.0.123 (not shown)

```
switch:admin> portcmd --ping lan.dp0 -s 10.0.0.4 -d 10.0.0.1

PING 10.0.0.1 (10.0.0.4) with 64 bytes of data.
64 bytes from 10.0.0.1: icmp_seq=1 ttl=255 time=1 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=255 time=1 ms
64 bytes from 10.0.0.1: icmp_seq=3 ttl=255 time=1 ms
64 bytes from 10.0.0.1: icmp_seq=4 ttl=255 time=1 ms

--- 10.0.0.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 691 ms
rtt min/avg/max = 1/1/1 ms
```

Enabling IP Extension on a Tunnel

A WAN side tunnel requires IP Extension to be enabled; it is not enabled by default. This tunnel attribute should not be confused with hybrid mode. A tunnel enabled for IP Extension has a protocol bandwidth distribution setting and IP Extension specific QoS priority percentages. Tunnels can carry both FCIP and IP Extension traffic, which is recommended because the tunnel manages the flow control for both protocols.

Staging is a configuration method for tunnels and circuits. When staging tunnels and circuits, create a basic tunnel using the `portcfg fciptunnel <ve> modify` command. Add features in stages. When the tunnel has been configured,

add the circuits using the `portcfg fcipcircuit create` command. Finally, add features to the circuits using the `portcfg fcipcircuit modify` command.

Create a tunnel or modify an existing tunnel with IP Extension enabled by using the `portCfg fciptunnel <ve> [create|modify] --ipext enable` command.

> **NOTE**
> When not using staging, circuit options can be specified on the `portcfg fciptunnel <ve> create` command and the `portcfg fciptunnel <ve> modify` command; however, the options apply only to circuit-0 because they are tunnel operands and not `fcipcircuit` operands. When additional circuits are added, circuit options must be applied per circuit using the `portcfg fcipcircuit create or portcfg fcipcircuit modify` commands.

1. Connect to the switch and log on as admin.
2. Use the `portcfg fciptunnel <ve> create` command to create an IP Extension capable tunnel.
   This example creates and enables an IP Extension tunnel on VE12.

```
switch:admin> portcfg fciptunnel 24 create --ipext enable

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fciptunnel specified for a brief period of time. This
 operation will bring the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n]
```

3. Use the `portcfg fciptunnel <ve> modify` command to modify an existing tunnel to disable IP Extension.

This example modifies an existing tunnel to disable IP Extension on VE12.

```
switch:admin> portcfg fciptunnel 12 modify --ipext disable

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fciptunnel specified for a brief period of time. This
 operation will bring the existing tunnel down (if tunnel is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [ n]
```

Tunnel Compression for IP Extension

The compression algorithm can be set generally for a tunnel, or if IP Extension is enabled, the algorithm can be enabled and disabled per protocol (FCIP and IP Extension). Protocol-specific compression settings override the general compression algorithm.

> **NOTE**
> IP Extension and Brocade 7810 Extension Switch do not support the fast-deflate compression.

Only deflate and aggressive-deflate are allowed with the `--ip-compression` option. If a general tunnel compression is set to fast-deflate, the IP Extension compression is set to none. The protocol-based compression modes can be set to default, which causes the compression setting of the protocol to be inherited from the general setting of a tunnel.

Perform the following steps to set compression specific to IP Extension:

1. Connect to the switch and log on as admin.
2. Use the `portcfg fciptunnel modify --ip-compression` command to configure the compression for IP Extension.

```
switch:admin> portcfg fciptunnel 12 modify --ip-compression deflate
```

3. Verify that deflate compression is enabled for IP Extension.

```
switch:admin> portshow fciptunnel 12
```

```
Tunnel: VE-Port:12 (idx:0, DP0)
 ==================================================
  Oper State           : Online
  TID                  : 12
  Flags                : 0x00000000
  IP-Extension         : Enabled
  Compression          : None
  FC-Compression       : None (Inherited)
  IP-Compression       : Deflate (Override)
...
```

4. With `portcfg fciptunnel modify` , you can restore compression values to their default values, which mean the setting is inherited from the general configuration.

```
switch:admin> portcfg fciptunnel 12 modify \
--fc-compression default \
--ip-compression default
```

IP Extension QoS

IP Extension has three QoS priorities; ip-high, ip-medium, and ip-low. Each protocol can configure a custom QoS distribution; the default distribution is 50%, 30%, 20% for high, medium, and low respectively. Minimum allocation for a single QoS type (high, medium, or low) is 10%. QoS allocations must total 100%.

FCIP and IP Extension protocol distribution percentages and circuit QoS percentages are rate-limiting to bandwidth. When any bandwidth is utilized in a distribution or QoS path, the entire percentage becomes reserved and cannot be shared with another distribution or QoS path.

IP Extension QoS Marking

IP Extension traffic can be marked with DSCP and L2 CoS (802.1P). Work with your Networking Administrators to determine if QoS marking can be used in the network.

Perform the following steps to mark the IP Extension traffic:

1. Connect to the switch and log on as admin.
2. Enter the `portcfg fcipcircuit --help` command.

```
switch:admin> portcfg fcipcircuit --help

Usage:   portCfg fcipcircuit [<slot>/]<port> <option> <cid> [<args>]

Option:    create - Create the specified tunnel/circuit
           modify - Modify the specified tunnel/circuit
           delete - Delete the specified tunnel/circuit

Optional Arguments:
  -a,--admin-status <enable|disable> -
                                -  Set the admin-status of the circuit.
  -S,--local-ip <ipaddr>|none  -  Set local IP address.
  -D,--remote-ip <ipaddr>|none -  Set remote IP address.
     --local-ha-ip <ipaddr>|none -
                                -  Set local HA IP address. This allows for HCL
                                   operations on local switch. [7840 / SX6 only]
     --remote-ha-ip <ipaddr>|none -
                                -  Set remote HA IP address. This allows for HCL
```

```
                                operations on remote switch. [7810 / 7840 /
                                SX6 only]
   -x,--metric <0|1>          -  Set the metric. 0=Primary 1=Failover.
   -g,--failover-group <0-9>  -  Set the failover group ID.
   -b,--min-comm-rate <kbps>  -  Set the minimum guaranteed rate.
   -B,--max-comm-rate <kbps>  -  Set the maximum rate.
      --arl-algorithm <mode>  -  Set the ARL algorithm. Allowable modes are
                                [auto] [reset] [step-down] [timed-step-down].
                                [7810 / 7840 / SX6 only]
   -C,--connection-type <type> -  Set the connection type. Allowable types are
                                [default] [listener] [initiator].
   -k,--keepalive-timeout <ms> -  Set keepalive timeout in ms.
      --dscp-f-class <dscp>   -  Set DSCP marking for Control traffic.
      --dscp-high <dscp>      -  Set DSCP marking for FC-High priority traffic.
      --dscp-medium <dscp>    -  Set DSCP marking for FC-Medium priority
                                traffic.
      --dscp-low <dscp>       -  Set DSCP marking for FC-Low priority traffic.
      --dscp-ip-high <dscp>   -  Set DSCP marking for IP-High priority traffic.
                                [7810 / 7840 / SX6 only]
      --dscp-ip-medium <dscp> -  Set DSCP marking for IP-Medium priority
                                traffic. [7810 / 7840 / SX6 only]
      --dscp-ip-low <dscp>    -  Set DSCP marking for IP-Low priority traffic.
                                [7810 / 7840 / SX6 only]
      --l2cos-f-class <l2cos> -  Set L2CoS value for Control priority traffic.
      --l2cos-high <l2cos>    -  Set L2CoS value for FC-High priority traffic.
      --l2cos-medium <l2cos>  -  Set L2CoS value for FC-Medium priority
                                traffic.
      --l2cos-low <l2cos>     -  Set L2CoS value for FC-Low priority traffic.
      --l2cos-ip-high <l2cos> -  Set L2CoS value for IP-High priority traffic.
      --l2cos-ip-medium <l2cos> -  Set L2CoS value for IP-Medium priority
                                traffic.
      --l2cos-ip-low <l2cos>  -  Set L2CoS value for IP-Low priority traffic.
      --sla <sla-name>|none   -  Set the SLA name for this circuit. [7810 /
                                7840 / SX6 only]
      --help                  -  Print this usage statement.

   Example:
     portcfg fcipcircuit 24 create 0 --min-comm-rate 100000
```

3. Set the DSCP marking for VE24 circuit-0 IP Extension high-priority traffic to an example value of 46.

   ```
   portcfg fcipcircuit 24 modify 0 --dscp-ip-high 46
   ```

4. Set the DSCP marking for VE24 circuit-0 IP Extension high-priority traffic back to the default value of 0.

   ```
   portcfg fcipcircuit 24 modify 0 --dscp-ip-high 0
   ```

5. Set the L2CoS marking for VE24 circuit-0 IP Extension low-priority traffic to an example value of 4.

   ```
   portcfg fcipcircuit 24 modify 0 --l2cos-ip-low 4
   ```

6. Set the L2CoS marking for VE24 circuit-0 IP Extension low-priority traffic back to the default value of 0.

   ```
   portcfg fcipcircuit 24 modify 0 --l2cos-ip-low 0
   ```

Traffic Control List (TCL)

A traffic control list (TCL) defines how the ingress LAN side traffic is mapped to a tunnel. A TCL is a list of rules that identifies specific LAN traffic characteristics based on fields that are found in the IP, TCP, and Ethernet headers. The TCL

identifies flows using protocols, ports, source and destination IP addresses, source and destination subnets, QoS values (DSCP or 802.1P), and 802.1Q VLAN tags. Each rule can either allow or deny a matched flow, acting as an input filter to an IP Extension tunnel. Multiple TCL rules arranged by priority provide a high level of control over the LAN side traffic directed to a particular tunnel.

Each TCL rule is identified by a name that is assigned when created. TCL rules are modified or deleted by name. A TCL name contains 31 or fewer alphanumeric characters. Each name must be unique within the Extension platform. Rules are local to each platform and are not shared across platforms.

When traffic matches an allow rule, the rule specifies which tunnel and which QoS to assign for that traffic. In a chassis system with multiple Brocade SX6 Extension Blades, when a rule denies traffic, it applies to all DPs across all blades unless a specific DP is configured.

TCL rules have an order of precedence and each rule has a TCL priority number within the TCL list. Smaller numbers are higher priority, and larger numbers are a lower priority. The TCL priority number must be a unique integer less than 65535. 65535 is the lowest priority (highest number); therefore, it is the last rule enforced. The TCL priority number can be modified to reposition the rule within the list. Gaps between TCL priority numbers are encouraged. When removing a rule, or when creating rules, leave gaps between numbers to allow potential subsequent entries.

NOTE

- The TCL priority number must be unique across all active TCL rules within an IP Extension platform. For example, if a chassis has multiple Brocade SX6 Extension Blades installed, the TCL priority number must be unique across all blades. If a TCL is defined as priority 10, that same priority cannot be used for another TCL rule on a different blade in the same chassis. A check is performed when the rule is enabled to ensure that the priority value is unique.
- The Brocade 7850 Extension Switch and Brocade SX6 Extension Blade provide 1024 defined TCLs with 128 active. The Brocade 7810 Extension Switch provides 256 defined TCLs with 32 active.
- A TCL rule must be configured. One default rule exists, which is to deny all traffic; therefore, if no other rule is created, all traffic is denied. The default rule cannot be removed or modified. To allow traffic through an IP Extension tunnel, one or more TCL rules must be configured that directs matching traffic to a specific tunnel.

A TCL rule has one of two actions, allow or deny. The allow rule specifies the target tunnel into which matching traffic should be directed. Upon the first match, further TCL processing terminates. For the TCL action, the following actions are possible:

- Allow, matching traffic is directed into the target
- Deny, matching traffic is dropped
- Target must be an IP Extension enabled tunnel
- Target is the VE_Port number of the tunnel
- Allow rules are activated on the DP of the target VE_Port
- Deny rules are pushed to all DPs in the chassis
- Target is not specified in a deny rule

  No target should be supplied in a deny priority rule, because the traffic will ultimately not be sent to any tunnel. Deny rules are pushed to all DPs in the chassis, which accounts for traffic that might arrive at either DP. Traffic not destined for a particular IP Extension Gateway (LAN side IPIF unicast address) cannot be forwarded to a specific DP. For example, Broadcast, Unknown unicast, and Multicast (BUM) traffic has no specific known destination, which means to forward to a specific IP Extension Gateway is unknown. Because the LAN side interfaces cannot determine which LAN IPIF to forward such traffic to, all DPs receive the TCL deny rule and all DPs must be capable of handling such traffic.

  If the deny rule has clear matching criteria for deterministic flows, the deny rule can be assigned to a specific DP. The rule is pushed to the specified DP, denying any matching traffic on that DP only.

  After traffic is matched to a TCL rule configured for the allow action, it is sent to the tunnel, and further TCL processing stops. The TCL target parameter specifies the VE_Port tunnel to which the matched traffic will be routed.

Optionally, you can specify a traffic priority (high, medium, low). For example, when configuring the TCL rule, specify `--target 24` or `--target 24-high`. If no priority is specified, the input traffic is sent over the IP Extension medium-priority QoS on the target tunnel. IP Extension traffic priorities are scheduled separately from the FCIP traffic priorities. Each traffic type has its own set of QoS priorities in the egress scheduler.

The TCL is evaluated only once for allow and deny actions. When a TCP stream performs its initial handshake, the TCL is evaluated. If changes are made to TCL rules, including changes to priority rules or enabled/disabled rules, those changes have no effect on existing streams. There is no effect on an existing stream because that stream is already completed its TCP handshake and the TCL is not referenced again. If you do not see any traffic changes after changing a TCL, it might be because the traffic streams have already formed. Newly established streams that do match changed priority rules will be affected.

Before TCL changes can take effect on established traffic streams, one of the following actions must occur:

- The end-device must reset its existing TCP streams, and must form new streams. Disabling and enabling an end-device interface resets its TCP streams.
- The tunnel VE_Port must be disabled and re-enabled. Disabling a VE_Port disrupts all traffic passing through the tunnel (FCIP and IP Extension).
- The LAN interfaces must be disabled and re-enabled. Disruptive to only the IP Extension traffic.

  A best practice approach is to configure the fewest number of allow rules that pass the required traffic. Unmatched traffic encounters the default deny-all rule. You can prevent a subset of the allowed traffic by first using a deny rule before the allow rule. For example, you can deny a smaller IP subnet that is within a larger allowed subnet.

  When IP routes on the end-device are configured correctly, no undesired traffic should appear at the TCL. Nonetheless, you can configure specific deny rules to ensure that certain devices or addresses cannot communicate across the IP Extension tunnel. When configuring and troubleshooting a complex rule set, there is an increased possibility of error.

  In some situations, all traffic that appears at the IP Extension LAN interface is intended to pass through the tunnel. The simplest solution in these situations is to configure an allow rule that simply passes the source and destination subnets.
- Directly connected end-device ports in which all the traffic from the end-device is intended for the remote data center.
- IP routes on the end-devices are closely managed such that only desired traffic is allowed to use the IP Extension LAN gateway.

NOTE

When using Virtual Fabric Logical Switches (VF LS), IP Extension LAN side GE interfaces must be in the Default LS. IP Extension LAN side GE interfaces cannot be assigned to different logical switches. The lan.dp0 and lan.dp1 IPIFs behind the LAN side GE interfaces are not part of VF LS contexts and cannot be assigned to an LS. When an IP datagram arrives at an IP Extension Gateway, it is processed by the TCL.

The LAN side GE interfaces must remain in the Default LS. Even though a VE_Port can reside within an LS other than the Default LS, there is no requirement for LAN side GE interfaces to reside in that LS. The TCL directs traffic to a specified VE_Port regardless of the LS in which it resides.

The target tunnel that is specified in a TCL allow rule must be native to the DP processing the TCL. IP storage traffic arrives on a deterministic DP based on the IP Extension Gateway, which is a LAN IPIF that is on either lan.dp0 or lan.dp1. If the TCL on the DP processing the incoming traffic matches a tunnel, the tunnel must be owned by that DP; otherwise, the TCL finds no valid match and the IP storage traffic is dropped.

Matching Traffic to a Tunnel

When there are multiple data centers, multiple tunnels are used. Between the local and remote IP Extension platforms, normally, only one tunnel is used. A single tunnel uses BET for the aggregated bandwidth by using multiple circuits, failover and failback, lossless link loss (LLL), and other benefits. TCL rules direct traffic to a specific tunnel. When more than one tunnel exists, the various rules match to specific traffic designating the appropriate tunnel for that traffic. TCL rule order counts because processing stops after the first match, the same traffic cannot be forwarded to more than one tunnel.

The following figure shows three local IP storage applications communicating to two remote data centers. The DB application (10.10.10.100) is destined for DC-B. The NAS and tape applications (10.10.10.101 and 10.10.10.102 respectively) are destined for DC-C. The target that is specified in the matching TCL rule directs traffic to the correct tunnel. Extension tunnels are point-to-point; therefore, pointing matched traffic to a particular tunnel sends that traffic to the desired data center. When traffic encounters the first matching TCL allow rule, that action is performed and no additional TCL processing occurs for that particular traffic stream. Any subsequent rules in the TCL are not evaluated.

As shown in the figure, the first rule is for a specific host source IP address of 10.10.10.100. When there is a match, all traffic that is sourced from 10.10.10.100 is sent to tunnel 24. The TCL looks just for this specific host IP address because the prefix length has been set to 32 (subnet mask 255.255.255.255), which indicates that all bits in the address must match. It is a host address and not a subnet address. If the traffic is not sourced from 10.10.10.100, then it will fall through to the next priority in the TCL.

The IP addresses 10.10.10.101 and 10.10.10.102 do not match priority rule 10 (as shown in the following figure). Priority rule 20 is evaluated next. That rule allows IP address 10.10.10.0 with a prefix length of 24 (subnet mask 255.255.255.0), which means that the first 24 bits of the IP address are significant and must match. The last 8 bits are not significant and can vary. If the incoming traffic is sourced from 10.10.10.<any>, it matches this rule and is sent to tunnel 25.

All traffic that does not match the first two priority rules (10 and 20) encounters the final priority rule, priority 65535. The final rule, which cannot be altered or removed, is to deny all traffic. Any traffic reaching the final rule is dropped.

Figure 19: Configuring TCL Rules with Multiple Tunnels Example

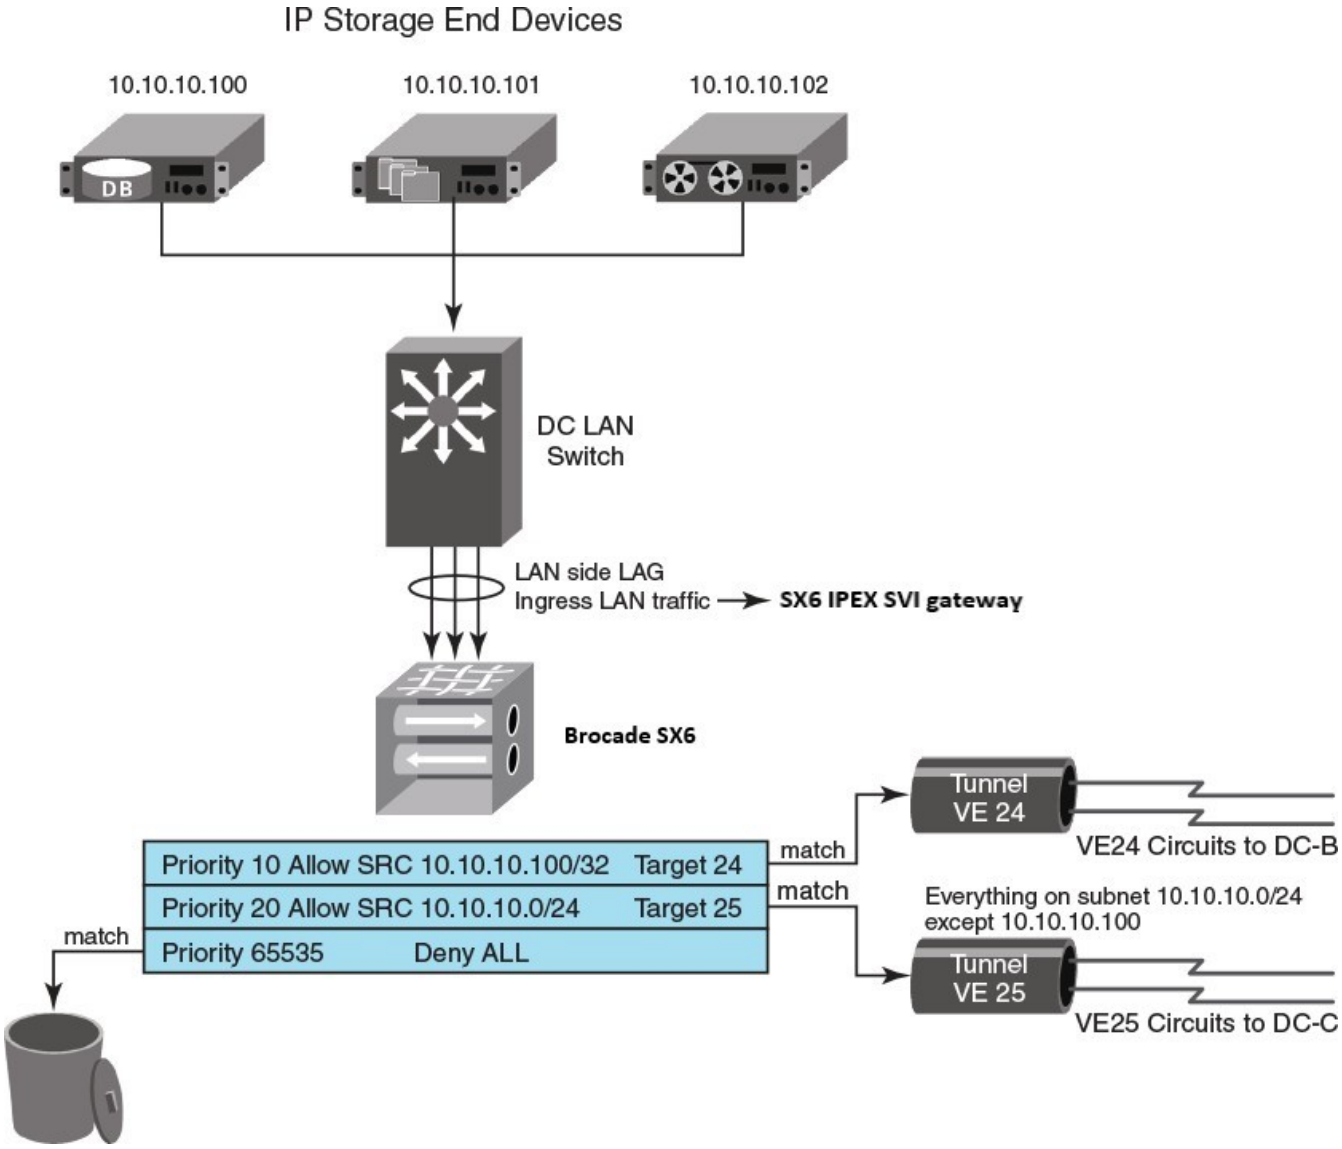

You can use the `portshow tcl` command to display the configured TCL rules as follows:

```
switch:admin> portshow tcl

Pri     Name          Flgs  Target    L2COS  VLAN DSCP Proto Port Hit
                      Src-Addr                Dst-Addr
------------------------------------------------------------------------
*10     DC1           AI--- 24-Med    ANY    ANY  ANY  ANY   ANY  0
                      10.10.10.100/32         ANY
*20     DC2           AI--- 25-Med    ANY    ANY  ANY  ANY   ANY  0
                      10.10.10.0/24           ANY
*65535 default        D---- -         ANY    ANY  ANY  ANY   ANY  0
                      ANY                     ANY
------------------------------------------------------------------------
Flags: *=Enabled ..=Name Truncated (see --detail for full name)
       A=Allow D=Deny I=IP-Ext P=Segment Preservation
```

```
R=End-to-End RST Propagation N=Non-terminated.
```

Non-Terminated Traffic

A TCL rule can be created to prevent traffic from being processed by IP Extension, which is referred to as a non-terminated traffic. The non-terminated TCL option allows TCP traffic to be sent as-is to the remote endpoint through the tunnel. As non-terminated traffic does not add to the consumed TCP streams, non-terminated streams do not consume streams from the maximum supported number of streams.

The non-terminated TCL option is meant to be used by low-bandwidth control-type connections; generally, these connections do not play a role in optimizing the performance of the applications. Control traffic is typically matched using the source and destination IP addresses and the TCP port of the protocol. Often control traffic can be differentiated from data traffic by the TCP port. For non-terminated traffic, use the lowest TCL priority number, which is the highest priority, because it must match before associated data traffic that is terminated. Any TCL rule can be administratively enabled or disabled, including rules for matching non-terminated traffic. However, a rule that changes administrative status requires TCP sessions of the traffic to be reformed; stop and restart the traffic from the application or host to ensure that proper handling after a TCL rule is altered.

The following configuration steps show how to create a TCL named FromSubnetA, enable it, specify a target tunnel, match source and destination subnets, and the rule's priority:

1. Connect to the switch and log on as admin.
2. Use the `portcfg tcl` command to create a TCL.

```
switch:admin> portcfg tcl FromSubnetA create \
--admin enable \
--target 7/24 \
--src-addr 10.10.0.0/24 \
--dst-addr 10.20.0.0/24 \
--priority 10
Operation Succeeded
```

3. Use the `portshow tcl` command to display the TCL information. The `portshow tcl --help` command displays the options available for displaying the TCL information.

```
switch:admin> portshow tcl

Pri    Name         Flgs   Target      L2COS VLAN DSCP Proto Port Hit
                    Src-Addr                Dst-Addr
------------------------------------------------------------------------
*10       FromSubnetA  AI---  7/24-Med    ANY   ANY  ANY  ANY   ANY  10457
                    10.10.0.0/24            10.20.0.0/24
*65535 default      D----  -           ANY   ANY  ANY  ANY   ANY  1229384
                    ANY                     ANY
------------------------------------------------------------------------
Flags: *=Enabled ..=Name Truncated (see --detail for full name)
       A=Allow D=Deny I=IP-Ext P=Segment Preservation
       R=End-to-End RST Propagation N=Non-terminated.
```

Configuring TCL for Non-Terminated Traffic

You can create a TCL rule for traffic that does not terminate at the Brocade Extension platform. The non-terminated TCL option allows TCP traffic to be sent as-is to the other endpoint over the tunnel.

Configure non-terminated TCLs with a higher priority (lower number) than terminated TCLs. For example, a priority value of 10 is higher because it is evaluated before a priority value of 100. By doing so, non-terminated flows between the same

IP addresses match before the terminated flows. In this case, the differentiator when matching is not the IP addresses, it is a specific TCP port number (`--proto-port`) for the traffic.

Perform the following steps to configure a TCL rule for the non-terminated traffic. By default, the `--non-terminated` option is disabled.

1. Connect to the switch and log on as admin.
2. Use the `portcfg tcl <rule_name> create` command with the `--non-terminated` enable option.

```
switch:admin> portcfg tcl control_1 create \
--priority 500 \
--proto-port 3400-3500 \
--non-terminated enable \
--admin-status enable \
--target 25-High
```

NOTE
To disable non-terminated on a TCL rule, use the `portcfg tcl <rule_name> modify` command with the `--non-terminated disable` option.

```
switch:admin> portcfg tcl control_1 modify --non-terminated disable
```

3. Use the `portshow tcl` command to display a summary of TCL rule information.

```
switch:admin> portshow tcl

Pri    Name        Flgs   Target    L2COS  VLAN DSCP Proto Port        Hit
                   Src-Addr                Dst-Addr
-------------------------------------------------------------------------
*500   control_1   AI---N 7/25-High ANY    ANY  ANY  TCP   3400-3500   0
                   ANY                     ANY
*65535 default     D-----   -       ANY    ANY  ANY  ANY   ANY         0
                   ANY                     ANY
-------------------------------------------------------------------------
Flags: *=Enabled ..=Name Truncated (see --detail for full name)
       A=Allow D=Deny I=IP-Ext P=Segment Preservation
       R=End-to-End RST Propagation N=Non-terminated.
```

4. Use the `portcfgshow tcl control_1` command with the `--detail` option to display the detailed TCL rule information. In this example, the non-terminated function is disabled.

```
switch:admin> portshow tcl control_1 --detail

TCL: control_1
================================================================
Admin Status:        Enabled
Priority:            500
Target:              7/25-High (tid:8)
VLAN:                ANY
L2COS:               ANY
DSCP:                ANY
Source Address:      ANY
Destination Address: ANY
L4 Protocol:         ANY
Protocol Port:       ANY
Action:              Allow DP0
RST Propagation:     Disabled
Segment Preservation:Disabled
Non Terminated:      Disabled    <==
Cfgmask:             0x08cc3807
```

```
Hit Count:           85214
```

Configuring TCL App-Type

You can create an app-type to use for TCL matching. Create a human readable app-type when an application uses a port or range of ports in a TCL rule.

Perform the following steps to create an app-type that is used for TCL matching:

1. Connect to the switch and log on as admin.
2. Use the `portcfg app-type <app_type_name> create` command to configure an app-type with an associated port or range of ports.

```
switch:admin> portcfg app-type SpecialApp create \
--portrange 21500-21600 \
--description "Special application"
Operation Succeeded
```

3. Use the `portshow app-type` command to display the configured app-types. All app-types are displayed as shown in the following example:

```
switch:admin> portshow app-type

Application          Port Ranges              Description
-------------------------------------------------------------------------------
CIFS                 139,445
Data-Domain          2051                     EMC Data Domain FCIP   3225-3226
FTP                  20-21,989-990,115        Includes Control data FTPS and Simple FTP
HTTP                 80,8080,8000-8001,3128
HTTPS                443
Isilon-SyncIQ        5666-5667
LDAP                 389,8404,636             Includes LDAP secure
MS-SQL               1443
MySQL                3306
NETAPP-SNAP-MIRROR   10566
NFS                  2049
ORACLE-SQL           66,1525,1521
RSYNC                873
SRDF                 1748
SSH                  22
SSL-SHELL            614
SpecialApp           21500-21600              Special application  <<<<<
TELNET               23,107,513,992           Includes telnets connections
TFTP                 69                       UDP File Transfer
VERITAS-BACKUP       6101-6102,6106,3527,1125 Does not include Veritas Net
backup
VTS-GRID Control     1415-1416
VTS-GRID Data        350
iSCSI                3260
```

4. Use the `portcfg app-type <app_type_name> delete` command to remove a custom app-type.

```
switch:admin> portcfg app-type SSH delete
Unable to modify or delete default app-types
```

You cannot delete or modify a default app-type.

```
switch:admin> portcfg app-type SSH delete
Unable to modify or delete default app-types
```

IP Extension TCL Configuration

The following example configuration shows two data center LANs connected by IP Extension:

```
SiteA:admin> portcfg tcl SiteA_to_SiteB \
--priority 100 \
--target 24 \
--admin-status enable \
-S 10.10.0.0/24 \
-D 192.168.0.0/24

SiteB:admin> portcfg tcl SiteB_to_SiteA \
--priority 100 \
--target 24 \
--admin-status enable \
-S 192.168.0.0/24 \
-D 10.10.0.0/24
```

The following figure shows a configuration that shows two data center LANs connected by IP Extension:

Figure 20: IP Extension TCL Configuration Example

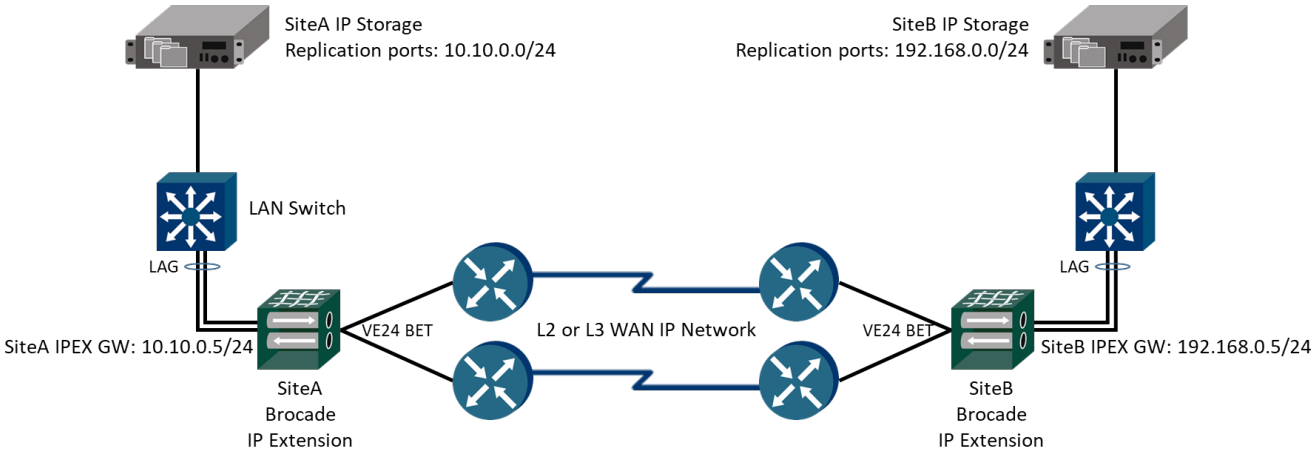

Tunnel Out-of-Order Delivery

IP Extension TCP stream of each end-device is a logically independent stream within the tunnel. Head of line blocking (HoLB) is mitigated for each TCP flow through the tunnel. Packet loss recovery is isolated to the impacted stream that the data belongs to. Data is always delivered in-order to each end-device; however, data can be delivered out-of-order relative to different end-devices. HoLB is mitigated by allowing free flowing data flows to pass stalled data flows.

Putting It All Together

This section describes all three sides of a Brocade Extension platform: FC and FICON, WAN, and LAN.

The WAN side and LAN side deployments are L3 deployments and are independent. This means that because one is L3 deployment does not imply that the other should be as well. The WAN side can be L2 (no routing) while the LAN side is L3 (that is, uses PBR), and in reverse, or both can be L2.

The platform in the example is a Brocade 7850 Extension Switch. The LAN side GE interfaces are connected to a data center switch that supports 25GbE and a portchannel is configured. The WAN side GE interfaces are connected to a data center switch that supports 100GbE. The `ping` command is used to confirm the connectivity. IPsec and eHCL are configured. The WAN and data center network support 9216 byte MTU.

The following figure shows that the replication traffic traverses the Brocade Extension platforms:

Figure 21: Example FC and IP Replication SAN

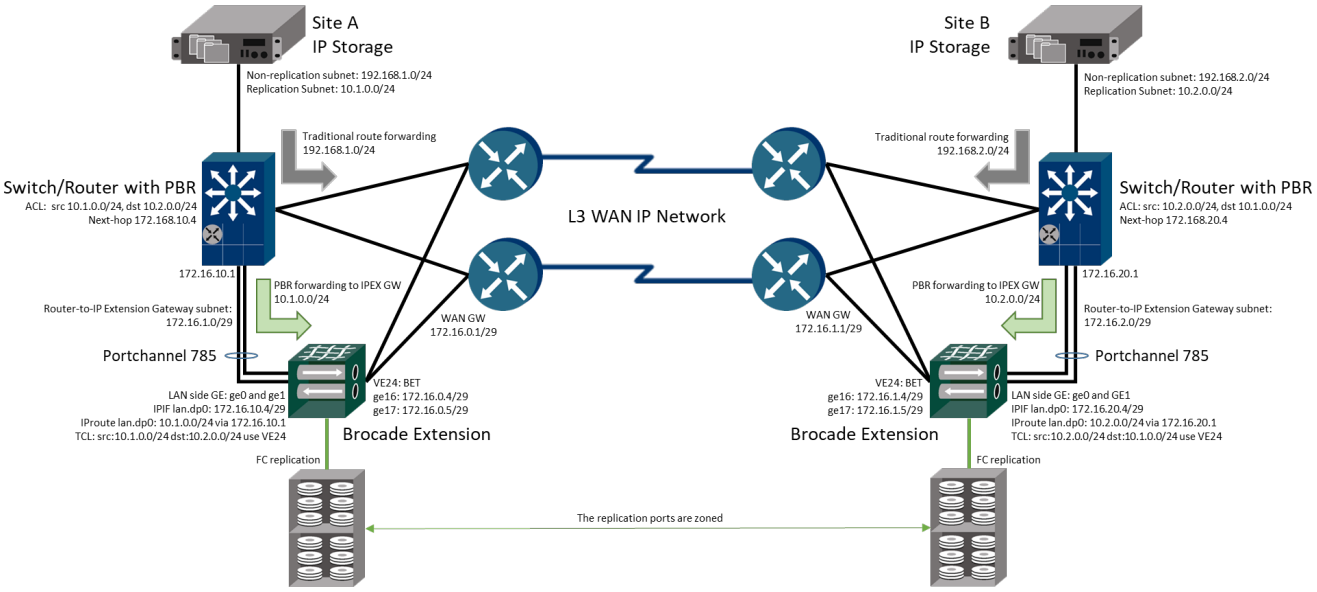

GE Interfaces

Site A

portcfgge ge0 --set -speed 25G

portcfgge ge1 --set -speed 25G

portcfgge ge0 --set -lan

portcfgge ge1 --set -lan

portchannel --create MyLAG -type dynamic -key 785

portchannel --add MyLAG -port ge0-1

portenable ge0-1

Site B

portcfgge ge0 --set -speed 25G

portcfgge ge1 --set -speed 25G

portcfgge ge0 --set -lan

portcfgge ge1 --set -lan

portchannel --create MyLAG -type dynamic -key 785

portchannel --add MyLAG -port ge0-1

portenable ge0-1

WAN Side (Layer 3 WAN Network)

Site A

portcfg ipif ge16.dp0 create 172.16.0.4/29 mtu 9216

portcfg ipif ge17.dp0 create 172.16.0.5/29 mtu 9216

portcfg ipif ge16.dp1 create 172.16.0.6/29 mtu 9216

portcfg ipif ge17.dp1 create 172.16.0.7/29 mtu 9216

portshow ipif

portcfg iproute ge16.dp0 create 172.16.1.0/29 172.16.0.1

portcfg iproute ge17.dp0 create 172.16.1.0/29 172.16.0.1

portcfg iproute ge16.dp1 create 172.16.1.0/29 172.16.0.1

portcfg iproute ge17.dp1 create 172.16.1.0/29 172.16.0.1

portcfgshow iproute

portcmd --ping ge16.dp0 -s 172.16.0.4 -d 172.16.0.1

portcmd --ping ge17.dp0 -s 172.16.0.5 -d 172.16.0.1

portcmd --ping ge16.dp1 -s 172.16.0.6 -d 172.16.0.1

portcmd --ping ge17.dp1 -s 172.16.0.7 -d 172.16.0.1

portcfg ipsec-policy MyIPsec create -k <PSK>

portshow ipsec-policy -p

portcfg fciptunnel 24 create

portcfg fciptunnel 24 modify --ipext enable

portcfg fciptunnel 24 modify --ipsec MyIPsec

portcfg fciptunnel 24 modify --ip-compression aggr-deflate

portcfg fciptunnel 24 modify --fc-compression fast-deflate

portcfg fcipcircuit 24 create 0

portcfg fcipcircuit 24 create 1

portcfg fcipcircuit 24 modify 0 --local-ip 172.16.0.4 --remote-ip 172.16.1.4

portcfg fcipcircuit 24 modify 1 --local-ip 172.16.0.5 --remote-ip 172.16.1.5

portcfg fcipcircuit 24 modify 0 --local-ha-ip 172.16.0.6 --remote-ha-ip 172.16.1.6

portcfg fcipcircuit 24 modify 1 --local-ha-ip 172.16.0.7 --remote-ha-ip 172.16.1.7

portcfg fcipcircuit 24 modify 0 --min 5G --max 10G

portcfg fcipcircuit 24 modify 1 --min 5G --max 10G

portshow fciptunnel -c [<--config>|<--perf>|<--qos>|<--detail>|<--hcl-status>]

Site B

portcfg ipif ge16.dp0 create 172.16.1.4/29 mtu 9216

portcfg ipif ge17.dp0 create 172.16.1.5/29 mtu 9216

portcfg ipif ge16.dp1 create 172.16.1.6/29 mtu 9216

portcfg ipif ge17.dp1 create 172.16.1.7/29 mtu 9216

portshow ipif

portcfg iproute ge16.dp0 create 172.16.0.0/29 172.16.1.1

portcfg iproute ge17.dp0 create 172.16.0.0/29 172.16.1.1

portcfg iproute ge16.dp1 create 172.16.0.0/29 172.16.1.1

portcfg iproute ge17.dp1 create 172.16.0.0/29 172.16.1.1

portcfgshow iproute

portcmd --ping ge16.dp0 -s 172.16.1.4 -d 172.16.1.1

portcmd --ping ge17.dp0 -s 172.16.1.5 -d 172.16.1.1

portcmd --ping ge16.dp1 -s 172.16.1.6 -d 172.16.1.1

portcmd --ping ge17.dp1 -s 172.16.1.7 -d 172.16.1.1

portcmd --ping ge16.dp0 -s 172.16.1.4 -d 172.16.0.4

portcmd --ping ge17.dp0 -s 172.16.1.5 -d 172.16.0.5

portcmd --ping ge16.dp1 -s 172.16.1.6 -d 172.16.0.6

portcmd --ping ge17.dp1 -s 172.16.1.7 -d 172.16.0.7

portcfg ipsec-policy MyIPsec create -k <PSK>

portshow ipsec-policy -p

portcfg fciptunnel 24 create

portcfg fciptunnel 24 modify --ipext enable

portcfg fciptunnel 24 modify --ipsec MyIPsec

portcfg fciptunnel 24 modify --ip-compression aggr-deflate

portcfg fciptunnel 24 modify --fc-compression fast-deflate

portcfg fcipcircuit 24 create 0

portcfg fcipcircuit 24 create 1

portcfg fcipcircuit 24 modify 0 --local-ip 172.16.1.4 --remote-ip 172.16.0.4

portcfg fcipcircuit 24 modify 1 --local-ip 172.16.1.5 --remote-ip 172.16.0.5

portcfg fcipcircuit 24 modify 0 --local-ha-ip 172.16.1.6 --remote-ha-ip 172.16.0.6

portcfg fcipcircuit 24 modify 1 --local-ha-ip 172.16.1.7 --remote-ha-ip 172.16.0.7

```
portcfg fcipcircuit 24 modify 0 --min 5G --max 10G
portcfg fcipcircuit 24 modify 1 --min 5G --max 10G
portshow fciptunnel -c [<--config>|<--perf>|<--qos>|<--detail>|<--hcl-status>]
```

LAN Side (Layer 3 Deployment)

Site A

```
portcfg ipif lan.dp0 create 172.16.10.4/29 mtu 9216
portshow ipif lan
portcfg iproute lan.dp0 create 10.1.0.0/24 172.16.10.1
portcfgshow iproute lan
portcfg tcl SiteA_to_SiteB \
--priority 100 \
--admin-status enable \
--target 24 \
--src-addr 10.1.0.0/24 \
--dst-addr 10.2.0.0/24
portshow tcl
```

Site B

```
portcfg ipif lan.dp0 create 172.16.20.4/29 mtu 9216
portshow ipif lan
portcfg iproute lan.dp0 create 10.2.0.0/24 172.16.20.1
portcfgshow iproute lan
portcfg tcl SiteB_to_SiteA \
--priority 100 \
--admin-status enable \
--target 24 \
--src-addr 10.2.0.0/24 \
--dst-addr 10.1.0.0/24
portshow tcl
portshow lan-stats --per-flow -all
```

Extension and Virtual Fabrics

This section describes how to configure extensions with logical switches. The configuration of the Brocade Extension with Virtual Fabrics (VF) might differ from the configuration of Brocade Extension without Virtual Fabrics.

The following is a brief overview of VF Logical Switch (LS) concepts and terminology, followed by details on configuring extensions with logical switches.

NOTE
The Brocade 7810 Extension Switch does not support VFs.

Virtual Fabrics Considerations and Limitations

The following considerations and limitations apply to a VF:

- A Base Switch is a type of LS used for Fibre Channel Routing (FCR) and extended inter-switch links (xISL) between Base Switches.
- xISLs tag FC traffic so that multiple FIDs can utilize common physical links, rather than separate ISLs from each LS.
- Up to 16 LSs support FICON on a Brocade X6 or Brocade X7 Director; however, refer to your OEM qualification for specific limits.
- You cannot intermix GE interfaces and VE_Ports from different Brocade SX6 Extension Blades into the same LS. The GE interfaces and VE_Ports must be on the same blade.

When Brocade 7850 Extension Switch is in 6VE mode or Brocade SX6 Extension Blade is in 10VE mode, all unused VE_Ports must be in the Default LS. If the unused VE_Ports are not in the Default LS, the VE mode cannot be set. Disabled VE_Ports cannot be moved to other logical switches while in 6VE or 10VE mode.

Logical Switches

VFs are the umbrella technology under which logical switches (LS) operate. An LS allows the division of a physical chassis into multiple fabric elements. Each of these fabric elements is referred to as an LS. Another way to refer to an LS is fabric ID (FID). Multiple LSs connected by the same FID form a logical fabric (LF). Each LS functions as an independent self-contained FC switch. Each chassis can have multiple LSs, each with a unique FID. Up to 16 LSs are supported in directors that support the Brocade SX6 Extension Blades. The Brocade 7850 Extension Switch supports a total of four logical switches, while the Brocade 7810 Extension Switch does not support VF.

The following types of logical switches are supported:

- Default LS (created when enabling VF, only one Default LS exists)
- FC LS (default when creating an LS, multiple FC LS can exist)
- Base LS (must be specified, only one Base LS can exist)
- FICON LS (must be specified, multiple FICON LS can exist)

This section is not meant to be a comprehensive VF guide, but to discuss the specifics of VF and LSs for Brocade Extension. For more details on managing and configuring VFs and LSs, refer to the "Managing Virtual Fabrics" section in the *Brocade Fabric OS Administration Guide*.

Logical Fabrics

An LF contains at least two or more logical switches connected together forming a fabric. An LF is a fabric of logical switches. To connect logical switches into an LF, the FIDs must be the same. FID and DID are not the same. It is not possible to have the same FID more than once on the same platform.

The Brocade Extension platforms support the following types of VF ISL connectivity:

- DISL: Connecting an LS to another LS using an ISL.
- xISL: Connecting a Base Switch to another Base Switch with LS tagged traffic ISL.
- FCIP: Supports DISL (LS to LS) and xISL (LS tagged traffic) connectivity.

Fabric ID (FID)

When creating an LS, a FID must be assigned in the range from 1 to 128. Each FID uniquely identifies an LS on the platform and also indicates which fabric the LS belongs to. No two LSs on the platform can have the same FID. An LS in one platform can communicate with an LS in another if the FID is the same; this is referred to as a logical fabric (LF). Only LSs with the same FID can form a logical fabric. When two LSs with different FIDs are connected, the link between them is segmented. A platform with VF enabled can communicate with platforms that do not have VF enabled. The default LS is initially assigned FID 128, which can be changed if needed.

Create an LS using the `lsCfg` command. For details, refer to the instructions for creating a logical switch or base switch section in the *Brocade Fabric OS Administration Guide* and the `lsCfg` command in the *Brocade Fabric OS Command Reference Manual*.

Logical Switch Contexts

Configuring features and performing tasks in the context of an LS is the same as any FC switch. Enter the context of the FID and the command line will be prompted for that LS. The following items are the two methods for changing to an LS context so that tasks can be performed:

- Use the `setcontext` *`fabricID`* command, which changes the CLI context to a specific LS. The CLI prompt changes to reflect the new FID. Any commands entered are initiated on the LS with that FID.
- Use the `fosexec --fid` *`FID`* `-cmd` *`command -string`* command to initiate a specific command on a specific LS, where *command-string* is the command string that you want to perform; the `fosexec` command must be enabled on the remote platform.

Using the `fosexec` command to issue a command on an LS, runs the specified FOS command on the specified LS, whereas using the `setcontext` command changes the CLI context to the desired LS.

Default Logical Switch

Before using VF, it must first be enabled. Enabling VF creates a default LS with FID 128, and initially contains all FC ports and GE interfaces on the platform. The default switch always exists and cannot be deleted. Disabling VF deletes all LSs. Additional FC and FICON LSs can be added and deleted. The number of LSs that can be created depends on the platform. After creating LS, the chassis operates with multiple independently managed switches.

Extension GE Interfaces and Logical Switches

By default, extension GE interfaces live in the default switch. Extension GE interfaces in the Default LS can be shared among circuits originating from any LS. Until explicitly moved to another LS, all FC, FICON, and GE interfaces reside in the Default LS.

Ethernet Port Sharing

VE_Ports in different LS can share a single GE interface in the default LS. Keep GE interfaces in the default LS (FID 128), even if only one LS is used. The GE interface does not have to be moved back to the default switch when a new LS with GE interface is added, such as a Brocade 7850 Extension Switch 100GbE interface.

NOTE

- If a GE interface is in the default FID, it can be shared across multiple FIDs and VEs. If a GE interface is not in the default FID, it can only be used in that FID.
- When using shared GE interfaces, if the Default LS is disabled, the GE interfaces in the Default LS are disabled as well, which impacts all circuits from all LSs using the shared GE interfaces.
- A LAN side GE interface must be in the Default LS.

With Ethernet Port Sharing, you can have the following configuration, as an example:

- The default LS has interface GE16.
- LS1 has VE_Port 24 (tunnel 24), which has a circuit over GE16.
- LS2 has VE_Port 33 (tunnel 33), which also has a circuit over GE16.
- LAN side interfaces remain in the default LS.

All license committed-rate restrictions and bandwidth guidelines for GE interfaces remains the same for shared GE interfaces. VE_Ports using shared GE interfaces initiate as a regular tunnel from its respective LS.

IPIFs and IP routes are created in the LS in which the GE interface lives. If a non-Default LS is used for the tunnels and circuits, then the IPIFs and IP routes are created in the Default LS for a GE interface living in the Default LS. An IPIF and IP route can be used by any circuit that is created in a non-Default LS or in the Default LS. This means that VE_Ports from different LSs can use the same GE interface. Although multiple circuits from different LSs can use the same GE interface, these circuits can be differentiated to the IP network by using VLAN tags, QoS markings, or the source and destination IP addresses.

VE_Ports and Logical Switches

Initially, all ports start in the Default LS. When you create an additional LS, the new LS is empty, and ports can be assigned to the new LS. When assigning ports to an LS, the ports are removed from originating LS and moved into the destination LS.

Port assignment requires the following requirements:

- Ports are moved from one LS to another.
- A port cannot be deleted from the platform.
- A port can only exist in one LS at a time.
- An LS can have as many ports as there are ports on the chassis.
- Ports with configuration settings cannot be moved without first deleting the existing settings.

A VE_ Port cannot be moved from one LS to another as long as a configured tunnel exists. First, delete the tunnel that is using the VE_Port, then move the VE_Port to the desired LS. Similarly, a GE interface cannot be moved between LSs until the existing IPIFs and IP routes have been deleted; although, there is no reason to move GE interfaces out of the Default LS.

As a recommended best practice, keep GE interfaces in the Default LS; do not move them to another LS. There is no reason to move GE interfaces because of the Ethernet Port Sharing (EPS) feature. Circuits in any LS can use a GE interface in the Default LS. Moving a GE interface out of the Default LS prevents it from being available to circuits from other LSs. See Ethernet Port Sharing section.

Use the `lsCfg -config <i>slot</i>/<i>ge_port</i>` command to move ports from one FID to another. The FID is the fabric ID of the LS that you want to move the ports to. The ports are automatically removed from their current LS and placed into the designated LS.

Moving VE_Ports and GE Interfaces between LS

You can use the following command to move ports and interfaces between one LS and another:

```
lscfg --config FID port [-slot slot | slot range] [-port port | port range])
```

The following considerations apply when you move ports between LSs:

- The FID is the Fabric ID of the LS to which the port is moving.
- Extension GE interfaces can be added to any LS.
- Extension GE interfaces cannot be moved if they have any configuration other than speed and autonegotiation.
- Extension VE_Ports cannot be moved if they have any configuration.
- The slot number is required only for the Brocade SX6 Extension Blade.
- The port number is the FC, VE, or GE port number.
- GE interfaces are independent of each other, and must be moved in independently.

Perform the following steps to move ports between LSs:

1. Connect to the switch and log on as admin.
2. Enter the `Lscfg --help` command to view the help options.

```
switch:FID128:admin> lscfg --help

lscfg [--create | --delete | --config | --show | --change | --help]
------------------
lscfg --create FID [-base | -lisldisable | -ficon ] [-force]
        FID:          Required.  Fabric ID for this switch.
                          Must be unique.
        -base:        Optional.  Make this switch the base switch.
                                 Only a single switch may be the base
                                 switch in a chassis.
        -lisldisable: Optional.  Use this option to bring up the lisl
                                 port(s) in Offline state.
                                 When not used lisl port(s) are brought
                                 Online
        -ficon:       Optional.  Creates logical switch auto-configured
                                 with FICON configurations.
        -force:       Optional.  Foregoes all user prompts.

        Note: "base" and "ficon" options are mutually exclusive.
------------------
lscfg --delete FID [-force]
        FID: Required.  Fabric ID for the switch to be deleted.
                        All ports must be removed from this switch
                        prior to issuing the command.
        -force:  Optional.  Foregoes all user prompts.
------------------
lscfg --config FID [-port port | port range]
        FID:    Required.  Fabric ID for the switch.
        -port:  Required.  The port to be added to the switch.
                           A range may be supplied (example: -port 6-16).
        -force: Optional.  Foregoes all user prompts.
------------------
lscfg --show [-ge] [-instance]
        -ge:         Optional.  Display GE ports.
        -instance:      Optional.  Display LS Instance.
------------------
lscfg --change FID [-newfid fid] [-base] [-ficon]
        FID:     Required.  Fabric ID for the switch.
        -newfid: Optional.  The new fid for the switch.
```

```
            -base:   Optional.  Makes this switch the base switch.
                                If the current base switch is indicated,
                                this will remove the base switch
    attribute.
            -ficon:  Optional.  Makes this switch a FICON mode switch
            -force:  Optional.  Foregoes all user prompts.
```

3. Move the port or interface to the destination FID. The following example moves VE_Port 25 port to FID-1:

```
switch:FID128:admin> lscfg --config 1 -port 25

This operation requires that the affected ports be disabled.
Would you like to continue [y/n]?: y
Making this configuration change.  Please wait...
Configuration change successful.
Please enable your ports/switch when you are ready to continue.
```

4. Move the port or interface to the destination FID. The following example moves GE0 to FID-1:

```
switch:FID128:admin> lscfg --config 1 -port ge0

This operation requires that the affected ports be disabled.
Would you like to continue [y/n]?: y
Making this configuration change.  Please wait...
Configuration change successful.
Please enable your ports/switch when you are ready to continue.
```

Displaying the LS Configuration

You can display the logical switch configuration for a switch and the GE interfaces that are located in each LS using the `lscfg --show -ge` command.

The following output shows that all GE interfaces are located in the default LS (FID 128):

```
switch:admin> lscfg --show -ge

Created switches FIDs(Domain IDs): 128(ds)(80) 60(1)
Slot 1   2   3   4   5   6    7     8
GE    Port
0  |   |   |   |   |   |   | 128 | 128 |
1  |   |   |   |   |   |   | 128 | 128 |
2  |   |   |   |   |   |   | 128 | 128 |
3  |   |   |   |   |   |   | 128 | 128 |
4  |   |   |   |   |   |   | 128 | 128 |
5  |   |   |   |   |   |   | 128 | 128 |
6  |   |   |   |   |   |   | 128 | 128 |
7  |   |   |   |   |   |   | 128 | 128 |
8  |   |   |   |   |   |   | 128 | 128 |
9  |   |   |   |   |   |   | 128 | 128 |
10 |   |   |   |   |   |   | 128 | 128 |
11 |   |   |   |   |   |   | 128 | 128 |
12 |   |   |   |   |   |   | 128 | 128 |
13 |   |   |   |   |   |   | 128 | 128 |
14 |   |   |   |   |   |   | 128 | 128 |
15 |   |   |   |   |   |   | 128 | 128 |
16 |   |   |   |   |   |   | 128 | 128 |
17 |   |   |   |   |   |   | 128 | 128 |
```

Use the `lscfg --show` command to display the logical switches and the VE_Ports assigned to each. The following output shows that the Default LS (FID 128) has Domain ID 80, and there is an LS (FID 60) with Domain ID 1.

NOTE
The VE_Port 24 has been moved from FID 128 to FID 60 in this example.

```
switch:admin> lscfg --show

Created switches FIDs(Domain IDs): 128(ds)(80) 60(1)
Slot  1    2      3     4      5     6      7       8
Port
0   |    |    |  128  |    |  128  |    |  128  |  128  |
1   |    |    |  128  |    |  128  |    |  128  |  128  |
2   |    |    |  128  |    |  128  |    |  128  |  128  |
3   |    |    |  128  |    |  128  |    |  128  |  128  |
4   |    |    |  128  |    |  128  |    |  128  |  128  |
5   |    |    |  128  |    |  128  |    |  128  |  128  |
6   |    |    |  128  |    |  128  |    |  128  |  128  |
7   |    |    |  128  |    |  128  |    |  128  |  128  |
8   |    |    |  128  |    |  128  |    |  128  |  128  |
9   |    |    |  128  |    |  128  |    |  128  |  128  |
10  |    |    |  128  |    |  128  |    |  128  |  128  |
11  |    |    |  128  |    |  128  |    |  128  |  128  |
12  |    |    |  128  |    |  128  |    |  128  |  128  |
13  |    |    |  128  |    |  128  |    |  128  |  128  |
14  |    |    |  128  |    |  128  |    |  128  |  128  |
15  |    |    |  128  |    |  128  |    |  128  |  128  |
16  |    |    |  128  |    |  128  |    |  128  |  128  |
17  |    |    |  128  |    |  128  |    |  128  |  128  |
18  |    |    |  128  |    |  128  |    |  128  |  128  |
19  |    |    |  128  |    |  128  |    |  128  |  128  |
20  |    |    |  128  |    |  128  |    |  128  |  128  |
21  |    |    |  128  |    |  128  |    |  128  |  128  |
22  |    |    |  128  |    |  128  |    |  128  |  128  |
23  |    |    |  128  |    |  128  |    |  128  |  128  |
24  |    |    |  128  |    |  128  |    |  60   |  128  |
25  |    |    |  128  |    |  128  |    |  128  |  128  |
26  |    |    |  128  |    |  128  |    |  128  |  128  |
27  |    |    |  128  |    |  128  |    |  128  |  128  |
28  |    |    |  128  |    |  128  |    |  128  |  128  |
29  |    |    |  128  |    |  128  |    |  128  |  128  |
30  |    |    |  128  |    |  128  |    |  128  |  128  |
31  |    |    |  128  |    |  128  |    |  128  |  128  |
32  |    |    |  128  |    |       |    |  128  |  128  |
33  |    |    |  128  |    |       |    |  128  |  128  |
34  |    |    |  128  |    |       |    |  128  |  128  |
35  |    |    |  128  |    |       |    |  128  |  128  |
36  |    |    |  128  |    |       |    |       |       |
37  |    |    |  128  |    |       |    |       |       |
38  |    |    |  128  |    |       |    |       |       |
39  |    |    |  128  |    |       |    |       |       |
40  |    |    |  128  |    |       |    |       |       |
41  |    |    |  128  |    |       |    |       |       |
42  |    |    |  128  |    |       |    |       |       |
```

```
 43  |    |    |  128  |    |       |    |       |       |
 44  |    |    |  128  |    |       |    |       |       |
 45  |    |    |  128  |    |       |    |       |       |
 46  |    |    |  128  |    |       |    |       |       |
 47  |    |    |  128  |    |       |    |       |       |
 48  |    |    |       |    |       |    |       |       |
 49  |    |    |       |    |       |    |       |       |
 50  |    |    |       |    |       |    |       |       |
 51  |    |    |       |    |       |    |       |       |
 52  |    |    |       |    |       |    |       |       |
 53  |    |    |       |    |       |    |       |       |
 54  |    |    |       |    |       |    |       |       |
 55  |    |    |       |    |       |    |       |       |
 56  |    |    |       |    |       |    |       |       |
 57  |    |    |       |    |       |    |       |       |
 58  |    |    |       |    |       |    |       |       |
 59  |    |    |       |    |       |    |       |       |
 60  |    |    |       |    |       |    |       |       |
 61  |    |    |       |    |       |    |       |       |
 62  |    |    |       |    |       |    |       |       |
 63  |    |    |       |    |       |    |       |       |
```

WAN Side IPIF with Logical Switches

When using VF LS, an IPIF must be configured in the LS owning the GE interface. The `portcfg ipif` command must be issued in the LS context where the GE interface resides. If the GE interface is in the Default LS, then these commands are executed in the default LS CLI context. If the GE interface is living in another LS, issue the command in that context.

If the GE interface is not in the default LS, the GE interface cannot be shared by any circuits from another LS. Typically, the GE interfaces stay in the default LS; therefore, the IPIF is created in the default LS. Tunnels and circuits can still be created in different LSs. Other than issuing commands from the correct LS context, all aspects of extension configuration remain the same.

Perform the following tasks to configure an IPIF when using VF LS:

1. Connect to the switch and log on as admin.
2. Set the CLI LS context to the default LS, which is usually FID 128 by default.

```
switch:FID128:admin> setcontext 128
```

3. Verify that the GE interfaces are in the default LS.

```
switch:FID128:ascgmgmt> lscfg --show -ge

Created switches FIDs(Domain IDs):  128(ds)(1)  1(3)  2(5)

GE Port    0     1     2     3     4     5     6     7     8     9
-----------------------------------------------------------------
FID     128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 |

GE Port   10    11    12    13    14    15    16    17
-------------------------------------------------------
FID     128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 |
```

4. Create the IPIF in the default LS.

```
switch:FID128:admin> portcfg ipif ge0.dp0 create 10.10.10.10/24
Operation Succeeded
```

```
switch:FID128:admin> portcfg ipif ge1.dp0 create 10.10.10.11/24
Operation Succeeded
```

5. Show the IPIF while the CLI context has the default LS, usually FID 128 by default.

```
switch:FID128:ascgmgmt> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge0.dp0      10.10.10.10                   / 24   1500  0     U R M
 ge1.dp0      10.10.10.11                   / 24   1500  0     U R M
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

WAN Side IP Routes with Logical Switches

When using VF LS, the IP routes for circuits must be configured in the LS owning the GE interface. The `portcfg iproute` command must be issued in the LS context where the GE interface resides. If the GE interface is in the Default LS, then these commands are executed in the Default LS CLI context. If the GE interface is in another LS, issue the command in that context.

If the GE interface is not in the Default LS, the GE interface cannot be shared by any circuits from another LS. Typically, the GE interfaces stay in the Default LS; therefore, the IP route is created in the Default LS. Tunnels and circuits can still be created in different LSs. Other than issuing commands from the correct LS context, all aspects of extension configuration remain the same.

Perform the following tasks to configure an IP route when using VF LS:

1. Connect to the switch and log on as admin.
2. Set the CLI LS context to the default LS, normally FID 128.

```
switch:FID3:admin> setcontext 128
switch:FID128:admin>
```

3. Configure an IP route using the `portcfg iproute` command in the FID 128 context. The destination subnet is 10.0.0.0/24. The local gateway is 10.10.10.1, which is on the same subnet as the IPIF configured above. The IP route must be added to each GE interface and DP being used. For example, ge0.dp0 and ge0.dp1 both need an IP route added.

```
switch:FID128:admin> portcfg iproute ge0.dp0 create 10.0.0.0/24 10.10.10.1
Operation Succeeded
switch:FID128:admin> portcfg iproute ge1.dp0 create 10.0.0.0/24 10.10.10.1
Operation Succeeded
```

4. Use the `portcfgshow iproute` command in the default LS FID 128 context to display the IP route information.

```
switch:FID128:admin> portcfgshow iproute

 Port         IP Address                  / Pfx  Gateway                Flags
-------------------------------------------------------------------------------
 ge0.dp0      10.0.0.0                    / 24   10.10.10.1             S
 ge1.dp0      10.0.0.0                    / 24   10.10.10.1             S
-------------------------------------------------------------------------------
 Flags: S=Static X=Crossport
```

Tunnels and Circuits with Logical Switches

The process of configuring a tunnel and circuit remains the same, except for moving the VE_Port and creating the tunnels and circuits in the correct LS context.

Perform the following steps to configure a tunnel and circuit on an LS other than the Default LS:

1. Create the IPIF in the context in which the GE interface lives.
2. Create the IP routes in the context in which the GE interface lives.
3. Move the VE_Port to the desired LS from the Default LS.
4. Create the tunnel and circuit in the context in which the VE_Port lives.

Enter the `portcfg fciptunnel` command in the context of the logical switch where the VE_Port resides. In the following example, FID 60 is created. The VE_Port originates in the Default LS (FID 128) and is moved to FID 60. Currently, no tunnel has been configured using the VE24, so it is free to be moved. The context is set to FID 60. From FID 60, the tunnel and circuit are configured.

1. Connect to the switch and log on as admin.
2. Enter the `lscfg` command to create an LS with FID 60.

```
switch:FID128:admin> lscfg --create 60

A Logical switch with FID 60 will be created with default configuration.
Would you like to continue [y/n]?: y
About to create switch with fid=60. Please wait...
Logical Switch with FID (60) has been successfully created.

Logical Switch has been created with default configurations.
Please configure the Logical Switch with appropriate switch
and protocol settings before activating the Logical Switch.
```

3. Move the VE_Port from the default LS to FID 60.

```
switch:FID128:admin> lscfg --config 60 -port 24

This operation requires that the affected ports be disabled.
Would you like to continue [y/n]?: y
Making this configuration change.
Please wait... Configuration change successful.
Please enable your ports/switch when you are ready to continue.
```

4. Set the CLI context to the logical switch with FID 60 using the following command:

```
switch:FID128:admin> setcontext 60
```

5. Re-enable the VE_Port because the VE_Port was disabled when it was moved.

```
switch:FID60:admin> portenable 24
```

6. Create a tunnel on VE24 in FID 60. Add features to the tunnel using the staging method.

```
switch:FID60:admin> portcfg fciptunnel 24 create
Operation Succeeded

switch:FID60:admin> portcfg fciptunnel 24 modify --ipext enable
Operation Succeeded

switch:FID60:admin> portcfg fciptunnel 24 modify --ipsec MyIPsecPolicy
Operation Succeeded

switch:FID60:admin> portcfg fciptunnel 24 modify --ip-compression deflate
Operation Succeeded
```

```
switch:FID60:admin> portcfg fciptunnel 24 modify --fc-compression fast-deflate
Operation Succeeded
```

7. Create two circuits on VE24 in FID 60.

```
switch:FID60:admin> portcfg fcipcircuit 24 create 0
Operation Succeeded

switch:FID60:admin> portcfg fcipcircuit 24 create 1
Operation Succeeded
```

8. Add local and remote IP addresses to the circuits using the IPIF created in the Default LS for interface ge0.dp0 and ge1.dp0.

```
switch:FID60:admin> portcfg fcipcircuit 24 modify 0 --local-ip 192.168.2.2 --remote-ip 192.168.12.2
Operation Succeeded

switch:FID60:admin> portcfg fcipcircuit 24 modify 1 --local-ip 192.168.2.3 --remote-ip 192.168.12.3
Operation Succeeded
```

9. Add the min and max bandwidth to each circuit. The ARL minimum and maximum rates are equal; therefore, ARL is not used, CIR is used instead.

```
switch:FID60:admin> portcfg fcipcircuit 24 modify 0 --min 4000000 --max 4000000
Operation Succeeded

switch:FID60:admin> portcfg fcipcircuit 24 modify 1 --min 4000000 --max 4000000
Operation Succeeded
```

10. Show the status of the tunnel and circuits. Verify that the min and max bandwidth settings are correct.

```
switch_60:FID60:admin> portshow fciptunnel -c

Tunnel Circuit  OpStatus Flags     Uptime TxMBps RxMBps ConnCnt CommRt Met/G
--------------------------------------------------------------------------------
 24    -        Up       --i----PI      9s   0.00   0.00    1       -      -
 24    0 ge0    Up       ----a--i4      9s   0.00   0.00    1   4000/4000 0/-
 24    1 ge1    Up       ----a--i4      9s   0.00   0.00    1   4000/4000 0/-
--------------------------------------------------------------------------------
Flags (tunnel): l=Legacy QOS Mode
                i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                I=IP-Ext
     (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                ARL a=Auto r=Reset s=StepDown t=TimedStepDown S=SLA
```

Enabling VF xISL, Base Switch, and VE_Ports

Two logical switches can be connected together using an extended inter-switch link (xISL). xISL only connect between Base Switches. When you divide a chassis or fixed-port switch into logical switches, one of the switches can be designated as a Base Switch, which is a special LS used for interconnecting logical switches.

A xISL connection between Base Switches can be used instead of separate ISLs. The base provides the physical connectivity across which the logical connectivity is established. xISLs carry the combined traffic from multiple logical fabrics while maintaining fabric separation.

To use xISLs, refer to the "Configuring a Logical Switch" section to use xISLs in the *Brocade Fabric OS Administration Guide*.

1. Connect to the switch and log on as admin.

2. Use the `switchshow` command to see if xISL use is enabled or disabled. In this example, xISL use is disabled (OFF). Notice, when in an LS context, the `switchshow` command shows only the ports and interfaces that live in that context.

```
switch:FID1:admin> switchshow
switchName:     switch
switchType:     148.0
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   3
switchId:       fffc03
switchWwn:      10:00:c4:f5:7c:3b:25:07
zoning:         ON (MyZoneCfg)
switchBeacon:   OFF
FC Router:      OFF
Fabric Name:    EXTENSION_FID1
HIF Mode:       OFF
Allow XISL Use: OFF
LS Attributes:  [FID: 1, Base Switch: No, Default Switch: No, Ficon Switch: No, Address Mode 0]

Index Port Address  Media Speed   State       Proto
==================================================
   0   0   030000   --    N16     No_Module   FC
   1   1   030100   id    N16     No_Light    FC
   2   2   030300   id    N16     No_Light    FC
   3   3   030400   --    N16     No_Module   FC
   4   4   030b00   --    N16     No_Module   FC
   5   5   030c00   id    N16     No_Light    FC
  24  24   030200   --    --       Offline     VE
  25  25   030500   --    --       Offline     VE  Disabled
      ge0           --    40G     No_Module   FCIP  Disabled
```

3. Use the `configure` command to change the xISL setting. The switch does not have to be disabled to change the xISL setting.

```
switch:FID1:admin> configure

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.

Configure...

  Fabric parameters (yes, y, no, n): [no] y

    WWN Based persistent PID (yes, y, no, n): [no]
    Allow XISL Use (yes, y, no, n): [no] y
    Dynamic Portname (on, off): [off]
    Edge Hold Time(Low(80ms), Medium(220ms), High(500ms): (80..500) [220]
    Remote Fosexec feature: (on, off): [off]
    High Integrity Fabric Mode (yes, y, no, n): [no]
  D-Port Parameters (yes, y, no, n): [no]
  RDP Polling Cycle(hours)[0 = Disable Polling]: (0..24) [1]
  System services (yes, y, no, n): [no]
```

4. Verify that xISL is now permitted on the Base Switch.

```
switch:FID1:admin> switchshow
switchName:     switch
switchType:     148.0
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   3
switchId:       fffc03
switchWwn:      10:00:c4:f5:7c:3b:25:07
zoning:         ON (MyZoneCfg)
switchBeacon:   OFF
FC Router:      OFF
Fabric Name:    EXTENSION_FID1
HIF Mode:       OFF
Allow XISL Use: ON
LS Attributes:  [FID: 1, Base Switch: No, Default Switch: No, Ficon Switch: No, Address Mode 0]

Index Port Address  Media Speed   State       Proto
==================================================
   0   0   030000   --    N16     No_Module   FC
   1   1   030100   id    N16     No_Light    FC
   2   2   030300   id    N16     No_Light    FC
   3   3   030400   --    N16     No_Module   FC
   4   4   030b00   --    N16     No_Module   FC
   5   5   030c00   id    N16     No_Light    FC
  24  24   030200   --    --       Offline     VE
  25  25   030500   --    --       Offline     VE  Disabled
      ge0           --    40G     No_Module   FCIP  Disabled
```

5. Create the Base Switch. In the example below, FID 64 is used.

```
switch:FID1:admin> lscfg --create 64 -base

Creation of a base switch requires that the proposed new base switch on this system be disabled.
Would you like to continue [y/n]?: y
About to create switch with fid=64. Please wait...
Logical Switch with FID (64) has been successfully created.

Logical Switch has been created with default configurations.
Please configure the Logical Switch with appropriate switch
and protocol settings before activating the Logical Switch.
```

6. Change context to FID 64.

```
switch:FID1:admin> setcontext 64
```

7. Use `switchshow` to show the current state of the Base Switch. There are no VE_Ports in the FID.

```
switch:FID64:admin> switchshow

switchName:     switch_64
switchType:     148.0
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:c4:f5:7c:3b:25:09
```

```
zoning:         OFF
switchBeacon:   OFF
FC Router:      ON
HIF Mode:       OFF
LS Attributes:  [FID: 64, Base Switch: Yes, Default Switch: No, Ficon Switch: No, Address Mode 0]

Index Port Address  Media Speed   State       Proto
==================================================
================================
No ports found in the system!!!
```

8. Show the logical switches and verify where the VE_Ports currently are.

```
switch_64:FID64:ascgmgmt> lscfg --show

Created switches FIDs(Domain IDs):  128(ds)(1)  1(3)  2(5)  64(bs)(1)

Port      0     1     2     3     4     5     6     7     8     9
-------------------------------------------------------------------
FID       1 |   1 |   1 |   1 |   1 |   1 | 128 | 128 | 128 | 128 |

Port     10    11    12    13    14    15    16    17    18    19
-------------------------------------------------------------------
FID     128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 |

Port     20    21    22    23    24    25    26    27    28    29
-------------------------------------------------------------------
FID     128 | 128 | 128 | 128 |   1 |   1 | 128 | 128 | 128 | 128 |

Port     30    31    32    33    34    35    36    37    38    39
-------------------------------------------------------------------
FID     128 | 128 | 128 | 128 |   2 | 128 | 128 | 128 | 128 | 128 |

Port     40    41    42    43
--------------------------------
FID     128 | 128 | 128 | 128 |
```

9. Move a VE_Port to the Base Switch. In this example, VE26 is moved to FID 64.

```
switch_64:FID64:admin> lscfg --config 64 -port 26

This operation requires that the affected ports be disabled.
Would you like to continue [y/n]?: y
Making this configuration change.  Please wait...
Configuration change successful.
Please enable your ports/switch when you are ready to continue.
```

10. Use `switchshow` to verify that VE26 is in the Base Switch (FID 64). Notice that VE26 was disabled when it was moved.

```
switch_64:FID64:ascgmgmt> switchshow

switchName:     switch_64
switchType:     148.0
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   1
```

```
switchId:       fffc01
switchWwn:      10:00:c4:f5:7c:3b:25:09
zoning:         OFF
switchBeacon:   OFF
FC Router:      ON
HIF Mode:       OFF
LS Attributes:  [FID: 64, Base Switch: Yes, Default Switch: No, Ficon Switch: No, Address Mode 0]

Index Port Address  Media Speed   State       Proto
==================================================
  26  26   010000   --    --       Offline     VE  Disabled
```

11. Enable VE26 and verify its state on the Base Switch. Because VE26 is not yet configured as an extension tunnel, it is in an offline state; however, it is not shown as disabled.

```
switch_64:FID64:ascgmgmt> portenable 26

switch_64:FID64:ascgmgmt> switchshow
switchName:     switch_64
switchType:     148.0
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:c4:f5:7c:3b:25:09
zoning:         OFF
switchBeacon:   OFF
FC Router:      ON
HIF Mode:       OFF
LS Attributes:  [FID: 64, Base Switch: Yes, Default Switch: No, Ficon Switch: No, Address Mode 0]

Index Port Address  Media Speed   State       Proto
==================================================
  26  26   010000   --    --       Offline     VE
```

Brocade Extension Configuration Manager (ECM)

ECM is a CLI-based menu-driven interactive utility that simplifies configuration of the Brocade Extension platforms.

ECM is intended to simplify the configuration for building a tunnel and circuits for FCIP and IP Extension protocols using the following Brocade Extension platforms:

- Brocade 7850 Extension Switch
- Brocade 7810 Extension Switch
- Brocade SX6 Extension Blade

ECM examines the configuration at each stage, checks for errors, detects missing objects, and alerts users to make corrections or add specific parameters. ECM can help you prevent potential mistakes and allow you to correct any mistakes promptly.

For example, when a destination address is entered for a given IPIF, the utility automatically determines if an associated IP route must be configured for the source and destination subnets. Similarly, if a LAN interface, TCL, or IP Extension Gateway is missing from a tunnel configured for IP Extension, the utility displays a warning about an incomplete configuration.

ECM Assumptions and Dependencies

The following assumptions and dependencies apply to ECM:

- ECM executes from the Fabric OS CLI.
- ECM only runs within the logical switch context.
- Configuration parameters are limited to those needed to create a tunnel and circuit.
- You can still use a limited set of available ECM parameters through CLI commands.

ECM Prerequisites

You must gather information before you begin configuring ECM. Use this section to prepare for the configuration process. It is assumed that one tunnel is configured at a time.

Table 35: ECM Prerequisites

Required Information		Your Information
Are there Brocade SX6 Extension Blade slots or not applicable?		
Is there a VE_Port?		
Are there WAN side GE interfaces?		
(IP Extension deployment) Are there LAN side GE interfaces?		
(IP Extension deployment) Is there LAN side Portchannels?		
Information for each GE interface to be used:	Interface speed (platform dependent)? Available options are 1G, 10G, 25G, 40G, and 100G.	
	(1GE only) Auto-negotiate? Available options are enabled or disabled.	
	GE interface mode? Available options are WAN or LAN.	

Required Information		Your Information
Is there a WAN side IPIF to configure? (For example, ge#.dp#)		
Are there WAN side IP addresses for each IPIF?		
Are there WAN side subnet masks or prefix lengths (CIDR) for each IPIF?		
What is the WAN side IPIF MTU? The default value is 1500 bytes.		
What is the WAN side IPIF VLAN (802.1Q)? The default value is none = not tagged.		
(IP Extension deployment [lan.dp#]) Is there a LAN side IPIF to configure?		
(IP Extension Gateway) LAN side IP addresses for each IPIF?		
What are the LAN side subnet masks or prefix lengths (CIDR) for each IPIF gateway?		
Information for the WAN side IP route ECM will not configure a route if it is not needed.	Destination IP subnet and subnet mask	
	Gateway IP address (only if needed)	
(FCIP only mode) Which algorithm for main compression?		
For FC compression (FCIP in hybrid mode), which algorithm?		
For IP compression (IP Extension in hybrid mode), which algorithm?		
How many circuits are for the tunnel?		
For each circuit:	Local IP address?	
	Remote IP address?	
	Local HA address?	
	If eHCL is supported?	
	Remote HA address?	
	If eHCL is supported?	
	Minimum bandwidth rate (ARL)?	
	Maximum bandwidth rate (ARL)?	
Existing or new application type? (If existing, skip this section)		
New application type	App-type name and description	
	Port or port range	
Existing or new TCL? (If existing, skip to the next segment)		
TCL	Rule priority number?	
	Match criteria?	
	Source subnet or IP address?	
	Destination subnet or IP address?	
	VLAN tag? LAN interface or portchannel has been configured for VLAN tagging?	
	Application or protocol?	
	QoS (DSCP or L2CoS)?	
	Layer 4 protocol? TCP/UDP/ICMP/ICMP6?	
Traffic handling	Action (allow or deny)?	

Required Information		Your Information
	Target? allow: requires a specified VE deny: default is “all” or specifies a particular DP	
	Reset propagation Enable or disable?	
	Segment Preservation Enable or disable?	
	Non-terminated Traffic Enable or disable?	
IPsec policy: PreShared Key (PSK) or Public Key Infrastructure (PKI)?		
IPsec PSK, random 32 char alpha-numeric sequence or not applicable? No need to record or save the PSK.		
IPsec PKI requires prior installed certificates or not applicable?		
IPsec PKI key-pair name? This value is your human readable name.		

Configurations Supported by ECM

The following tables describe the objects that are configurable through ECM. Not all parameters are configurable for every object. Object parameters that are not available through ECM must be configured outside of the ECM utility. For more information about configurations that are not supported by ECM, see Configurations Not Covered by ECM.

The following table describes the System Configuration objects that you can configure through ECM.

Table 36: System Configurations

System Configurations	CLI Command	Menu/Submenu
App-Mode (Brocade SX6 Extension Blade and Brocade 7850 Extension Switch only)	`extncfg --app-mode`	System menu
GE-Mode (Brocade 7810 Extension Switch only)	`extncfg --ge-mode`	System menu
Extension Config Default/Clear	`extncfg --config -default` `extncfg --config -clear`	System menu
WAN side IPsec Policy	`portcfg ipsec-policy`	Security menu
IP Extension TCL	`portcfg tcl`	Traffic Profile menu
IP Extension App-Type	`portcfg app-type`	Traffic Profile menu

The following table describes the GE Interface Configuration objects that you can configure through ECM.

Table 37: GE Interface Configuration

Port Configurations	CLI Command	Menu/Submenu
GE Interface LAN/WAN Mode	`portcfgge --set -lan | -wan`	GE Port menu (WAN is the default)
GE Interface Speed	`portcfgge --set -speed`	GE Port menu (SFP+ ports 10GE is the default)

Port Configurations	CLI Command	Menu/Submenu
GE Interface Autoneg (1GbE only)	`portcfgge --enable | --disable -autoneg`	GE Port menu (Enabled is the default)

The following table describes the IP Configuration objects that you can configure through ECM.

Table 38: IP Configuration Parameters

Port Configuration	CLI Command	Menu/Submenu
WAN side IP Address Create	`portcfg ipif <ge#.dp#> create` `portcfg ipif <slot#/ge#.dp#> create`	IP menu or Circuit menu
WAN side IP Address Delete	`portcfg ipif <ge#.dp#> delete` `portcfg ipif <slot#/ge#.dp#> delete`	IP menu or Circuit menu
LAN side IP Extension Gateway Create	`portcfg ipif <lan.dp#> create` `portcfg ipif <slot#/lan.dp#> create`	IP menu or Circuit menu
LAN side IP Extension Gateway Delete	`portcfg ipif <lan.dp#> delete` `portcfg ipif <slot#/lan.dp#> delete`	IP menu or Circuit menu
IPIF MTU	`portcfg ipif ... mtu <Number_of_Bytes>`	IP menu or Circuit menu
IPIF VLAN	`portcfg ipif ... vlan <VLAN_ID_Number>`	IP menu or Circuit menu
WAN side IP Route Create	`portcfg iproute <ge#.dp#> create portcfg iproute <slot#/ge#.dp#> create`	IP menu or Circuit menu
LAN side IP Route Create	`portcfg iproute <lan.dp#> create portcfg iproute <slot#/lan.dp#> create`	IP menu or Circuit menu
WAN side IP Route Delete	`portcfg iproute <ge#.dp#> delete portcfg iproute <slot#/ge#.dp#> delete`	IP menu or Circuit menu
LAN side IP Route Delete	`portcfg iproute <lan.dp#> delete portcfg iproute <slot#/lan.dp#> delete`	IP menu or Circuit menu

The following table describes the Tunnel Configuration objects that you can configure through ECM.

Table 39: Tunnel Configuration Parameters

Tunnel Configuration	CLI Command	Menu/Submenu
Tunnel Create	`portcfg fciptunnel <ve> create`	Tunnel menu
Tunnel Delete	`portcfg fciptunnel <ve> delete`	Tunnel menu
Tunnel IP Extension	`portcfg fciptunnel <ve> --ipext <enable|disable>`	Tunnel menu

Tunnel Configuration	CLI Command	Menu/Submenu
Compression	`portcfg fciptunnel <ve> --compression <algorithm>`	Tunnel/Compression menu
FC-Compression	`portcfg fciptunnel <ve> --fc-compression <algorithm>`	Tunnel/Compression menu
IP-Compression	`portcfg fciptunnel <ve> --ip-compression <algorithm>`	Tunnel/Compression menu
Description	`portcfg fciptunnel <ve> --description <description>`	Tunnel menu
IPsec	`portcfg fciptunnel <ve> --ipsec <IPsec_Policy_Name>`	Tunnel menu

The following table describes the Circuit Configuration objects that you can configure through ECM.

Table 40: Circuit Configuration Parameters

Circuit Configuration	CLI Command	Menu/Submenu
Circuit Create	`portcfg fcipcircuit <ve> create <cir#>`	Circuit menu
Circuit Delete	`portcfg fcipcircuit <ve> delete <cir#>`	Circuit menu
Circuit Admin State	`portcfg fcipcircuit <ve> modify <cir#> --admin-status`	Circuit menu
Local IP Address	`portcfg fcipcircuit <ve> modify <cir#> --local-ip`	Circuit menu
Remote IP Address	`portcfg fcipcircuit <ve> modify <cir#> --remote-ip`	Circuit menu
Local HA IP Address	`portcfg fcipcircuit <ve> modify <cir#> --local-ha-ip`	Circuit menu
Remote HA IP Address	`portcfg fcipcircuit <ve> modify <cir#> --remote-ha-ip`	Circuit menu
Min/Max Rates	`portcfg fcipcircuit <ve> modify <cir#> --min <min_rate_in_kbps>` `portcfg fcipcircuit <ve> modify <cir#> --max <max_rate_in_kbps>`	Circuit menu

The following table describes the Verification and Troubleshooting Configuration objects that you can configure through ECM.

Table 41: Verification and Troubleshooting Configuration Parameters

WAN Tools	CLI Command	Menu/Submenu
WAN side Ping	`portcmd --ping <ge#.dp#>` `portcmd --ping <slot#/ge#.dp#>`	IP menu
LAN side Ping	`portcmd --ping <lan.dp#>` `portcmd --ping <slot#/lan.dp#>`	IP menu

Configurations Not Covered by ECM

Some advanced configuration tasks might not be covered by ECM because it is primarily intended to expedite and simplify the implementation. The following advanced configuration tasks can still be performed using the CLI outside of the ECM utility:

- **General Fabric OS configuration:**
  - DNS, NTP, SNMP, and IP filter
  - Firmware updates
  - FTrace
  - Zoning
- **Port configuration:**
  - GE SPF Channel ID
  - Portchannel (LAG/LACP)
  - LLDP
- **Tunnel configuration:**
  - FICON related configuration
  - FastWrite
  - OSTP
  - FCIP:IP Extension Distribution
  - FCIP H/M/L Distribution
  - IP Extension H/M/L Distribution
  - Legacy QoS Mode
  - Load Leveling Algorithm (spillover)
- **Circuit configuration:**
  - Failover metrics
  - Failover groups
  - ARL algorithm
  - SLA
  - NAT connection type
  - KAToV
  - DSCP marking
  - L2CoS (802.1P) marking
- **WAN Tools:**
  - WTool
  - Traceroute
  - PMTU discovery

Compatibility and Other Considerations

The following table lists the factors that you must consider when you use ECM.

Table 42: ECM Compatibility Considerations

Factors	Reason for Concern?
Backward Compatibility	ECM only provides configuration support for existing options, so backward compatibility is not an issue. Configurations that are created or managed using ECM use only the features and functionality of the Fabric OS installed on the platform.
Limitations	Not all configuration parameters are available through the ECM menus. The CLI might still be used to set features such as portchannel (LAG/LACP) or FICON. For more details about which parameters are included and which are not, see Configurations Supported by ECM and Configurations Not Covered by ECM.
High Availability	Any HA operations are conducted using the existing HA infrastructure, and configurations are generated in real-time. All configuration requests are made at the time that a parameter is entered, except for some parameters requiring multiple inputs. For example, when you enter the tunnel QoS ratio (high/medium/low), a minimum of three parameters are required before any take effect. If only two parameters are entered and the configuration is not completed, those two parameters will not persist when the platform is rebooted. After an HA event, you can re-enter ECM and can resume building the configuration. The context of the menu is rebuilt, so you only must navigate to the last submenu you used. However, ECM does not maintain a menu context across an unexpected event, such as a switch panic.
Upgrade and Downgrade	ECM configurations are only based on Fabric OS capabilities. There are no upgrade or downgrade considerations to be considered. Fabric OS upgrades and downgrades might affect any features that were introduced or removed from a specific release. Fabric OS configurations that are constructed using ECM work with the version where they were created. Fabric OS version limitations apply to the features themselves, not to the configuration utility that originally configured them.

Using ECM

ECM consists of a new CLI interface that uses interactive menus to perform various configuration options. The interface reads standard input from the user and exits to the Fabric OS shell environment upon terminating the utility.

ECM is started using the `extnCfg --config -manager` command.

```
ExtnPlatform:admin> extncfg --help Usage: extncfg <action> [options]

action:
  --ve-mode 10VE|20VE                   - Set VE-Mode to 10VE or 20VE mode. (SX6 only)
  --ve-mode { 6VE | 18VE }              - Set VE-Mode to 6VE or 18VE. (7850 only
  --ge-mode copper|optical              - Set GE-Mode to copper or optical. (7810 only)
  --app-mode fcip|hybrid                - Set APP-Mode to FCIP or HYBRID (FCIP with IPEXT - SX6 only).
  --show                                - Display APP & VE mode details.
  --config -default                     - Default the Extension configuration.
  --config -manager                     - Run the extension configuration management utility.
  --fwdl-prep [-version #.#.#] [-abort] - Prepare the switch for firmware download to the target version.
  --help                                - Print this usage statement.
```

ECM Menu Hierarchy

The following procedure shows the menu hierarchy and demonstrates the general use of the ECM menu:

NOTE
Root menu means the top-level menu. On the switch, the menu is not associated with any user account, such as root, admin, or user.

1. Root menu

```
--------------------------------------------------------------------------------
Extension Configuration Management
--------------------------------------------------------------------------------
0    : System Configurations
1    : Security Configurations
2    : Port Configurations
3    : IP Configurations
4    : Tunnel Configurations
5    : Traffic Profile Configurations
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
```

   If you select 0, go to step 2 or 3 depending on the platform.
   If you select 1, go to step 9.
   If you select 2, go to step 14.
   If you select 3, go to step 16.
   If you select 4, go to step 25.
   If you select 5, go to step 34.

2. Root/System Configurations (Brocade 7810 Extension Switch or Brocade 7850 Extension Switch)

```
--------------------------------------------------------------------------------
System Configurations
--------------------------------------------------------------------------------
0   : 7850 Switch - AppMode:HYBRID VE-Mode:6VE
1   : Clear Extension Configurations
--------------------------------------------------------------------------------
<Enter>=Select <ESC>=Back
```

   If you select 1, go to step 4.

3. Root/System Configurations (Brocade SX6 Extension Blade)

```
--------------------------------------------------------------------------------
System Configurations
--------------------------------------------------------------------------------
0    : Slot 3 Mode Configurations [App-Mode:FCIP|Hybrid (FCIP + IPExt)]
1    : Clear Extension Configuration
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
```

   If you select 1, go to step 5.

4. Root/System/Clear Extension Configuration (Brocade 7810 Extension Switch and Brocade 7850 Extension Switch)
   Use with caution as this menu completely clears all extension-related configurations.

```
--------------------------------------------------------------------------------
```

```
This operation will clear all extension configurations and reset all extension
blades into the default state. This operation cannot be undone.
--------------------------------------------------------------------------------
Do you with to proceed?

<Enter> to confirm
<ESC> to go back
--------------------------------------------------------------------------------
```

5. Root/System/Clear Extension Configuration (Brocade SX6 Extension Blade)

   The dynamically built display, lists all available extension-based slots, including enabled and powered-off slots. When selected, the extension configurations for the specified slot, or all slots, are cleared. Any powered-on blades might be reset during this process. Wait for the slot to fully boot before entering ECM again. Any powered-off blades remain powered off.

```
--------------------------------------------------------------------------------
Clear Extension Configuration
--------------------------------------------------------------------------------
0    : SX6 Slot 7 - AppMode:HYBRID, VE-Mode:10VE
1    : SX6 Slot 8 - (Inserted)
2    : <all>
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
```

   Selection 0 – Blade Powered-On, go to step 6. Clear all extension-related configurations on a slot. Any global configurations such as IPsec policies, flow definitions remain.

   Selection 1 – Blade Powered-Off, go to step 7. Clear all extension-related configurations on a slot. Any global configurations such as IPsec policies, flow-definitions, and so on remain.

   Selection 2 – All Blades, go to step 8. where you can clear all extension-related configurations on the chassis.

6. Root/System/Clear Extension Configuration/Slot (Brocade SX6 Extension Blade - blade powered-on)

```
--------------------------------------------------------------------------------
This operation will clear all extension configurations from all logical
switches for slot 7 and reset the blade into the default state. All switch
wide configurations will remain unaffected. This operation cannot be undone.
The configuration manager will be terminated once complete. Please wait for
system to be ready again before entering the configuration manager.

Do you wish to proceed?
--------------------------------------------------------------------------------
<ENTER> to confirm
<ESC> to go back
--------------------------------------------------------------------------------
```

7. Root/System/Clear Extension Configuration/Slot (Brocade SX6 Extension Blade - blade powered-off)

```
--------------------------------------------------------------------------------
This operation will clear all extension configurations from all logical
switches for slot 8 and leave the blade disabled. All switch wide
configurations will remain unaffected. This operation cannot be undone. The
configuration manager will be terminated once complete. Please wait for system
to be ready again before entering the configuration manager.

Do you wish to proceed?
--------------------------------------------------------------------------------
```

```
<ENTER> to confirm
<ESC> to go back
--------------------------------------------------------------------------------
```

8. Root/System/System Configuration/All (Brocade SX6 Extension Blade - all blades)

```
--------------------------------------------------------------------------------
This operation will clear all extension configurations from all logical
switches and reset all extension blades into the default state. This operation
cannot be undone. The configuration manager will be terminated once complete.
Please wait for system to be ready again before entering the configuration
manager.

Do you wish to proceed?
--------------------------------------------------------------------------------
<ENTER> to confirm
<ESC> to go back
--------------------------------------------------------------------------------
```

9. Root/Security Configurations

```
Security Configurations
--------------------------------------------------------------------------------
0    : IPSec Policy Configurations
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
Security Configurations
--------------------------------------------------------------------------------
0    : IPSec Policy Configurations
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
```

10. Root/Security/IPsec Policy Configurations

```
--------------------------------------------------------------------------------
IPSec Policy Configurations
--------------------------------------------------------------------------------
0    : myPskPolicy
1    : <new>
2    : <delete>
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
```

Selection 0 – go to step 13, depending on the type of security (PSK or PKI), one of the following two outputs is shown:
Selection 1 – go to step 11.
Selection 2 – go to step 12.

11. Root/Security/IPsec Policies/<new>

```
--------------------------------------------------------------------------------
<new>
--------------------------------------------------------------------------------
Please enter a name for the new IPSec policy:
```

12. Root/Security/IPsec Policies/<delete>

```
--------------------------------------------------------------------------------
DELETE IPSec Policy - This cannot be undone
--------------------------------------------------------------------------------
0    : myPskPolicy
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
```

13. Root/Security/IPsec Policies/<policy>

    The configuration menu for the IPsec policy instance is built dynamically based on the selected profile. For a PSK (preshared key) profile, only the PSK value is configurable. For a PKI (Public Key Infrastructure) profile, only the key-pair Name is configurable. When selecting the key-pair name for a PKI policy, a list of available key-pairs displays. You can select only from this list. New key-pair creation is not supported through ECM.

    a. Preshared Key

    ```
    --------------------------------------------------------------------------------
    IPSec Policy: myPskPolicy
    --------------------------------------------------------------------------------
    0    : Profile                [Preshared Key]
    1    : Preshared Key          [     ]
    2    : Restart authentication
    --------------------------------------------------------------------------------
    <Enter>=Select    <ESC>=Back

    Please select menu option:
    ```

    b. PKI

    ```
    --------------------------------------------------------------------------------
    IPSec Policy: myPskPolicy
    --------------------------------------------------------------------------------
    0    : Profile                [PKI]
    1    : Key-Pair Name          [     ]
    2    : Restart authentication
    --------------------------------------------------------------------------------
    <Enter>=Select    <ESC>=Back

    Please select menu option:
    ```

14. Root/Port Configurations

    The following example applies to Brocade 7810 Extension Switch and Brocade 7850 Extension Switch. Brocade SX6 Extension Blade is similar except for the additional slot information. The following menu lists the GE interfaces available in the default and current logical switch. In case of a logical switch, the FID info is displayed.

    ```
    --------------------------------------------------------------------------------
    Port Configurations
    --------------------------------------------------------------------------------
    0    : GE-Port ge2 10G / WAN
    1    : GE-Port ge3 10G / LAG (MyLAGname)
    2    : GE-Port ge4 10G / WAN
    3    : GE-Port ge5 10G / WAN
    4    : GE-Port ge6 10G / LAN
    5    : GE-Port ge7 10G / WAN
    --------------------------------------------------------------------------------
    <Enter>=Select    <ESC>=Back
    ```

```
Please select menu option:
```

15. Root/Port/Port#

```
--------------------------------------------------------------------------------
GE-Port # Configuration
--------------------------------------------------------------------------------
0    : Speed              [1G|10G|25G|40G|100G]
1    : Auto-negotiate     [Enabled|Disabled]
2    : Port Mode          [WAN|LAN]
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
```

16. Root/IP Configurations

```
--------------------------------------------------------------------------------
IP Configurations
--------------------------------------------------------------------------------
0    : IP Interfaces
1    : IP Routes
2    : Ping
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
```

17. Root/IP/IP Interface

This dynamically generated menu lists all configured IP interfaces. All IP interfaces of the default and current logical switches are included in this configuration.

```
--------------------------------------------------------------------------------
IP Interface Configurations
--------------------------------------------------------------------------------
0    : ge0.dp0    192.168.0.1 / 24    MTU:1500    VLAN:0
1    : ge0.dp0    192.168.0.2 / 24    MTU:1500    VLAN:0
2    : ge2.dp0    2002:2::10  / 64    MTU:1500    VLAN:100
3    : lan.dp0    10.0.0.1    / 24    MTU:1500    VLAN:0
4    : <new>
5    : <delete>
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
```

Selection 5, a list is presented of available interfaces to delete.

18. Root/IP/IP Interface/<new>

This option requests the arguments that are required to uniquely identify an IP interface. It triggers a cascade of menu options where entering each option brings you to the next menu.

NOTE

- The DP selection screen only exists on platforms with multiple DPs such as Brocade 7850 Extension Switch and Brocade SX6 Extension Blade.
- When entering the IP address, if a CIDR notation is entered, for example: 10.0.0.100/24), then the subsequent *subnet mask*/*prefix length* prompt is omitted.
- If using the default VLAN and MTU settings, use ESC to back out. If using the default VLAN and MTU settings, use ESC to back out.

Example input shown.

```
--------------------------------------------------------------------------------
<new>
--------------------------------------------------------------------------------
Please choose an available GE port

0    : GE-Port  ge2   10G  / WAN
1    : GE-Port  ge3   10G  / WAN
2    : GE-Port  ge4   1G   / WAN
3    : GE-Port  ge5   10G  / WAN
4    : SVI-Port lan   ---  / LAN
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option: 0

--------------------------------------------------------------------------------
<new>
--------------------------------------------------------------------------------
Please choose an available DP

0    : DP0
1    : DP1
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option: 0

--------------------------------------------------------------------------------
Create IP Address: ge2.dp0

IP Address
--------------------------------------------------------------------------------
Please enter an IP Address: 10.10.10.10

--------------------------------------------------------------------------------
Create IP Address: ge2.dp0    10.10.10.10

Subnet Mask / Prefix length:
--------------------------------------------------------------------------------
Please enter a Netmask / Prefix length: 24

--------------------------------------------------------------------------------
 ge2.dp0     10.10.10.10                     / 24    MTU:1500  VLAN:0
--------------------------------------------------------------------------------
 0  : VLAN                    [none]
 1  : MTU                     [1500]
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

19. Root/IP/IP Interface/<delete>

This option provides a dynamic list of all available IP interfaces.

```
--------------------------------------------------------------------------------
DELETE IP Interface - This cannot be undone
--------------------------------------------------------------------------------
0    : ge0.dp0     192.168.0.1           / 24     MTU:1500     VLAN:0
1    : ge0.dp0     192.168.0.2           / 24     MTU:1500    VLAN:0
2    : ge2.dp0     2002:2::10            / 64     MTU:1500    VLAN:100
3    : lan.dp0     10.0.0.1              / 24     MTU:1500    VLAN:0
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select a menu option:
```

20. Root/IP/IP Interface /<ip-interface>

    This menu displays after you select an IP interface or after entering the subnet mask/prefix for a new IP interface.

```
--------------------------------------------------------------------------------
 ge0.dp0     192.168.0.1                  / 24    MTU:1500  VLAN:0
--------------------------------------------------------------------------------
 0  : VLAN                     [none]
 1  : MTU                      [1500]
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

21. Root/IP/IP Route

    The menu provides a dynamic list of all statically configured IP routes.

```
--------------------------------------------------------------------------------
IP Route Configurations
--------------------------------------------------------------------------------
0    : ge0.dp0    192.168.2.1 / 24    192.168.0.250
1    : <new>
2    : <delete>
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Exit

Please select a menu option:
```

    Selection 2 displays a list of available statically configured IP routes to delete.

22. Root/IP/IP Route/<new >

    DP selection is not prompted for Brocade 7810 Extension Switch because it only has DP0. DP selection is prompted for Brocade 7850 Extension Switch and Brocade SX6 Extension Blade. When providing the destination IP address, if CIDR (/prefix length) notation is entered, the subsequent subnet mask/prefix length prompt is omitted. An IP route cannot be modified once it is created; therefore, the menu returns to the Root/IP/IP Route menu.

    Example entries are shown.

```
--------------------------------------------------------------------------------
<new>
--------------------------------------------------------------------------------
Please choose an available GE port
0    : GE-Port 0   10G     / WAN
1    : GE-Port 1   10G     / WAN
2    : GE-Port 2   1G      / WAN
3    : GE-Port 3   10G     / WAN
4    : GE-Port lan         / LAN
```

```
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option: 2

--------------------------------------------------------------------------------
Please choose an available DP
--------------------------------------------------------------------------------
0    : DP0
1    : DP1
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option: 0

--------------------------------------------------------------------------------
 IP Route Create: ge2.dp0

 Configured IP Interfaces:Interface Addresses:
   ge2.dp0 192.168.10.1 / 30
   ge2.dp0 10.0.0.5 / 24

 Destination IP Address
--------------------------------------------------------------------------------
 Please enter a Destination IP Address or network: 10.10.10.0/24

--------------------------------------------------------------------------------
 IP Route Create: ge2.dp0     10.10.10.0                     / 24

 Configured IP Interfaces:Interface Addresses:
   ge2.dp0 192.168.10.1 / 30
   ge2.dp0 10.0.0.5 / 24

 Gateway
--------------------------------------------------------------------------------
 Please enter a Gateway IP Address: 10.0.0.1
```

23. Root/IP/IP Route/<delete>

```
--------------------------------------------------------------------------------
DELETE IP Route - This cannot be undone
--------------------------------------------------------------------------------
0 : ge0.dp0    192.168.2.1 / 24    192.168.0.250
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option:
```

24. Root/IP/Ping

    This menu lists all available IP addresses (already configured) and asks you to select a source address. Then, the ping utility requests the destination address and initiates the ping request using all default values.

```
--------------------------------------------------------------------------------
Source IP Address
--------------------------------------------------------------------------------
0 : ge0.dp0    192.168.0.10                      / 24      MTU:1500   VLAN:0
```

```
1 : ge0.dp0    192.168.0.20                    / 24      MTU:1500   VLAN:0
2 : ge2.dp0    2002:2::10                      / 64      MTU:1500   VLAN:100
3 : lan.dp0    10.0.0.10                       / 24      MTU:1500   VLAN:0
--------------------------------------------------------------------------------
<Enter>=Select    <ESC>=Back

Please select menu option: 3

--------------------------------------------------------------------------------
Destination IP Address
--------------------------------------------------------------------------------
Please enter a destination address: 10.0.0.1

<ping command output>
```

25. Root/Tunnel

    If a tunnel or circuit has an incomplete configuration, the items remaining to be configured are listed. Incomplete tunnel and circuit configuration is a normal part of the configuration staging method.

```
--------------------------------------------------------------------------------
 Tunnel Configurations

 WARNING:
 * Circuit 12.1: missing LocalIP RemoteIP MinCommRate MaxCommRate
   configuration.
--------------------------------------------------------------------------------
*0  : Tunnel: 12
 1  : <new>
 2  : <delete>
```

    Selection 1 – go to step 26.

    Selection 2 – go to step 27.

26. Root/Tunnel/<new>

    This menu option lists the available VE ports. An example entry is shown.

```
--------------------------------------------------------------------------------
 VE Port
--------------------------------------------------------------------------------
 0  : VE-Port 13
 1  : VE-Port 14
 2  : VE-Port 15
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option: 0
```

27. Root/Tunnel/<delete>

```
--------------------------------------------------------------------------------
 DELETE Tunnel - This cannot be undone
--------------------------------------------------------------------------------
 0  : Tunnel: 12
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

28. Root/Tunnel/<tunnel>

The tunnel configuration menu dynamically changes based on the type field.

– If the tunnel is configured for FCIP, then the compression field appears as only a single field.
– If the tunnel is configured for IP Extension, then the compression option presents as a submenu to configure the various modes (see step 29).

```
--------------------------------------------------------------------------------
 Tunnel: 13

 WARNING:
 * Tunnel: 13: No circuits configured.
--------------------------------------------------------------------------------
 0  : Description             [                        ]
 1  : Tunnel Type             [FCIP]
 2  : Compression             [None]
 3  : IP-Sec Policy           [                        ]
 4  : Circuit Configurations
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option: 1

--------------------------------------------------------------------------------
 Tunnel Type

 Modification of this parameter may be disruptive.
--------------------------------------------------------------------------------
 0  : FCIP
 1  : Hybrd (FCIP + IPExt)
--------------------------------------------------------------------------------
 <Enter>=Select  <ESC>=Cancel

Input: 1
```

29. Root/Tunnel/<tunnel>/Compression

This menu is reachable only when the tunnel is configured for IP Extension.

NOTE

Fast-Deflate is not supported with IP Extension on any platform, or on Brocade 7810 Extension Switch.

```
--------------------------------------------------------------------------------
 Compression

 Modification of this parameter may be disruptive.
--------------------------------------------------------------------------------
 0  : Main Compression        [None]
 1  : FC-Compression          [Default (None)]
 2  : IP-Compression          [Default (None)]
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

30. Root/Tunnel/<tunnel>/Circuit

```
--------------------------------------------------------------------------------
Circuit Configurations
--------------------------------------------------------------------------------
```

```
0    : Circuit <ve>.0 (Incomplete)
1    : Circuit <ve>.1 (Offline)
2    : <new>
3    : <delete>
--------------------------------------------------------------------------------
 <Enter>=Select    <ESC>=Exit

Please select menu option:
```

Selection 2 – you see the following menu, Step 31.

Selection 3 – you see the available circuits to delete, Step 32.

31. Root/Tunnel/<tunnel>/Circuit/<new>

```
--------------------------------------------------------------------------------
 Circuit ID
--------------------------------------------------------------------------------
 Please enter a circuit ID for the new circuit:
```

32. Root/Tunnel/<tunnel>/Circuit/<delete>

```
--------------------------------------------------------------------------------
DELETE Circuit - This cannot be undone
--------------------------------------------------------------------------------
0    : Circuit <ve>.0 (Incomplete)
1    : Circuit <ve>.1 (Offline)
--------------------------------------------------------------------------------
 <Enter>=Select    <ESC>=Exit

Please select menu option:
```

33. Root/Tunnel/<tunnel>/<circuit>

At the circuit configuration menu, once one of the two local IP address fields is selected, a list of configured IP addresses available for the VE, based on the DP, is presented. VE_Ports and IPIFs live on a specific DP; therefore, ECM shows only IP addresses a VE can use based on the association the IP addresses to DP. When you enter one of the two remote IP address fields, you are prompted for one of those addresses.

Once you enter a required address, the utility performs an IP route validation check. If a corresponding IP route has not been configured, you are prompted to enter the gateway IP address. An IP route is automatically generated.

NOTE

For Brocade 7810 Extension Switch, only the Remote-HA is configurable, the Local-HA is not shown. The Local-HA-IP and Remote-HA-IP fields are visible on Brocade 7850 Extension Switch and Brocade SX6 Extension Blade, which supports eHCL.

```
--------------------------------------------------------------------------------
 Circuit 13.0

 WARNING:
 * Circuit 13.0: missing LocalIP RemoteIP MinCommRate MaxCommRate
   configuration.
--------------------------------------------------------------------------------
 0  : Admin Status             [Enable]
 1  : Local IP Address         [0.0.0.0]
 2  : Remote IP Address        [0.0.0.0]
 3  : Remote HA IP Address     [0.0.0.0]
 4  : Minimum Committed Rate   [0]
 5  : Maximum Committed Rate   [0]
 6  : Test Connection
--------------------------------------------------------------------------------
```

```
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

34. Root/Traffic Profile

    Traffic Profile is specific to IP Extension. If a tunnel exists that has IP Extension enabled, a warning is shown if there is no TCL allow rules to match traffic into the tunnel.

```
--------------------------------------------------------------------------------
 Traffic Profile Configurations

 WARNING:
 * Tunnel: 13: has IPExt enabled but no TCLs targeting it.
--------------------------------------------------------------------------------
 0  : App-Type Configurations
*1  : TCL Configurations
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

    Selection 0 – go to step 35.

    Selection 1 – go to step 39.

35. Root/Traffic Profile/App-Type

```
--------------------------------------------------------------------------------
 App-Type Configurations - S=Statically Defined (cannot be modified / deleted).
--------------------------------------------------------------------------------
 0  : CIFS                    S  139,445
 1  : Data-Domain             S  2051
 2  : FCIP                    S  3225-3226
 3  : FTP                     S  20-21,989-990,115
 4  : HTTP                    S  80,8080,8000-8001,3128
 5  : HTTPS                   S  443
 6  : Isilon-SyncIQ           S  5666-5667
 7  : LDAP                    S  389,8404,636
 8  : MS-SQL                  S  1443
 9  : MySQL                   S  3306
 10 : NETAPP-SNAP-MIRROR      S  10566
 11 : NFS                     S  2049
 12 : ORACLE-SQL              S  66,1525,1521
 13 : RSYNC                   S  873
 14 : SRDF                    S  1748
 15 : SSH                     S  22
 16 : SSL-SHELL               S  614
 17 : TELNET                  S  23,107,513,992
 18 : TFTP                    S  69
 19 : VERITAS-BACKUP          S  6101-6102,6106,3527,1125
 20 : VTS-GRID Control        S  1415-1416
 21 : VTS-GRID Data           S  350
 22 : iSCSI                   S  3260
 23 : <new>
 24 : <delete>
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back
```

```
Please select menu option:
```

36. Root/Traffic Profile/App-Type/<new>

```
--------------------------------------------------------------------------------
 <new>
--------------------------------------------------------------------------------
 Please enter a name for the new App-Type: MyApp


--------------------------------------------------------------------------------
 MyApp                    -
--------------------------------------------------------------------------------
 0  : Description              [                         ]
 1  : Port Range               [                         ]
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option: 0

--------------------------------------------------------------------------------
 Description
--------------------------------------------------------------------------------
 Please enter a description: MyEnterpise Application


--------------------------------------------------------------------------------
 MyApp                    -
--------------------------------------------------------------------------------
 0  : Description              [MyEnterpise Application]
 1  : Port Range               [                         ]
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option: 1

--------------------------------------------------------------------------------
 Port Range - Examples: 100   100-200   100,200   100,200-300
--------------------------------------------------------------------------------
 Please enter a port range value: 2345


--------------------------------------------------------------------------------
 MyApp                    -  2345
--------------------------------------------------------------------------------
 0  : Description              [MyEnterpise Application]
 1  : Port Range               [2345]
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

37. Root/Traffic Profile/App-Type/<delete>

    Only the user added App-Types are listed.

```
--------------------------------------------------------------------------------
```

```
 DELETE App-Type - This cannot be undone
--------------------------------------------------------------------------------
 0  : MyApp                    -  2345
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

38. Root/Traffic Profile/App-Type/<app-type>

```
--------------------------------------------------------------------------------
 App-Type Configurations - S=Statically Defined (cannot be modified / deleted).
--------------------------------------------------------------------------------
 0  : CIFS                     S  139,445
 1  : Data-Domain              S  2051
 2  : FCIP                     S  3225-3226
 3  : FTP                      S  20-21,989-990,115
 4  : HTTP                     S  80,8080,8000-8001,3128
 5  : HTTPS                    S  443
 6  : Isilon-SyncIQ            S  5666-5667
 7  : LDAP                     S  389,8404,636
 8  : MS-SQL                   S  1443
 9  : MyApp                    -  2345
 10 : MySQL                    S  3306
 11 : NETAPP-SNAP-MIRROR       S  10566
 12 : NFS                      S  2049
 13 : ORACLE-SQL               S  66,1525,1521
 14 : RSYNC                    S  873
 15 : SRDF                     S  1748
 16 : SSH                      S  22
 17 : SSL-SHELL                S  614
 18 : TELNET                   S  23,107,513,992
 19 : TFTP                     S  69
 20 : VERITAS-BACKUP           S  6101-6102,6106,3527,1125
 21 : VTS-GRID Control         S  1415-1416
 22 : VTS-GRID Data            S  350
 23 : iSCSI                    S  3260
 24 : <new>
 25 : <delete>
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option: 9

--------------------------------------------------------------------------------
 MyApp                    -  2345
--------------------------------------------------------------------------------
 0  : Description              [MyEnterpise Application]
 1  : Port Range               [2345]
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

39. Root/Traffic Profile/TCL

```
--------------------------------------------------------------------------------
TCL Configurations
--------------------------------------------------------------------------------
0 : default
1 : <new>
2 : <delete>
--------------------------------------------------------------------------------
 <Enter>=Select    <ESC>=Back

Please select menu option:
```

40. Root/Traffic Profile/TCL/<new>

```
--------------------------------------------------------------------------------
<new>
--------------------------------------------------------------------------------
 Please enter a name for the new TCL:
```

41. Root/Traffic Profile/TCL/<delete>

```
--------------------------------------------------------------------------------
DELETE TCL - This cannot be undone
--------------------------------------------------------------------------------
0    : HostA_to_HostB
--------------------------------------------------------------------------------
 <Enter>=Select    <ESC>=Back

Please select menu option:
```

42. Root/Traffic Profile/TCL/<tcl>

```
--------------------------------------------------------------------------------
 <TCL_Name>
--------------------------------------------------------------------------------
 0  : Admin-Status             [Disable]
 1  : Priority                 [0]
 2  : Match Criteria
 3  : Traffic Handling
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

43. Root/Traffic Profile/TCL/<tcl>/Match Criteria

```
--------------------------------------------------------------------------------
 Match Criteria
--------------------------------------------------------------------------------
 0  : Source IP Address        [<any>]
 1  : Destination IP Address   [<any>]
 2  : VLAN                     [<any>]
 3  : Application / Protocol
 4  : QoS (DSCP / L2COS)
 5  : Layer 4 Protocol         [<any>]
--------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

44. Root/Traffic Profile/TCL/<tcl>/Application/Protocol

These options are mutually exclusive. Setting one of them (TCL: 0 or 1) automatically unsets the other.

If you choose the known app-type, you are prompted with an App-Type list. From there, you choose an existing app-type or create one.

```
 ------------------------------------------------------------------------------
  Application / Protocol
 ------------------------------------------------------------------------------
  0  : Known App-Type           [                          ]
  1  : Port Range               [<any>]
 ------------------------------------------------------------------------------
  <Enter>=Select   <ESC>=Back

 Please select menu option:
```

45. Root/Traffic Profile/TCL/<tcl>/QoS

```
 ------------------------------------------------------------------------------
  QoS (DSCP / L2COS)
 ------------------------------------------------------------------------------
  0  : DSCP Marking             [<any>]
  1  : L2CoS                    [<any>]
 ------------------------------------------------------------------------------
  <Enter>=Select   <ESC>=Back

 Please select menu option:
```

46. Root / Traffic Profile / TCL / <tcl> / Traffic Handling

For the traffic handling menu, the fields for reset propagation, segment preservation, and non-terminated traffic are visible only when Action is set to `Allow` .

The Target is a list of available VE ports when Action is set to `Allow` , and when the Action is set to `Deny` , it is where the TCL rule is enforced: DP0, DP1, or ALL.

```
 ------------------------------------------------------------------------------
  Traffic Handling
 ------------------------------------------------------------------------------
  0  : Action                   [Deny]
  1  : Target DP                [<all>]
 ------------------------------------------------------------------------------
  <Enter>=Select   <ESC>=Back

 Please select menu option: 0

 ------------------------------------------------------------------------------
  Action
 ------------------------------------------------------------------------------
  0  : Allow
  1  : Deny
 ------------------------------------------------------------------------------
  <Enter>=Select  <ESC>=Cancel

 Input: 0

 ------------------------------------------------------------------------------
  Traffic Handling
 ------------------------------------------------------------------------------
  0  : Action                   [Allow]
```

```
 1  : Target Tunnel            [                        ]
 2  : Reset Propagation        [Disable]
 3  : Segment Preservation     [Disable]
 4  : Non-Terminated Traffic   [Disable]
------------------------------------------------------------------------------
 <Enter>=Select   <ESC>=Back

Please select menu option:
```

IP Extension Flow Monitor

The IP Extension Flow Monitor allows you to monitor and view LAN side statistics.

By viewing the traffic statistics of an IP pair or a specific port type, you can determine if the IP Extension traffic flows are functioning as expected. IP Extension Flow Monitor is a separate feature from Brocade Flow Vision and Flow Monitoring. You can access IP Extension Flow Monitor through the Fabric OS CLI commands. The following table lists the types of IP traffic information that you can monitor.

NOTE
Because there is a less available memory on a Brocade 7810 Extension Switch, the maximum upper limit for the VE port-based monitoring of FC is lower. The maximum learned VE flows and static VE flows are 4 (1 per VE port) and 256, respectively. However, a Brocade 7810 Extension Switch provides similar support for IP Extension flow monitoring. The maximum supported flows (TCP connections, the values for the Brocade SX6 Extension Blade [4000 flows per DP for current and historical stats]) and the maximum supported IP-PAIR (32 IP pairs) are unchanged.

Table 43: IP Extension Flow Monitor IP Traffic

Layer	Monitored Information
L2	VLAN ID Non-IP traffic is not monitored
L3	Source and Destination IP addresses Non-TCP traffic is not monitored UDP traffic monitoring is not supported
L4	TCP port information based on the TCP port number TCP port information based on the TCP port range TCP port information based on the TCP application type Non-TCP traffic is not monitored UDP traffic monitoring is not supported

Monitoring Traffic Flows

The primary use of the IP Extension Flow Monitor is to display the statistical information about the flow of L2, L3, and L4 traffic on the ports or addresses. The type of traffic information that is displayed by the IP Extension Flow Monitor is controlled by user-defined flows. When you configure the flow monitor, you define a flow name and the type of flow that you want to monitor. Using flow names, you can define filters for the following categories:

- GE port
- VE port
- Slot (in chassis)
- DP
- IP address
- TCP port, by port number, port range, or application
- VLAN ID
- Byte values
- Retransmit values

Limited Boolean operations are allowed so that you can further refine the definition for a named flow. Up to 32 flow names can be created on a chassis or platform. You can view the dynamic information and historical information. The dynamic information is displayed when you enter the CLI command.

In the view of historical information, statistics are collected from the saved list of connections from the DP. The saved list includes only closed connections. When an active connection is closed, that connection is moved to the saved list. Statistics for closed connections include information about the reason for closure, the time of closure, and optional detailed information.

You can freeze and thaw the historical statistics. When you freeze historical data, that action freezes the saved list in the DP. Connections that are closed after the freeze are not added to the saved list. When you thaw the historical statistics, the saved list returns to its normal operating state, which means that closed connections are again added. During the freeze, no additional information is added to the statistics.

Monitoring IP Pairs

The second use of the IP Extension Flow Monitor is to view IP pair statistics. IP pair statistics are useful for tracking traffic between local and remote nodes. For example, you configure the storage replication traffic which flows over one or more TCP connections between the nodes and you want to see traffic bandwidth patterns.

Current IP pair information and historical IP pair information are available. You can view the historical IP pair information as a total of the past 24 hours that are captured once every 24-hour period. Up to seven consecutive days of information are retained. The day 1 data drops off when day 8 data is stored. The IP pair information is automatically configured and maintained for each TCP connection and stored when a TCP connection closes on the DP. With the historical information, you can monitor the traffic patterns for each day to analyze the replication bandwidth and view the transmitted and received data for each closed TCP connection during a 24-hour period.

When there is no TCP connection on an IP pair for seven days, the IP pair is removed from the list. New TCP connections initiate the creation of a new IP pair. Statistics are updated from the creation time and not necessarily from the time traffic started.

IP pair statistics can be reset with a CLI command. In addition, the following situations also reset IP pair statistics and result in a loss of current and historical information:

- Firmware upgrade or downgrade
- HCL operations
- DP reset

For more information about use and configuration, see Using IP Extension Flow Monitor.

Using IP Extension Flow Monitor

Using the IP Extension Flow Monitor feature, you can view LAN traffic statistics.

Before you can use IP Extension Flow Monitor, you must configure IP Extension on the platform. The IP Extension Flow Monitor is useful only when the IP traffic is flowing through the extension tunnel. To use the commands that display the IP Extension Flow Monitor statistics, you must connect to the switch and log on using an account assigned to the admin role. For information on all command options, refer to the *Brocade Fabric OS Command Reference Manual*.

NOTE
Non-Terminated-TCP (that is, Wtool) and UDP traffic is not included in the IP pair statistics.

IP Extension Flow Monitor provides default views for traffic flow statistics. You can define and name additional flows so that you can view a specific set of details. You can define up to 32 named flows on a platform.

- In addition to named flows, you can use IP Extension Flow Monitor to view statistics for IP pairs. An IP pair is established automatically for each TCP connection created between source and destination IP addresses. A historical

snapshot of the IP pair information is automatically created every 24 hours and retained for seven days. By using command options, you can filter or expand the information displayed for IP pairs.

The following steps show how to choose your default view for flows and IP pairs.

1. To view the default view of flow statistics, use the `portshow lan-stats --per-flow` command.

```
switch:admin> portshow lan-stats --per-flow

*** Displaying Top 5 connections by throughput ***

DP    Idx  Src-Address      Dst-Address     Sport Dport Pro   Tx(B/s) Rx(B/s)
-------------------------------------------------------------------------------
DP0   142  192.168.10.76    192.168.20.38   63040 49883 TCP   6.6m    6.6m
DP0   120  192.168.10.76    192.168.20.38   63018 49883 TCP   6.6m    6.6m
DP0   164  192.168.10.76    192.168.20.38   63062 49883 TCP   6.6m    6.6m
DP0   170  192.168.10.76    192.168.20.38   63068 49883 TCP   6.6m    6.6m
DP0   150  192.168.10.76    192.168.20.38   63048 49883 TCP   6.6m    6.6m
-------------------------------------------------------------------------------
Sport=Source-Port Dport=Destination-Port Pro=Protocol


DP   ActTCP  ExdTCP  TCLDeny TCLFail
-----------------------------------------------------
DP0  101     0       0       0
-----------------------------------------------------
ActTCP=Active TCP Conns ExdTCP=Exceeded TCP Conn Cnt
```

   The default display shows up to the top 25 connections sorted by throughput. Each connection is identified by a unique index number as shown in the Idx column.

2. To view the default view of IP pair statistics, use the `portshow lan-stats --ip-pair` command.

```
switch:admin> portshow lan-stats --ip-pair

DP       Idx SrcAddr           DstAddr           Active   TxB       RxB
-------------------------------------------------------------------------------
DP0      1   192.168.20.38     192.168.10.76    101      13.2t        13.2t
-------------------------------------------------------------------------------
```

   The display shows summary information for all IP pairs. In the example output, DP0 has 101 active TCP connections between the source IP address and the destination IP address. DP1 has no IP pair connections, so no information is displayed.

3. Select and filter statistics by using available options for the `portshow lan-stats` command without creating flow names. To display all available options, use the `portshow lan-stats --help` command.

```
switch:admin> portshow lan-stats --help[Output is not shown]
```

For information on all command options, refer to the *Brocade Fabric OS Command Reference Manual*.

Configuring a Port-Based Flow

You can create a named flow that filters the IP Extension Flow Monitor output based on a specific port, such as a VE_Port or GE port.

NOTE
The flow name is case-sensitive. GE3 is not the same as ge3.

Perform the following steps to name and create port-based flows.

1. Use the `portcfg lan-stats --flow create` command to create a port-based flow. The following command creates a flow for a GE port with the name GE3.

```
switch:admin> portcfg lan-stats --flow GE3 create --port ge3.dp0
Operation Succeeded.
```

   The port is identified by both its port ID and DP number.

2. This step shows how to create a port-based flow for VE_Port 35, which is already configured as an IP Extension tunnel.
   a. Use the `portcfgshow fciptunnel` command to display VE_Port tunnels; this step is optional.

```
switch:admin> portcfgshow fciptunnel

Tunnel Circuit   AdminSt Flags
-------------------------------------------------------------------------------
7/25   -         Enabled --i----PI
7/35   -         Enabled --i---DPI
-------------------------------------------------------------------------------
Flags (tunnel): l=Legacy QOS Mode
                i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                I=IP-Ext
```

   b. To create the flow for a VE_Port on a blade in slot 7, use the `portcfg lan-stats --flow create` command.

```
switch:admin> portcfg lan-stats --flow test create --port 7/25
Operation Succeeded.
```

3. To view the flow names that you created on the platform, use the `portshow lan-stats --flow` command. The Flow Info column shows the flow definition that is associated with the flow name.

```
switch:admin> portshow lan-stats --flow

Name             Flow Info
-------------------------------------------------------------------------------
APP-flow         (TCP:3225-3226)
host223          (port 7/25)
test             (port 7/25)
-------------------------------------------------------------------------------
```

4. To view the flows that you created, use the `portshow lan-stats --per-flow -flow` *name* command where *name* is one of the flow names you created.

```
switch:admin> portshow lan-stats --per-flow -flow test

Aggregate Info:
Flow       Pro Cnt   TX(B/s)   RX(B/s) CTx(B) CRx(B) CR  ReTx
-------------------------------------------------------------------------------
Test       TCP 31    591.8m    0       -      -      -   198
-------------------------------------------------------------------------------
Pro=Protocol Cnt=Connection Count
CTx(B)=Post-Compression bytes CRx(B)=Pre-Compression bytes
CR=Compression-Ratio
ReTx=ReTransmission

*UDP flows not considered for Compression stats*

Individual Info:
DP      Idx  Src-Address     Dst-Address     Sport Dport Pro Tx(B/s) Rx(B/s)
```

```
-------------------------------------------------------------------------------
7/DP0   20   192.168.10.76   192.168.20.38   53023 59780 TCP 19.6m   0
7/DP0   23   192.168.10.76   192.168.20.38   53012 59780 TCP 19.7m   0
7/DP0   21   192.168.20.38   192.168.10.76   59779 52994 TCP 0       0
7/DP0   25   192.168.10.76   192.168.20.38   53016 59780 TCP 19.7m   0
7/DP0   26   192.168.10.76   192.168.20.38   53025 59780 TCP 19.7m   0
7/DP0   30   192.168.10.76   192.168.20.38   53021 59780 TCP 19.9m   0
7/DP0   27   192.168.10.76   192.168.20.38   53020 59780 TCP 19.8m   0
7/DP0   13   192.168.10.76   192.168.20.38   53018 59780 TCP 19.7m   0
7/DP0   8    192.168.10.76   192.168.20.38   53003 59780 TCP 19.7m   0
7/DP0   15   192.168.10.76   192.168.20.38   53001 59780 TCP 19.6m   0
7/DP0   12   192.168.10.76   192.168.20.38   52997 59780 TCP 19.7m   0
7/DP0   18   192.168.10.76   192.168.20.38   53008 59780 TCP 19.8m   0
7/DP0   16   192.168.10.76   192.168.20.38   53024 59780 TCP 19.7m   0
7/DP0   10   192.168.10.76   192.168.20.38   53006 59780 TCP 19.7m   0
7/DP0   22   192.168.10.76   192.168.20.38   53011 59780 TCP 19.6m   0
7/DP0   14   192.168.10.76   192.168.20.38   53007 59780 TCP 19.3m   0
7/DP0   7    192.168.10.76   192.168.20.38   53013 59780 TCP 19.7m   0
7/DP0   6    192.168.10.76   192.168.20.38   53002 59780 TCP 19.7m   0
7/DP0   4    192.168.10.76   192.168.20.38   52999 59780 TCP 19.6m   0
7/DP0   1    192.168.10.76   192.168.20.38   53004 59780 TCP 19.7m   0
7/DP0   3    192.168.10.76   192.168.20.38   52998 59780 TCP 19.8m   0
7/DP0   28   192.168.10.76   192.168.20.38   53010 59780 TCP 19.4m   0
7/DP0   0    192.168.10.76   192.168.20.38   53022 59780 TCP 19.8m   0
7/DP0   9    192.168.10.76   192.168.20.38   53017 59780 TCP 19.8m   0
7/DP0   17   192.168.10.76   192.168.20.38   53014 59780 TCP 19.7m   0
7/DP0   29   192.168.10.76   192.168.20.38   53026 59780 TCP 19.7m   0
7/DP0   11   192.168.10.76   192.168.20.38   53009 59780 TCP 19.8m   0
7/DP0   5    192.168.10.76   192.168.20.38   53005 59780 TCP 19.7m   0
7/DP0   2    192.168.10.76   192.168.20.38   53000 59780 TCP 19.7m   0
7/DP0   24   192.168.10.76   192.168.20.38   53019 59780 TCP 19.3m   0
7/DP0   19   192.168.10.76   192.168.20.38   53015 59780 TCP 19.7m   0
-------------------------------------------------------------------------------
Sport=Source-Port Dport=Destination-Port Pro=Protocol

DP     ActTCP  ExdTCP  TCLDeny TCLFail
-------------------------------------------------------------------------------
7/DP0  31      0       0       0
7/DP1  21      0       0       0
-------------------------------------------------------------------------------
ActTCP=Active TCP Conns ExdTCP=Exceeded TCP Conn Cnt
```

Configuring an IP Address Flow

You can create a named flow that filters the IP Extension Flow Monitor output based on a specific IP address. The IP address can be either a source address or destination address that is associated with a TCP connection. This section provides examples on naming and creating an IP address flow.

You can view the IP addresses associated with the TCP connections using the `portshow lan-stats --per-flow` command. This step is optional. The following example displays the IP addresses associated with the TCP connections:

```
switch:admin> portshow lan-stats --per-flow
```

```
*** Displaying Top 25 connections by throughput ***
--------------------------------------------------------------------------------
DP     Idx  Src-Address    Dst-Address    Sport   Dport   Pro   Tx(B/s)  Rx(B/s)
--------------------------------------------------------------------------------
7/DP0  9    192.168.10.76  192.168.20.38  53017   59780   TCP   19.9m    0
7/DP0  12   192.168.10.76  192.168.20.38  52997   59780   TCP   19.9m    0
7/DP0  27   192.168.10.76  192.168.20.38  53020   59780   TCP   19.9m    0
7/DP0  2    192.168.10.76  192.168.20.38  53000   59780   TCP   19.8m    0
7/DP0  19   192.168.10.76  192.168.20.38  53015   59780   TCP   19.8m    0
[Output is truncated.]
```

You can create a flow with the IP address 192.168.20.38 using the `portcfg lan-stats --flow create` command as shown in the following example:

```
switch:admin> portcfg lan-stats --flow host create --ipaddr 192.168.20.38
Operation Succeeded.
```

You can view the flow names that you created on the platform using the `portshow lan-stats --flow` command. The Flow Info column displays the flow definition that is associated with the flow name, as shown in the example below:

```
switch:admin> portshow lan-stats --flow
Name              Flow Info
--------------------------------------------------------------------------------
APP-flow          (TCP:3225-3226)
Host              (ipaddr:192.168.10.76/32)
host223           (port 7/25)
test              (port 7/25)
--------------------------------------------------------------------------------
```

You can view the flows that you created using the `portshow lan-stats --per-flow -flow` *`name`* command where the name is one of the flow names you created. The following example shows the flows where the name is one of the flow names you created:

```
switch:admin> portshow lan-stats --per-flow -flow host

Aggregate Info:
Flow   Pro      Cnt    TX(B/s)     RX(B/s)    CTx(B)     CRx(B)    CR      ReTx
--------------------------------------------------------------------------------
Host   TCP      31     591.7m      0          -          -         -       198
--------------------------------------------------------------------------------
Pro=Protocol Cnt=Connection Count
CTx(B)=Post-Compression bytes CRx(B)=Pre-Compression bytes
CR=Compression-Ratio
ReTx=ReTransmission
*UDP flows not considered for Compression stats*

Individual Info:
--------------------------------------------------------------------------------
DP     Idx  Src-Address    Dst-Address    Sport   Dport   Pro   Tx(B/s)  Rx(B/s)
--------------------------------------------------------------------------------
7/DP0  20   192.168.10.76  192.168.20.38  53023   59780   TCP   19.6m    0
7/DP0  23   192.168.10.76  192.168.20.38  53012   59780   TCP   19.7m    0
7/DP0  21   192.168.20.38  192.168.10.76  59779   52994   TCP   0        0
7/DP0  25   192.168.10.76  192.168.20.38  53016   59780   TCP   19.7m    0
```

```
7/DP0  26   192.168.10.76  192.168.20.38  53025   59780   TCP   19.8m    0
7/DP0  19   192.168.10.76  192.168.20.38  53015   59780   TCP   19.9m    0
--------------------------------------------------------------------------------
Sport=Source-Port Dport=Destination-Port Pro=Protocol
[Output is truncated.]

DP      ActTCP  ExdTCP  TCLDeny TCLFail
--------------------------------------------------------------------------------
7/DP0   31      0       0       0
--------------------------------------------------------------------------------
ActTCP=Active TCP Conns ExdTCP=Exceeded TCP Conn Cnt
```

Configuring a TCP Port Flow

You can create a named flow that filters the IP Extension Flow Monitor output based on a specific TCP port. The examples in this section show how to create and name a TCP port flow.

You can view the TCP ports in use using the `portshow lan-stats --per-flow` command. The following example shows the TCP port numbers in the Sport and Dport columns:

```
switch:admin> portshow lan-stats --per-flow
*** Displaying Top 5 connections by throughput ***

--------------------------------------------------------------------------------
DP     Idx  Src-Address    Dst-Address    Sport   Dport   Pro   Tx(B/s)  Rx(B/s)
--------------------------------------------------------------------------------

DP0    142  192.168.10.76  192.168.20.38  63040   49883   TCP   6.6m     6.6m
DP0    120  192.168.10.76  192.168.20.38  63018   49883   TCP   6.6m     6.6m
DP0    164  192.168.10.76  192.168.20.38  63062   49883   TCP   6.6m     6.6m
DP0    170  192.168.10.76  192.168.20.38  63068   49883   TCP   6.6m     6.6m
DP0    150  192.168.10.76  192.168.20.38  63048   49883   TCP   6.6m     6.6m
--------------------------------------------------------------------------------
Sport=Source-Port Dport=Destination-Port Pro=Protocol
DPActTCP ExdTCP TCLDeny TCLFail
--------------------------------------------------------------------------------
DP0101000
DP10000
--------------------------------------------------------------------------------
ActTCP=Active TCP Conns ExdTCP=Exceeded TCP Conn Cnt
```

You can view a list that shows TCP ports associated with known application ports using the `portshow lan-stats --known-apps` command. This step is optional. You can use the Port-ID(s) to create a TCP port flow. You can also use the application information to create an application-based flow.

The following example shows the TCP ports associated with known application ports:

```
switch:admin> portshow lan-stats --known-apps

App                   Port-Id(s)
--------------------------------------------------------------------
 CIFS                 139,445
 Data-Domain          2051
 FCIP                 3225-3226
 FTP                  20-21,989-990,115
```

```
 HTTP                 80,8080,8000-8001,3128
 HTTPS                443
 Isilon-SyncIQ        5666-5667
 LDAP                 389,8404,636
 MS-SQL               1443
 MySQL                3306
 NETAPP-SNAP-MIRROR   10566
 NFS                  2049
 ORACLE-SQL           66,1525,1521
 RSYNC                873
 SRDF                 1748
 SSH                  22
 SSL-SHELL            614
 TELNET               23,107,513,992
 TFTP                 69
 VERITAS-BACKUP       6101-6102,6106,3527,1125
 VTS-GRID Control     1415-1416
 VTS-GRID Data        350
 iSCSI                3260
-------------------------------------------------------------------------
```

You can create a TCP flow with the port address 49883 using the `portcfg lan-stats --flow create` command. The following example shows the TCP flow with the port address 49883:

```
switch:admin> portcfg lan-stats --flow TCP-flow create --tcp 49883
Operation Succeeded.
```

You can view the flow names that you created on the platform using the `portshow lan-stats --flow` command. The following example shows the flow definition associated with the flow name in the Flow Info columns:

```
admin> portshow lan-stats --flow

Name        Flow Info
-------------------------------------------------------------------------------
GE3         (port:ge3.DP0)
TCP-flow    (TCP:49883)
VE35        (port:35)
host38      (ipaddr:192.168.20.38/32)
-------------------------------------------------------------------------------
```

You can view the flows that you created using the `portshow lan-stats --per-flow -flow` *`name`* command where the name is one of the flow names you created.

```
switch:admin> portshow lan-stats --per-flow -flow TCP-flow
No Active Per-flow stats to Display
```

You can view a listing of all TCP ports that attempted to connect using the `portshow lan-stats --per-flow -tcp` command. The following example shows a list of all TCP ports:

```
switch:admin> portshow lan-stats --per-flow -tcp
 *** Displaying Top 5 connections by throughput ***

 DP Idx Src-Address Dst-Address Sport Dport Pro Tx(B/s) Rx(B/s)
 TCP TxPkt RxPkt TxDrp RxDrp ReTx DpAck OOO RTT FlwCtrl
-------------------------------------------------------------------------------
 DP0 1105 10.75.18.183 10.89.16.51 53620 22 TCP 441 109.2k
```

```
 880.2k 2.8m 0 0 23 31 24 16 0
 DP0 4794 10.187.98.30 10.75.20.72 47496 11111 TCP 10 1
 8 3 0 0 0 0 0 1 0
 DP0 324 10.75.20.63 10.187.98.45 53634 11111 TCP 0 0
 2 2 0 0 0 0 0 0 0
 DP0 809 10.75.20.128 10.187.98.45 49464 11111 TCP 0 0
 2 2 0 0 0 0 0 0 0
 DP0 1519 10.187.98.26 10.75.20.71 38400 20000 TCP 0 0
 4 1 0 0 0 0 0 1 0
-------------------------------------------------------------------------------
 Sport=Source-Port Dport=Destination-Port Pro=Protocol
 TxPkt=Tx-Packets Rxpkt=Rx-Packets
 TxDrp=TX-Drops RxDrp=RX-Drops
 ReTx=ReTransmission OOO=out-of-order
 DpAck=Duplicate-Acks RTT=Round-Trip-Time(milliseconds)
 FlwCtrl=Number of Flow-controls
 DP ActTCP ExdTCP TCLDeny TCLFail
-----------------------------------------------------------
 DP0 144 421364611047724 0
 DP1 0 23796 224816520
-----------------------------------------------------------
 ActTCP=Active TCP Conns ExdTCP=Exceeded TCP Conn Cnt
```

Configuring a Flow Using Logical Operators

You can use the logical operators AND and OR to create flows that meet specific criteria. Logical operators are restricted by the following conditions:

- A flow definition can contain a maximum of 30 elements.
- You cannot combine AND operators with OR operators in a flow definition.

You can use the `portcfg lan-stats -- flow create` command to combine a specific DP with an IP address and VE_Port using the AND operator. The following example shows that DP1 is selected along with an IP address on a specific port. Traffic flow statistics that meet all criteria are displayed for the dp1_225 flow.

```
switch:admin> portcfg lan-stats --flow dp1_225 create --dp dp1 --and --ipaddr 192.168.10.76/24 --port 24
Operation Succeeded.
```

You can use the `portcfg lan-stats --flow create` command to combine a DP with an IP address and VE_Port using the OR operator. The following example shows the statistics for all TCP traffic on DP0 are displayed, along with all traffic on VE24 or IP address 192.168.10.76:

```
switch:admin> portcfg lan-stats --flow dp0_ve35 create --dp dp0 --or --ipaddr 192.168.10.76/24 --port 24
Operation Succeeded.
```

You can use the `portshow lan-stats --flow` command to view the flow names that you created on the platform. The following example shows the flow definition that is associated with the flow name in the Flow Info column.

```
switch:admin> portshow lan-stats --flow

Name        Flow Info
-----------------------------------------------------------
dp0_ve35    (dp:DP0 or ipaddr:192.168.10.76/24 or port 24)
dp1_225     (dp:DP1 and ipaddr:192.168.10.76/24 and port 24)
-----------------------------------------------------------
```

Displaying Historical Flow Statistics

You can display historical statistics for the LAN traffic in the IP Extension Flow Monitor feature. The historical display is a snapshot of information for closed connections that are stored in the DP for both DP0 and DP1. You can freeze or thaw the statistics. Freezing the statistics means that no additional updates are made. All activity that occurs during the freeze is not captured. When you thaw the statistics, activity is updated and information about closed connections is again captured.

To filter information based on specific criteria, you can apply options to the command. For example, you can filter on a DP or you can define additional filters. The following examples show various options to view the historical flow statistics.

The default display for historical flow statistics shows a summary view of the first five and last five entries. When you select a detailed view, first freeze the historical statistics.

You can view the default view of historical statistics using the `portshow lan-stats --hist-stats` command. The following CLI output shows the first and last five entries and summary information for each DP:

```
switch:admin> portshow lan-stats --hist-stats

DP0 Connection Summary: (Thawed)
--------------------------------------------------------------------------------
 Idx  Src-Address         Dst-Address         Sport  Dport  Pro  Tx(B)    Rx(B)
--------------------------------------------------------------------------------
First 5 Connections:
 4878 10.97.8.250         10.75.246.251       5667   57118  TCP  180      168
 540  10.89.122.250       10.75.246.251       5667   56854  TCP  180      168
 669  10.187.98.37        10.75.20.71         38326  11111  TCP  452      272
 1390 10.187.98.26        10.75.20.71         47640  20000  TCP  1.0k     516
 1427 10.187.98.45        10.75.20.139        11111  57406  TCP  228      356
Last 5 Connections:
 398  10.187.98.30        10.75.20.72         54200  11111  TCP  592      268
 1644 10.187.98.26        10.75.20.74         45162  20000  TCP  616      324
 1883 10.187.98.30        10.75.20.72         53320  11111  TCP  592      268
 833  10.187.98.26        10.75.20.74         44314  20000  TCP  616      324
 1274 10.187.98.45        10.75.20.128        11111  51682  TCP  840      2.0k
--------------------------------------------------------------------------------

Total Connection count:   21
Oldest Entry:             12/02/2022 11:34:47 PST
Newest Entry:             12/02/2022 11:42:09 PST
Close RX/TX FIN:          7 / 14
Close RX/TX RST:          0 / 0
Total TX Errors
 Slow Starts:             0
 FastRetrans/RetransTO:   0 / 0
Total RX Errors
 Out of Orders/Dup Ack:   0 / 0

DP1 Connection Summary: (Thawed)
--------------------------------------------------------------------------------
 Idx  Src-Address         Dst-Address         Sport  Dport  Pro  Tx(B)    Rx(B)
--------------------------------------------------------------------------------
First 5 Connections:
 5844 10.86.16.251        10.75.246.251       5667   34498  TCP  1.2k     1.8k
 5489 10.86.16.250        10.75.246.252       5667   46632  TCP  1.4k     1.5k
```

```
 7295 10.86.16.252        10.75.246.250       5667   56607  TCP  2.0k     14.8k
 7472 10.86.16.250        10.75.246.250       5667   56599  TCP  1.2k     1.2k
 8083 10.110.41.250       10.75.246.251       5667   50858  TCP  2.0k     12.6k
Last 5 Connections:
 9758 10.110.41.250       10.75.246.252       5667   54200  TCP  668      1.1k
 8531 10.110.41.251       10.75.246.250       5667   31389  TCP  588      991
 6690 10.110.41.251       10.75.246.250       5667   31386  TCP  588      991
 8933 10.110.41.250       10.175.58.251       5667   53688  TCP  1.5k     1.3k
 8189 10.110.41.251       10.75.246.250       5667   31141  TCP  1.1k     1.2k
--------------------------------------------------------------------------------

Total Connection count:   1038
Oldest Entry:             12/02/2022 04:00:30 PST
Newest Entry:             12/02/2022 11:39:01 PST
Close RX/TX FIN:          1015 / 23
Close RX/TX RST:          0 / 0
Total TX Errors
 Slow Starts:             57
 FastRetrans/RetransTO:   0 / 57
Total RX Errors
 Out of Orders/Dup Ack:   0 / 0
```

You can freeze the historical flow data using the `portshow lan-stats --hist-stats -freeze` command as shown in the following CLI output:

```
switch:admin> portshow lan-stats --hist-stats -freeze
Operation Succeeded.
DP0 Hist Status: Frozen
DP1 Hist Status: Frozen
```

You can thaw the historical flow data using the `portshow lan-stats --hist-stats -thaw` command as shown in the following CLI output:

```
switch:admin> portshow lan-stats --hist-stats -thaw
Operation Succeeded.
DP0 Hist Status: Thawed
DP1 Hist Status: Thawed
```

You can view details of the first and last five entries in details using the `portshow lan-stats --hist-stats -detail` command. When 10 or fewer entries are stored, all connection details are displayed. Use the `-all` option to display details for all entries that are stored in the DP instead of the first and last five entries. Freeze the table before displaying detailed statistics.

```
switch:admin> portshow lan-stats --hist-stats -detail

DP0 Connection Summary: (Frozen)
--------------------------------------------------------------------------------
 Idx  Src-Address         Dst-Address         Sport  Dport  Pro  Tx(B)    Rx(B)
--------------------------------------------------------------------------------
First 5 Connections:
--------------------------------------------------------------------------------
 834  10.187.98.45        10.75.20.127        11111  42786  TCP  228      356
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 1517 10.187.98.45        10.75.20.128        11111  39620  TCP  232      496
--------------------------------------------------------------------------------
```

```
--------------------------------------------------------------------------------
 1986 10.187.98.26        10.75.20.73         46410  20000  TCP  616      324
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 1193 10.187.98.45        10.75.20.129        11111  47616  TCP  228      356
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 3347 10.187.98.45        10.75.20.134        11111  49996  TCP  228      356
--------------------------------------------------------------------------------

Last 5 Connections:
--------------------------------------------------------------------------------
 397  10.187.98.42        10.75.20.71         34736  11111  TCP  452      272
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 3181 10.187.98.26        10.75.20.74         48936  20000  TCP  616      324
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 1965 10.187.98.45        10.75.20.136        11111  36542  TCP  228      356
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 751  10.187.98.45        10.75.20.128        11111  58236  TCP  232      496
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 1844 10.187.98.30        10.75.20.72         32906  11111  TCP  1.7k     716
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------

Total Connection count:   65
Oldest Entry:             12/02/2022 11:39:05 PST
Newest Entry:             12/02/2022 11:47:21 PST
Close RX/TX FIN:          38 / 27
Close RX/TX RST:          0 / 0
Total TX Errors
 Slow Starts:             0
 FastRetrans/RetransTO:   0 / 0
Total RX Errors
 Out of Orders/Dup Ack:   0 / 0

DP1 Connection Summary: (Frozen)
--------------------------------------------------------------------------------
 Idx  Src-Address         Dst-Address         Sport  Dport  Pro  Tx(B)    Rx(B)
--------------------------------------------------------------------------------
First 5 Connections:
--------------------------------------------------------------------------------
```

```
 5489 10.86.16.250        10.75.246.252       5667   46632  TCP  1.4k     1.5k
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 5844 10.86.16.251        10.75.246.251       5667   34498  TCP  1.2k     1.8k
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 7295 10.86.16.252        10.75.246.250       5667   56607  TCP  2.0k     14.8k
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 7472 10.86.16.250        10.75.246.250       5667   56599  TCP  1.2k     1.2k
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 9530 10.110.41.251       10.75.246.252       5667   30412  TCP  1.6k     1.7k
--------------------------------------------------------------------------------

Last 5 Connections:
--------------------------------------------------------------------------------
 8531 10.110.41.251       10.75.246.250       5667   31389  TCP  588      991
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 6690 10.110.41.251       10.75.246.250       5667   31386  TCP  588      991
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 6392 10.110.41.251       10.75.246.252       5667   54210  TCP  588      991
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 8933 10.110.41.250       10.175.58.251       5667   53688  TCP  1.5k     1.3k
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
 8189 10.110.41.251       10.75.246.250       5667   31141  TCP  1.1k     1.2k
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------

Total Connection count:   1038
Oldest Entry:             12/02/2022 04:00:30 PST
Newest Entry:             12/02/2022 11:39:01 PST
Close RX/TX FIN:          1015 / 23
Close RX/TX RST:          0 / 0
Total TX Errors
 Slow Starts:             57
 FastRetrans/RetransTO:   0 / 57
Total RX Errors
 Out of Orders/Dup Ack:   0 / 0
```

You can view all options that you can use to filter the historical statistics using the `portshow lan-stats --hist-stats -filter -help` command. This command returns a list of all options available for the `portshow` command as shown in the following CLI output:

```
switch:admin> portshow lan-stats --hist-stats -help

Usage:
  portshow lan-stats --hist-stats [<hargs>]

  Hist-Stats Args: Displays the historical TCP connection info for
  closed TCP IP-Extension flows.
    -all              Display all stored connections. Default 10
                      entries if omitted.
    -detail           Displays the connection stats in detailed view
    -freeze           Freeze the historical TCP connection table to
                      prevent existing entries from being removed and
                      new entries from being added.
    -thaw             Thaw the historical TCP connection table allowing
                      existing entries to be removed and new entries to
                      be added.
    -dp [<slot>/]dp<#> [-index <index id>]  Specify a target DP
                      to get the historical TCP stats from.
    -filter <args>    Limit the output to specific filter criteria.
                      Use portShow lan-stats --hist-stats -filter -help
                      for details.
```

Displaying IP Pair Details

IP pair statistics are useful for understanding the traffic flow that occurs across all TCP connections that are created between a source and destination IP address. See Using IP Extension Flow Monitor for information on the default view of IP pair statistics. IP pair details can provide information about the TCP connections on a specific DP and index value. The following example shows detailed information on IP pairs.

You can view the detailed information for IP pairs, use the `portshow lan-stats --ip-pair -detail` command. The output shows detailed IP pair information. Each IP pair has a unique index value that is used to filter the results. In the following example, only DP0 has an active IP pair connection.

```
switch:admin> portshow lan-stats --ip-pair -detail

 DP0: Index:13 10.187.98.25 <-> 10.75.20.74
-------------------------------------------------------------------------------
 No of active connections:  3
 No of closed connections:  759694
 Creation time              03/08/2022 06:04:31 PST
 Last snapshot time:        12/02/2022 08:36:43 PST
                            TxB       RxB
 Current stats:             76.9m     229.3k
 Thu 12/01 stats:           533.9m    1.5m
 Wed 11/30 stats:           464.7m    1.5m
 Tue 11/29 stats:           1.0g      1.8m
 Mon 11/28 stats:           586.7m    1.4m
 Sun 11/27 stats:           910.2m    3.1m
 Sat 11/26 stats:           315.6m    1.0m
 Fri 11/25 stats:           663.5m    1.5m
```

```
 Total stats:              14.0g     2.3g

[Output truncated…]
```

If the DP and index number are already known, to display IP pair details for a specific DP and index value, use the `portshow lan-stats --ip-pair [slot/dp] [index] -detail` command.

Displaying IP Pair History

IP pair history provides a snapshot of activity every 24 hours. The snapshot is saved for seven days unless the history is reset. The time of the last snapshot is displayed in the detailed view.

A snapshot takes several seconds to accumulate, so the 24-hour timestamp might vary from day to day. For information on how to display historical flow history, see Displaying Historical Flow Statistics.

You can view the IP pair history using the `portshow lan-stats --ip-pair -hist` command. The values in the *Today* column are a snapshot of when the command was entered.

```
switch:admin> portshow lan-stats --ip-pair -hist

 DP       Idx SrcAddr             DstAddr             Active  Closed
             Today   Thu     Wed     Tue     Mon     Sun     Sat     Fri
-------------------------------------------------------------------------------
 DP0      22  10.89.122.251       10.75.244.251       4       57513
   TxB       27.1k   131.3k  167.6k  121.1k  267.1k  212.3k  127.5k  123.8k
   RxB       99.7m   363.5k  2.0g    383.3k  1.8g    3.9g    388.1k  356.6k

 DP0      9   10.187.98.38        10.75.20.74         7       762436
   TxB       73.6m   721.4m  1.4g    1.4g    1.0g    1.9g    182.2m  367.5m
   RxB       227.7k  1.7m    1.3m    2.8m    2.5m    6.8m    720.9k  961.2k

 DP0      11  10.187.98.26        10.75.20.74         7       787880
   TxB       274.9m  2.4g    870.8m  2.0g    1.8g    1.1g    650.2m  661.8m
   RxB       723.9k  2.4m    2.9m    2.5m    4.0m    3.9m    1.5m    2.0m

 DP0      13  10.187.98.25        10.75.20.74         0       759729
   TxB       76.9m   533.9m  464.7m  1.0g    586.7m  910.2m  315.6m  663.5m
   RxB       229.3k  1.5m    1.5m    1.8m    1.4m    3.1m    1.0m    1.5m

[Output truncated]
-------------------------------------------------------------------------------

There are connections available which does not belong to any IP pair
```

You can specify a DP and index value for the history display, use the `portshow lan-stats --ip-pair [slot/port] [index] -hist` command. Add the `-detail` option to obtain a detailed history.

```
switch:admin> portshow lan-stats --ip-pair dp0 13 -hist -detail

 DP0: Index:13 10.187.98.25 <-> 10.75.20.74
-------------------------------------------------------------------------------
 No of active connections:  0
 No of closed connections:  759750
 Creation time              03/08/2022 06:04:31 PST
 Last snapshot time:        12/02/2022 08:36:43 PST
                            TxB       RxB
```

```
  Current stats:              76.9m     229.3k
  Thu 12/01 stats:            533.9m    1.5m
  Wed 11/30 stats:            464.7m    1.5m
  Tue 11/29 stats:            1.0g      1.8m
  Mon 11/28 stats:            586.7m    1.4m
  Sun 11/27 stats:            910.2m    3.1m
  Sat 11/26 stats:            315.6m    1.0m
  Fri 11/25 stats:            663.5m    1.5m
  Total stats:                14.0g     2.3g

 There are connections available which does not belong to any IP pair
```

Resetting IP Pair Statistics

You can reset the IP pair statistics to start the data gathering from zero. Current statistics are set to zero and then updated only with what has transpired since the reset was entered. When there is no active connection with a saved IP pair, that IP pair is removed and not displayed. Conditions other than a command entry can reset the IP pair statistics. Those conditions are as follows:

- Firmware upgrade
- Firmware downgrade
- Platform reboot

NOTE
When IP pair statistics are reset, no warning is displayed or confirmation is required.

You can control which IP pair statistics are reset. For example, you can reset statistics on a specific DP, or you can reset statistics for a particular index number. The following examples show how to reset IP pair statistics:

You can reset all IP pair statistics using the `portshow lan-stats --ip-pair reset` command as shown in the following example:

```
admin> portshow lan-stats --ip-pair -reset
IP-Pair reset status:Success
```

You can specify values for the DP and IP pair index by using the `portshow lan-stats --ip-pair reset` command with the DP option and index option to reset only certain IP pair statistics. Use the `portshow lan-stats --ip-pair` command to identify a DP and an index, and then use the `portshow lan-stats --ip-pair reset` command.

The following CLI output shows an example with the IP pair information:

```
switch:admin> portshow lan-stats --ip-pair -hist

  DP       Idx SrcAddr             DstAddr             Active  Closed
              Today   Thu     Wed     Tue     Mon     Sun     Sat     Fri
-------------------------------------------------------------------------------
  DP0      22  10.89.122.251       10.75.244.251       0       57517
    TxB       36.5k   131.3k  167.6k  121.1k  267.1k  212.3k  127.5k  123.8k
    RxB       287.9m  363.5k  2.0g    383.3k  1.8g    3.9g    388.1k  356.6k

  DP0      9   10.187.98.38        10.75.20.74         0       762478
    TxB       73.6m   721.4m  1.4g    1.4g    1.0g    1.9g    182.2m  367.5m
    RxB       227.7k  1.7m    1.3m    2.8m    2.5m    6.8m    720.9k  961.2k

  DP0      11  10.187.98.26        10.75.20.74         1       787923
    TxB       274.9m  2.4g    870.8m  2.0g    1.8g    1.1g    650.2m  661.8m
    RxB       724.3k  2.4m    2.9m    2.5m    4.0m    3.9m    1.5m    2.0m
```

```
 DP0      13  10.187.98.25        10.75.20.74         4       759764
   TxB        76.9m   533.9m  464.7m  1.0g    586.7m  910.2m  315.6m  663.5m
   RxB        229.3k  1.5m    1.5m    1.8m    1.4m    3.1m    1.0m    1.5m
[output truncated]
```

You can reset the IP pair statistics for DP0 and index 13 as shown in the following example:

```
switch:admin> portshow lan-stats --ip-pair dp0 13 -reset
IP-Pair reset status:Success
```

You can confirm that the selected IP pair statistics are reset. Any IP pair that has no active connection is removed from the display table when the statistics are reset.

```
switch:admin> portshow lan-stats --ip-pair dp0 25 -hist -summary

 DP       Idx SrcAddr             DstAddr             Active  Closed
              Today   Thu     Wed     Tue     Mon     Sun     Sat     Fri
-------------------------------------------------------------------------------
 DP0      13  10.187.98.25        10.75.20.74         0       758083
   TxB        0       -       -       -       -       -       -       -
   RxB        0       -       -       -       -       -       -       -

-------------------------------------------------------------------------------

There are connections available which does not belong to any IP pair
```

Displaying IP Extension Performance

You can view the summary or aggregate of all slot or DP instances within the chassis. The `portshow lan-stats --status` command displays aggregate bandwidth that is utilized by IP Extension, the total number of connections, and a summary of errors. The `portshow lan-stats --status` command is supported only when the CP is running Fabric OS version 9.1.0 or higher.

The following CLI output shows an example of the `portshow lan-stats --status` command:

```
switch:maintenance-service> portshow lan-stats --status

Ip-Extension Status
===============================================================================
 Slot/DP                 0/dp0  0/dp1
 Active Emulated TCP
  Connection Count    :    102    102
  Tx Bps (30s avg)    :  56.9m  57.2m
  Rx Bps (30s avg)    :  56.8m  57.2m
  Fast Retransmits    :      0      0
  Slow Retransmits    :   4509  11.8k
  Out of Sequence     :      0      0
  Tx Window Closed    :    677      0
  Rx Window Closed    :      0      0
  Slow Drain          :      0      0
 Active Emulated UDP
  Connection Count    :      0      0
  Tx Bps (30s avg)    :      0      0
  Rx Bps (30s avg)    :      0      0
 Global Emulated
  Max TCP Connections :    104    104
```

```
  Max UDP Connections :     58     58
 Global Non-Terminated
  TCP Tx Packets      :      0      0
  TCP Rx Packets      :      0      0
  TCP Drop Packets    :      0      0
  UDP Tx Packets      :      0      0
  UDP Rx Packets      :   5490   5490
  UDP Drop Packets    :   3565   3564
  ICMP Tx Packets     :      0      0
  ICMP Rx Packets     :     18     18
  ICMP Drop Packets   :     18     18
 Global Rx Errors
  TCP Checksum        :      0      0
  IP Checksum         :      0      0
  Ethernet CRC        :      0      0
  IP Fragment Drop    :      0      0
--------------------------------------------------------------------------------
```

The following example shows the LAN side counters for each DP. In this example, only DP0 is shown.

```
switch:admin> portshow lan-stats --global -dp

LAN Global stats:
--------------------------------------------------------------------------------
 Active TCP conn                   :0
 Active TCP conn on Local Backup   :0
 Active TCP conn on Remote Backup  :0
 Establish TCP conn                :7
 Closed TCP conn                   :7
 TCP Tx-Bytes                      :0
 TCP Rx-Bytes                      :0
 TCP Tcl-Deny conn                 :0
 TCP Tcl lookup fail               :0
 Sync-Recv                         :1
 Sync-fail                         :0
 Drop-Bytes                        :0
 Drop-Pkts                         :0
 Stale Reset from Host             :4
 Number of times
 Max TCP conn exceeded as client   :0
 Number of times
 Max TCP conn exceeded as server   :0
 Number of times Max TCP conn
 per second exceeded as client     :0
 Number of times Max TCP conn
 per second exceeded as server     :0
 Number of times
 Max UDP conn exceeded on Egress   :0
  Number of UDP Pkts sent ASIS      :0
 Number of times
 TX PDU preserve ON                :0
 Number of times
 RX PDU preserve ON                :0
 Total IPV6 pkts                   :46489
```

```
FlowControl on                    :0
FlowControl off                   :0
Active UDP conn                   :0
Establish UDP conn                :11555
Closed UDP conn                   :11555
UDP route lookup fail             :52752
UDP PDU drops due to
PKO flow control                  :0
TX UDP PDUs                       :0
TX UDP PDU drops                  :27642
RX UDP PDUs                       :57760
RX UDP Tcl lookup fail PDUs       :96
RX UDP Tcl-Deny PDUs              :28867
Total RX UDP PDU drops            :28964
RX UDP PDU drops due to
stream flow control               :0
TX ICMP PDUs                      :0
TX ICMP PDU drops                 :80
RX ICMP PDUs                      :55
RX ICMP Tcl lookup fail PDUs      :9
RX ICMP Tcl-Deny PDUs             :23
Total RX ICMP PDU drops           :32
RX ICMP PDU drops due to
stream flow control               :0
TX ASIS IP PDUs                   :0
TX ASIS IP PDU drops              :25030
RX ASIS IP PDUs                   :49302
RX ASIS IP Tcl lookup fail PDUs   :68
RX ASIS IP Tcl-Deny PDUs          :24624
Total RX ASIS IP PDU drops        :24693
RX ASIS IP PDU drops due to
stream flow control               :0
RX Error IP Checksum              :0
RX Error TCP Checksum             :0
RX Error MAC                      :0
RX Error CRC                      :0
RX Error Parity                   :0
RX Error Length                   :0
TX UDP pkts < 64 bytes            :0
TX UDP pkts < 128 bytes           :0
TX UDP pkts < 256 bytes           :0
TX UDP pkts < 512 bytes           :0
TX UDP pkts < 1024 bytes          :0
TX UDP pkts < 1500 bytes          :0
TX UDP pkts < 3000 bytes          :0
TX UDP pkts < 4500 bytes          :0
TX UDP pkts < 6000 bytes          :0
TX UDP pkts < 9000 bytes          :0
RX UDP pkts < 64 bytes            :23554
RX UDP pkts < 128 bytes           :19804
RX UDP pkts < 256 bytes           :4446
RX UDP pkts < 512 bytes           :0
RX UDP pkts < 1024 bytes          :9956
```

```
 RX UDP pkts < 1500 bytes          :0
 RX UDP pkts < 3000 bytes          :0
 RX UDP pkts < 4500 bytes          :0
 RX UDP pkts < 6000 bytes          :0
 RX UDP pkts < 9000 bytes          :0
-------------------------------------------------------------------------------
```

NOTE
The unit suffix indicators of k, m, g, and t follow the typical metric values of kilo, mega, giga, and tera. These units are used to indicate the size of the counters.

The following example immediately resets the Global counters to zero. The `portshow lan-stats --global -reset` command resets only the Global counters. The Active Emulated TCP and Active Emulated UDP counters are not affected. The following CLI output is from a Brocade Director with four Brocade SX6 Extension Blades:

```
switch:admin> portshow lan-stats --global -reset

Ip-Extension Status (Global Stats Since: 11/02/2023-08:34:10)
================================================================================
Slot/DP             3/dp0  3/dp1  4/dp0  4/dp1  7/dp0  7/dp1  8/dp0  8/dp1
Active Emulated TCP
Connection Count:      41    495    453      0      0      0      0      0
Tx Bps (30s avg):       0 618.1m 872.3m      0      0      0      0      0
Rx Bps (30s avg):   42.0m 466.7m 872.5m      0      0      0      0      0
Fast Retransmits:       0 109.8m    524      0      0      0      0      0
Slow Retransmits:       0  1040m   1500      0      0      0      0      0
Out of Sequence:        0 207.5k   4913      0      0      0      0      0
Zero Window Tx:         0      0      0      0      0      0      0      0
Zero Window Rx:         0      0      0      0      0      0      0      0
Slow Drain:             0      0      0      0      0      0      0      0
Active Emulated UDP
Connection Count:       1      0      0      0      0      0      0      0
Tx Bps (30s avg):      13      0      0      0      0      0      0      0
Rx Bps (30s avg):      14      0      0      0      0      0      0      0
Global Emulated
Max TCP Connections:   41    495    453      0      0      0      0      0
Max UDP Connections:    1      0      0      0      0      0      0      0
Global Non-Terminated
TCP Tx Packets:         0      0      0      0      0      0      0      0
TCP Rx Packets:         0      0      0      0      0      0      0      0
TCP Drop Packets:       0      0      0      0      0      0      0      0
UDP Tx Packets:         0      0      0      0      0      0      0      0
UDP Rx Packets:         0      0      0      0      0      0      0      0
UDP Drop Packets:       0      0      0      0      0      0      0      0
ICMP Tx Packets:        0      0      0      0      0      0      0      0
ICMP Rx Packets:        0      0      0      0      0      0      0      0
ICMP Drop Packets:      0      0      0      0      0      0      0      0
Global Rx Errors
TCP Checksum:           0      0      0      0      0      0      0      0
IP Checksum:            0      0      0      0      0      0      0      0
Ethernet CRC:           0      0      0      0      0      0      0      0
IP Fragment Drop:       0      0      0      0      0      0      0      0
```

The `-lifetime` command displays the lifetime counters for all DPs in the chassis, if the `-reset` command was previously issued. This command affects only the Global counters. The Active Emulated TCP/UDP counters are not affected.

```
switch:maintenance-service> portshow lan-stats --status -lifetime

Ip-Extension Status
================================================================================
 Slot/DP                 0/dp0  0/dp1
 Active Emulated TCP
  Connection Count    :    102    102
  Tx Bps (30s avg)    :  56.6m  57.3m
  Rx Bps (30s avg)    :  56.6m  57.4m
  Fast Retransmits    :      0      0
  Slow Retransmits    :   4519  11.9k
  Out of Sequence     :      0      0
  Tx Window Closed    :    677      0
  Rx Window Closed    :      0      0
  Slow Drain          :      0      0
 Active Emulated UDP
  Connection Count    :      0      0
  Tx Bps (30s avg)    :      0      0
  Rx Bps (30s avg)    :      0      0
 Global Emulated
  Max TCP Connections :    104    104
  Max UDP Connections :     58     58
 Global Non-Terminated
  TCP Tx Packets      :      0      0
  TCP Rx Packets      :      0      0
  TCP Drop Packets    :      0      0
  UDP Tx Packets      :      0      0
  UDP Rx Packets      :   5498   5498
  UDP Drop Packets    :   3567   3566
  ICMP Tx Packets     :      0      0
  ICMP Rx Packets     :     18     18
  ICMP Drop Packets   :     18     18
 Global Rx Errors
  TCP Checksum        :      0      0
  IP Checksum         :      0      0
  Ethernet CRC        :      0      0
  IP Fragment Drop    :      0      0
--------------------------------------------------------------------------------
```

The portShow Command

This section provides information about the `portshow` command. You use this command to display information about Brocade Extension operations.

For more information about the `portshow` command syntax and options, refer to*Brocade Fabric OS Command Reference Manual*. The following sections identify a few specific outputs that are useful for maintenance and troubleshooting.

Displaying IP Interfaces (IPIF)

Use the `portshow ipif` command to display the IPIF information. The following example shows IPIF information:

```
switch:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.10.1                  / 30   9216  0     U R M I
 ge2.dp0      10.0.0.5                      / 24   AUTO  0     U R M
 lan.dp0      10.10.10.123                  / 24   1500  100   U R M
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

Displaying IP Routes

Use the `portcfgshow iproute` command to display the user configured routes. The following example shows the IP route information. Notice the `portcfgshow` command is used instead of the `portshow` command because it only shows that the iproutes added by the user. The `portshow` command outputs all iproutes including the routes from the operating system, which can be confusing. The iproutes added by the operating system cannot be deleted or altered by the user.

```
switch:admin> portcfgshow iproute

 Port         IP Address                / Pfx  Gateway               Flags
-------------------------------------------------------------------------------
 ge2.dp0      10.2.0.0                  / 24   10.1.0.1              S
 ge3.dp0      10.2.0.0                  / 24   10.1.0.1              S
 ge2.dp1      10.2.0.0                  / 24   10.1.0.1              S
 ge3.dp1      10.2.0.0                  / 24   10.1.0.1              S
-------------------------------------------------------------------------------
 Flags: S=Static X=Crossport
```

Displaying Extension Mode Information

Use the `extncfg --show` command to view the APP-Mode, VE-Mode, and GE-Mode for the extension platforms.

The following example shows the operating modes for the Brocade 7850 Extension Switch:

```
switch:admin> extncfg --show
APP Mode is HYBRID (FCIP with IPEXT)
VE-Mode: configured for 6VE mode.
GE-Mode: Not Applicable.
```

The following example shows the operating modes for the Brocade 7810 Extension Switch:

```
switch:admin> extncfg –show

APP Mode is HYBRID (FCIP with IPEXT)
VE-Mode: Not Applicable.
GE-Mode: Copper
```

Displaying GE Interface Information

Use the `portcfgge --show` command to verify that the GE interface is in LAN mode, show the speed, and to see if GE interfaces are part of a LAG. The following example shows the GE interface configuration for extension platforms:

```
switch:maintenance-service> portcfgge --show

 Port         Speed     Flags     Channel  FEC   Lag Name
-------------------------------------------------------------------------
 ge0          1G        ----      N/A      Off   -
 ge1          10G       --L-      N/A      Off   -
 ge2          10G       --L-      N/A      Off   -
 ge3          1G        ----      N/A      Off   -
 ge4          25G       ----      N/A      CL108 -
 ge5          25G       --L-      N/A      CL108 -
 ge6          25G       --L-      N/A      CL108 -
 ge7          10G       ----      N/A      Off   -
 ge8          10G       ----      N/A      Off   -
 ge9          10G       ----      N/A      Off   -
 ge10         25G       ----      N/A      CL108 -
 ge11         25G       ----      N/A      CL108 -
 ge12         25G       ----      N/A      CL108 -
 ge13         25G       ----      N/A      CL108 -
 ge14         10G       ----      N/A      Off   -
 ge15         10G       ----      N/A      Off   -
 ge16         100G      ----      N/A      CL91  -
 ge17         100G      ----      N/A      CL91  -
-------------------------------------------------------------------------

Flags: A:Auto-Negotiation Enabled  C:Copper Media Type
L:LAN Port G: LAG Member
```

Listing the MAC Addresses of LAN and GE Interfaces

Use the `portcfgge` *`ge#/lan`* `--show -lmac` command to view the MAC addresses of the LAN and GE ports.

The following example lists the MAC addresses of all the LAN and GE interfaces on the extension platform:

```
switch:admin> portcfgge --show -lmac

Port     dpid    MAC Address
-------------------------------------------------------------------------
7/ge0    dp0  c4:f5:7c:88:dc:9c
7/ge0    dp1  c4:f5:7c:88:dc:9c
7/ge1    dp0  c4:f5:7c:88:dc:9d
7/ge1    dp1  c4:f5:7c:88:dc:9d
7/ge2    dp0  c4:f5:7c:88:dc:9e
7/ge2    dp1  c4:f5:7c:88:dc:9e
7/ge3    dp0  c4:f5:7c:88:dc:9f
```

```
7/ge3    dp1  c4:f5:7c:88:dc:9f
7/ge4    dp0  c4:f5:7c:88:dc:a0
7/ge4    dp1  c4:f5:7c:88:dc:a0
7/ge5    dp0  c4:f5:7c:88:dc:a1
7/ge5    dp1  c4:f5:7c:88:dc:a1
7/ge6    dp0  c4:f5:7c:88:dc:a2
7/ge6    dp1  c4:f5:7c:88:dc:a2
7/ge7    dp0  c4:f5:7c:88:dc:a3
7/ge7    dp1  c4:f5:7c:88:dc:a3
7/ge8    dp0  c4:f5:7c:88:dc:a4
7/ge8    dp1  c4:f5:7c:88:dc:a4
7/ge9    dp0  c4:f5:7c:88:dc:a5
7/ge9    dp1  c4:f5:7c:88:dc:a5
7/ge10   dp0  c4:f5:7c:88:dc:a6
7/ge10   dp1  c4:f5:7c:88:dc:a6
7/ge11   dp0  c4:f5:7c:88:dc:a7
7/ge11   dp1  c4:f5:7c:88:dc:a7
7/ge12   dp0  c4:f5:7c:88:dc:a8
7/ge12   dp1  c4:f5:7c:88:dc:a8
7/ge13   dp0  c4:f5:7c:88:dc:a9
7/ge13   dp1  c4:f5:7c:88:dc:a9
7/ge14   dp0  c4:f5:7c:88:dc:aa
7/ge14   dp1  c4:f5:7c:88:dc:aa
7/ge15   dp0  c4:f5:7c:88:dc:ab
7/ge15   dp1  c4:f5:7c:88:dc:ab
7/ge16   dp0  c4:f5:7c:88:dc:ac
7/ge16   dp1  c4:f5:7c:88:dc:ac
7/ge17   dp0  c4:f5:7c:88:dc:ad
7/ge17   dp1  c4:f5:7c:88:dc:ad
7/lan    dp0  00:00:00:00:00:00
7/lan    dp1  00:00:00:00:00:00
7/lan1   dp0  00:00:00:00:00:00
7/lan1   dp1  00:00:00:00:00:00
7/ha0    dp0  00:00:00:00:00:01
7/ha0    dp1  00:00:00:00:00:01
-------------------------------------------------------------------------
```

The following example lists the MAC addresses of the LAN ports on slot 7 in the switch:

NOTE
LAN MAC addresses are all zeros because Brocade SX6 Extension Blade in slot 7 is in FCIP-only mode.

```
switch:admin> portcfgge lan --show -lmac

 Port       dpid    MAC Address
-------------------------------------------------------------------------
 lan        dp0     c4:f5:7c:bc:c1:62
 lan        dp1     c4:f5:7c:bc:c1:76
------------------------
```

The following example lists the MAC addresses of interface GE4:

```
switch:admin> portcfgge ge4 --show -lmac

 Port       dpid    MAC Address
```

```
------------------------------------------------------------------------
 ge4        dp0     c4:f5:7c:bc:c1:54
 ge4        dp1     c4:f5:7c:bc:c1:54
------------------------------------------------------------------------
```

Displaying LAG (Portchannel) Information

Use the `portchannel --show` command to view portchannel (LAG) information.

NOTE
The Brocade 7850 Extension Switch and the Brocade SX6 Extension Blade support a maximum of four portchannels, static and dynamic, whereas the Brocade 7810 Extension Switch supports a maximum of two portchannels. The following example displays LAG information for both static and dynamic LAGs:

```
switch:admin> portchannel --show -all

Name                    Type        Oper-State   Port-Count   Member Ports
--------------------------------------------------------------------------------
MyLAG1                  Dynamic     Online       2            ge0 ,ge1*
MyLAG2                  Static      Online       2            ge4 ,ge5
```

You can display more detailed information with the `-detail` option.

```
You can display more detailed information with the detail option.

switch:admin> portchannel --show -detail

 Name :MyLAG1
 Type :Dynamic
 Key  :784
 Speed  :10G
 Admin-state: Enable
 Oper-state : Online
  Admin Key: 0784 - Oper Key 0784
  LACP System ID: 0x8000,c4-f5-7c-cb-c1-8a
  PART System ID: 0x8000,00-be-75-79-38-40
  Portchannel Member count = 2
  Port          Oper state      Sync    Timeout         Auto-Negotiation
  ------------------------------------------------------------------
   ge4          Online          1       Long            Disabled
  *ge5          Online          1       Long            Disabled
```

Displaying eHCL HA Information

Use the `portshow fciptunnel --hcl-status` command to view the status of a correctly configured eHCL. The following example shows the tunnel eHCL status:

```
switch:admin> portshow fciptunnel --hcl-status

Checking FCIP Tunnel HA Status.
Current Status : Ready
CP Version : v9.2.0
DP0 Status:
State : Online - Inactive
Version : v9.2.0
```

```
Current FC HA Stage : IDLE
Current IP HA Stage : IDLE
IP SVI Swapped : NO
DP COMM Status : UP 7/DP1 Status:
DP1 Status:
State : Online - Inactive
Version : v9.2.0
Current FC HA Stage : IDLE
Current IP HA Stage : IDLE
IP SVI Swapped : NO
DP COMM Status : UP
Tunnel 24 (FID:128) FC:HA Ready   IP:Disabled   - FC traffic will be disrupted.
Tunnel 25 (FID:128) FC:HA Online  IP:Disabled   - Traffic will not be disrupted.
Tunnel 26 (FID:128) FC:HA Offline IP:HA Offline - FC and IP traffic will be disrupted.
Tunnel 33 (FID:128) FC:HA Ready   IP:Disabled   - FC traffic will be disrupted.
Tunnel 34 (FID:128) FC:HA Online  IP:HA Online  - Traffic will not be disrupted.
Tunnel 35 (FID:128) FC:HA Ready   IP:Disabled   - FC traffic will be disrupted.
```

Displaying IP Extension TCL Information

Use the `portshow tcl` command to configure and view the traffic control list (TCL) in the hybrid mode.

NOTE
The Brocade 7850 Extension Switch and the Brocade SX6 Extension Blade support a maximum of 1024 defined and 128 active TCL sessions, whereas the Brocade 7810 Extension Switch supports a maximum of 256 defined and 32 active TCL sessions.

Portshow tcl options:

```
portshow tcl [<name>] [<option>]
Options:
    -s,--summary            Displays summary view of TCLs.
    -d,--detail             Displays detailed view of TCLs.
    -p,--priority <pri>     Displays TCL(s) matching the specified priority
    -S,--sort <sort>        Sorts the TCL list based on specified sort field.
                            sort=[name|priority|src-addr|dst-addr]
       --filter <args>      Limit the output to specific filter criteria.
                            Use portShow tcl --filter -help for details.
```

The following example displays the TCL configuration. A summary table shows the current TCL consumption on a per-DP basis. The output is truncated.

```
switch:admin> portshow tcl

 Pri    Name                 Flgs   Target     L2COS  VLAN DSCP Proto Port Hit
                             Src-Addr                 Dst-Addr
-------------------------------------------------------------------------------
 *100   PBR_ICMP             AI---  24-Med     ANY    ANY  ANY  ICMP  ANY  71230
                             172.16.250.160/29        172.16.254.160/29

 *200   PDX_75_to_SJC_187    AI---  24-Med     ANY    ANY  ANY  ANY   ANY  17084
                             10.75.0.0/16             10.187.0.0/16

 *300   PDX_17_to_LAX_34     AI---  33-Med     ANY    ANY  ANY  ANY   ANY  37529
                             10.17.0.0/16             10.34.0.0/16
```

```
--------------------------------------------------------------------------------
Flags: *=Enabled ..=Name Truncated (see --detail for full name)
       A=Allow D=Deny I=IP-Ext P=Segment Preservation
       R=End-to-End RST Propagation N=Non Terminated.

Active TCL Limits:   Cur / Max
-------------------------------------------
  DP0                32 / 128
  DP1                21 / 128
-------------------------------------------
Configured Total:    52 / 1024
```

You can use the `--detail` option to display more detailed information.

```
switch:admin> portshow tcl --detail

TCL: PBR_ICMP
================================================================
 Admin Status:           Enabled
 Priority:               100
 Target:                 24-Medium (tid:175)
 VLAN:                   ANY
 L2COS:                  ANY
 DSCP:                   ANY
 Source Address:         172.16.250.160/29
 Destination Address:    172.16.254.160/29
 L4 Protocol:            ICMP
 Protocol Port:          ANY
 Action:                 Allow DP0
 RST Propagation:        Disabled
 Segment Preservation:   Disabled
 Non Terminated:         Disabled
 Cfgmask:                0x081fc3e67
 Hit Count:              0

TCL: PDX_75_to_SJC_187
================================================================
 Admin Status:           Enabled
 Priority:               200
 Target:                 24-Medium (tid:5)
 VLAN:                   ANY
 L2COS:                  ANY
 DSCP:                   ANY
 Source Address:         10.75.0.0/16
 Destination Address:    10.187.0.0/16
 L4 Protocol:            ANY
 Protocol Port:          ANY
 Action:                 Allow DP0
 RST Propagation:        Disabled
 Segment Preservation:   Disabled
 Non Terminated:         Disabled
 Cfgmask:                0x081fc3e27
 Hit Count:              170905204
```

```
TCL: PDX_17_to_LAX_34
================================================================
 Admin Status:           Enabled
 Priority:               300
 Target:                 33-Medium (tid:152)
 VLAN:                   ANY
 L2COS:                  ANY
 DSCP:                   ANY
 Source Address:         10.17.0.0/16
 Destination Address:    10.34.0.0/16
 L4 Protocol:            ANY
 Protocol Port:          ANY
 Action:                 Allow DP1
 RST Propagation:        Disabled
 Segment Preservation:   Disabled
 Non Terminated:         Disabled
 Cfgmask:                0x081fc3e27
 Hit Count:              37529
```

You can sort the output by using the `--sort` option. The following example sorts the configured TCLs by name:

```
switch:admin> portshow tcl --sort name
```

Use the `help` option to show all available options.

```
switch:admin> portshow tcl --help
Usage:
portshow tcl [<name>] [<option>]

Options:

     -s,--summary          Displays summary view of TCLs.
     -d,--detail           Displays detailed view of TCLs.
     -p,--priority <pri>   Displays TCL(s) matching the specified priority
     -S,--sort <sort>      Sorts the TCL list based on specified sort field.
                           sort=[name|priority|src-addr|dst-addr]
        --filter <args>    Limit the output to specific filter criteria.
                           Use portShow tcl --filter -help for details.
```

Displaying IP Extension LAN Statistics

Use the `portshow lan-stats` command to view the IP Extension LAN statistics in hybrid mode. The following example displays the global LAN statistics for a Brocade 7810 Extension Switch:

```
switch:admin> portshow lan-stats --global
```

For more information about displaying IP Extension LAN statistics, see Using IP Extension Flow Monitor.

Displaying Performance Statistics

Use the `portshow fciptunnel` command to view a summary of performance statistics for all tunnels and circuits. The following command displays a summary of performance statistics for all tunnels and circuits.

```
switch:admin> portshow fciptunnel --circuit --perf --summary
```

You can display a summary of performance statistics for the logical switch and circuits using the `portshow fciptunnel` command.

```
switch:admin> portshow fciptunnel --all-ls --circuit --perf --summary
```

Displaying QoS Statistics

Use the `portshow fciptunnel --qos` command to view QoS statistics for all tunnels.

```
switch:admin> portshow fciptunnel --qos --summary
```

Displaying Configuration Details

Use the `portshow fciptunnel --detail` command to view all configuration details.

```
switch:admin> portshow fciptunnel --detail
```

Filtering portshow Output

Use the `portshow <option> --filter` command to filter the display output by creating filter names or by specifying information that you want to show or hide. The filter rules can contain any combination of IP addresses, TCP ports, GE ports, Tunnel IDs, and other counters/IDs.

Portshow <option> --filter -help:

```
Usage: portshow <option> [<args>] --filter [<filter_args>]

Options:
  ipif                - Show list of IP addresses.
  iproute             - Show list of IP routes.
  arp                 - Show list of ARP / Neighbor Discovery entries.
  fciptunnel          - Show list of Tunnel entries.
  fcipcircuit         - Show list of Circuit entries.
  tcl                 - Show list of TCL entries.
  ipsec-policy        - Show list of IPSec Policies.
  lan-stats           - Show list of LAN stats entries.
  --wtool <arg> show  - Show list of WTool sessions.

Args: optional arguments for specified <option>

Filter_Args:
   --set <name>               - Specify the filter-set name to use.
   --port [<slot>/]<port>     - Filter on specified FC/VE/GE port.
   --slot <slot>              - Filter on specified slot number.
   --ipaddr <addr>[/<pfx>]    - Filter on specified an IP address / network.
   --dp [<slot>/]dp<#>        - Specify a DP ID to filter on.
   --circuit <cid>            - Specify a Circuit ID to filter on.
   --priority <pri>           - Specify a priority to filter on. Valid
                                priorities
                                [control|high|medium|low|IPHigh|IPMedium|IPLow].
   --ha-type <type>           - Specify an HA type to filter on. Valid ha-type:
                                [main|local-backup|remote-backup].
   --tcp-port <1-65535|app>   - Specify a TCP port/range/application type to
                                filter on (see 'portshow lan-stats --known-apps'
                                for list of application types).
```

```
    --retransmits <value>      - Filter on retransmits greater than or equal to
                                 <value>.
    --rtt <value>              - Filter on round-trip-time greater than or equal
                                 to <value>.
    --bytes <bytes>[k|m|g]     - Filter on bytes greater than or equal to
                                 <bytes>.
    --conn-cnt <value>         - Filter on connect-count greater than or equal to
                                 <value>.
    --vlan <vlanid>            - Filter on specified VLAN-ID.
    --oper-status <oper>       - Filter based on obj oper-status. Can specify
                                 exact oper-string or specify
                                 active|inactive|healthy|unhealthy states.
    --default-behavior show|hide - Set the default behavior for objects that
                                 don't support the specified filter criteria.
                                 Default:hide
    --show                     - Show objects matching the filter criteria.
    --hide                     - Hide objects matching the filter criteria.
    --and                      - Use 'AND' operation when using multiple filter
                                 criteria.
    --or                       - Use 'OR' operation when using multiple filter
                                 criteria.
    --help                     - Display this usage message.
```

The following example displays only per-flow stats that use both IP address 192.168.0.10 and TCP port 336:

```
switch:admin> portshow lan-stats --per-flow -all \
--filter \
-ipaddr 192.168.0.10 \
-and \
-tcp-port 336
```

The following example creates a filter set called tcpErrors that looks for any retransmits greater than 100 and a byte count greater than 1,000,000 bytes. The output shows objects that are moving traffic.

```
switch:admin> portcfg filter-set tcpErrors create --retransmits 100 --and --bytes 1000000
Operation Succeeded

switch:admin> portshow filter-set

Name         ACT/DEF Filter Statement
-------------------------------------------------------------
tcpErrors    SHOW/HIDE (retx:100 && bytes:1000000)
-------------------------------------------------------------
ACT: Action for objects matching filter
DEF: Default behavior for objects where filter doesn't apply
```

The filters of the `portshow` command provide built-in filter sets. The following example shows how to use the `--ipaddr` filter to show a specific IPIF address:

```
switch:admin> portshow ipif

Port        IP Address      / Pfx   MTU    VLAN   Flags
-----------------------------------------------------------------------------
4/ge0.dp0   192.168.20.10   / 24    1500   0      U R M
4/ge0.dp0   192.168.20.11   / 24    1280   0      U R M
4/ge0.dp0   192.168.20.12   / 24    1350   0      U R M
```

```
4/ge0.dp0   192.168.20.13   / 24    1500   0      U R M
4/ge0.dp0   192.168.20.14   / 24    3000   0      U R M
4/ge0.dp0   192.168.20.15   / 24    6000   0      U R M
4/ge0.dp0   192.168.20.16   / 24    9000   0      U R M
4/ge0.dp0   192.168.20.17   / 24    9216   0      U R M
-----------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport

switch:admin> portshow ipif --filter --ipaddr 192.168.20.11

Port        IP Address            / Pfx  MTU   VLAN  Flags
-----------------------------------------------------------------------------
4/ge0.dp0   192.168.20.11         / 24   1280  0     U R M
-----------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

For more information about the `portshow filter` command, refer to the *Brocade Fabric OS Command Reference Manual*.

Displaying Tunnel Status

Use the `portshow fciptunnel` command to view the tunnel status. The following example shows the status of the tunnel and its circuits:

```
switch:admin> portshow fciptunnel -c


Tunnel Circuit  OpStatus Flags    Uptime   TxMBps    RxMBps  ConnCnt CommRt    Met/G
-------------------------------------------------------------------------------------
12     -        Up       --if---d- 1d1h44m  0.00      0.00     2       -          -
12     0 ge2    Up       ----a--i4 1d1h44m  0.00      0.00     2      1250/1250  0/-
12     1 ge3    Up       ----a--i4 1d1h44m  0.00      0.00     2      1250/1250  0/-
-------------------------------------------------------------------------------------
Flags (tunnel): l=Legacy QOS Mode
                i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                I=IP-Ext
     (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                ARL a=Auto r=Reset s=StepDown t=TimedStepDown S=SLA
```

Displaying Tunnel Information

Use the `portshow fciptunnel` command to view the tunnel and circuit information. The following example displays general tunnel information associated with VE12:

```
switch:admin> portshow fciptunnel 12

 Tunnel: VE-Port:12 (idx:0, DP0)
 ================================================
  Oper State           : Online
  TID                  : 12
  Flags                : 0x00000000
```

```
  IP-Extension         : Disabled
  Compression          : Deflate
  QoS BW Ratio         : 50% / 30% / 20%
  Fastwrite            : Enabled
  Tape Pipelining      : Disabled
  IPSec                : Enabled
  IPSec-Policy         : MyIPsec
  Legacy QOS Mode      : Disabled
  Load-Level (Cfg/Peer): Failover (Failover / Failover)
  Local WWN            : 10:00:88:94:71:60:99:7e
  Peer WWN             : 10:00:88:94:71:60:97:5e
  RemWWN (config)      : 00:00:00:00:00:00:00:00
  Peer Platform        : 7810
  cfgmask              : 0x0000001f 0x4000020c
  Uncomp/Comp Bytes    : 8005081490104 / 4258266530693 / 1.88 : 1
  Uncomp/Comp Byte(30s): 0 / 0 / 1.00 : 1
  Flow Status          : 0
  ConCount/Duration    : 2 / 1d1h51m
  Uptime               : 1d1h46m
  Stats Duration       : 1d1h46m
  Receiver Stats       : 19210710502 bytes / 109883939 pkts /   65.00 Bps Avg
  Sender Stats         : 4258272614781 bytes / 293658095 pkts /   58.00 Bps Avg
  TCP Bytes In/Out     : 44405519066 / 4299972988456
  ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
  RTT (min/avg/max)    : 20 / 20 / 21 ms
  Wan Util             : 0.0%
  TxQ Util             : 0.0%
```

Displaying a Tunnel with Circuit Information

Use the `portshow fciptunnel -c` command to view the tunnel and its circuits. The following example shows how to use the `-c` option to add the circuit information to the `portshow fciptunnel` command output:

```
switch:admin> portshow fciptunnel 12 -c

 Tunnel: VE-Port:12 (idx:0, DP0)
 ================================================
  Oper State           : Online
  TID                  : 12
  Flags                : 0x00000000
  IP-Extension         : Disabled
  Compression          : Deflate
  QoS BW Ratio         : 50% / 30% / 20%
  Fastwrite            : Enabled
  Tape Pipelining      : Disabled
  IPSec                : Enabled
  IPSec-Policy         : MyIPsec
  Legacy QOS Mode      : Disabled
  Load-Level (Cfg/Peer): Failover (Failover / Failover)
  Local WWN            : 10:00:88:94:71:60:99:7e
  Peer WWN             : 10:00:88:94:71:60:97:5e
  RemWWN (config)      : 00:00:00:00:00:00:00:00
  Peer Platform        : 7810
```

```
 cfgmask              : 0x0000001f 0x4000020c
 Uncomp/Comp Bytes    : 8005081498552 / 4258266539141 / 1.88 : 1
 Uncomp/Comp Byte(30s): 0 / 0 / 1.00 : 1
 Flow Status          : 0
 ConCount/Duration    : 2 / 1d1h52m
 Uptime               : 1d1h47m
 Stats Duration       : 1d1h47m
 Receiver Stats       : 19210732714 bytes / 109884037 pkts /   59.00 Bps Avg
 Sender Stats         : 4258272625377 bytes / 293658149 pkts /   58.00 Bps Avg
 TCP Bytes In/Out     : 44406192038 / 4299973827228
 ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
 RTT (min/avg/max)    : 20 / 20 / 21 ms
 Wan Util             : 0.0%
 TxQ Util             : 0.0%

Circuit 12.0 (DP0)
==================================================
 Admin/Oper State     : Enabled / Online
 Flags                : 0x00000000
 IP Addr (L/R)        : 192.168.1.10 ge2 <-> 192.168.1.20
 HA IP Addr (L/R)     : 0.0.0.0  <-> 0.0.0.0
 Configured Comm Rates: 1250000 / 1250000 kbps
 Peer Comm Rates      : 1250000 / 1250000 kbps
 Actual Comm Rates    : 1250000 / 1250000 kbps
 Keepalive (Cfg/Peer) : 1000 (1000 / 1000) ms
 Metric               : 0
 Connection Type      : Default
 ARL-Type             : Auto
 PMTU                 : Disabled
 HA PMTU              : Disabled
 SLA                  : (none)
 Failover Group       : 0
 VLAN-ID              : NONE
 L2Cos (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
 L2Cos (IP:h/m/l)     : 0 / 0 / 0
 DSCP (FC:h/m/l)      : 0 / 0 / 0 (Ctrl:0)
 DSCP (IP:h/m/l)      : 0 / 0 / 0
 cfgmask              : 0x40000000 0x00010caf
 Flow Status          : 0
 ConCount/Duration    : 2 / 1d1h52m
 Uptime               : 1d1h47m
 Stats Duration       : 1d1h47m
 Receiver Stats       : 9607825198 bytes / 54957535 pkts /   48.00 Bps Avg
 Sender Stats         : 2129142261302 bytes / 146817826 pkts /   10.00 Bps Avg
 TCP Bytes In/Out     : 22383428298 / 2195226412420
 ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
 RTT (min/avg/max)    : 20 / 20 / 21 ms
 Wan Util             : 0.0%

Circuit 12.1 (DP0)
==================================================
 Admin/Oper State     : Enabled / Online
 Flags                : 0x00000000
```

```
  IP Addr (L/R)         : 192.168.1.11 ge3 <-> 192.168.1.21
  HA IP Addr (L/R)      : 0.0.0.0  <-> 0.0.0.0
  Configured Comm Rates: 1250000 / 1250000 kbps
  Peer Comm Rates       : 1250000 / 1250000 kbps
  Actual Comm Rates     : 1250000 / 1250000 kbps
  Keepalive (Cfg/Peer) : 1000 (1000 / 1000) ms
  Metric                : 0
  Connection Type       : Default
  ARL-Type              : Auto
  PMTU                  : Disabled
  HA PMTU               : Disabled
  SLA                   : (none)
  Failover Group        : 0
  VLAN-ID               : NONE
  L2Cos (FC:h/m/l)      : 0 / 0 / 0 (Ctrl:0)
  L2Cos (IP:h/m/l)      : 0 / 0 / 0
  DSCP (FC:h/m/l)       : 0 / 0 / 0 (Ctrl:0)
  DSCP (IP:h/m/l)       : 0 / 0 / 0
  cfgmask               : 0x40000000 0x00010caf
  Flow Status           : 0
  ConCount/Duration     : 2 / 1d1h52m
  Uptime                : 1d1h47m
  Stats Duration        : 1d1h47m
  Receiver Stats        : 9602907516 bytes / 54926502 pkts /   11.00 Bps Avg
  Sender Stats          : 2129130364075 bytes / 146840323 pkts /   47.00 Bps Avg
  TCP Bytes In/Out      : 22411554984 / 2195212431308
  ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
  RTT (min/avg/max)     : 20 / 20 / 21 ms
  Wan Util              : 0.0%
```

Displaying Tunnel Performance

Use the `portshow fciptunnel` command with the `--perf` option to display the performance statistics for a tunnel. The following example displays performance statistics for a tunnel that is associated with VE12:

```
switch:admin> portshow fciptunnel 12 --perf

 Tunnel: VE-Port:12 (idx:0, DP0)
 ================================================
  Oper State            : Online
  TID                   : 12
  Flags                 : 0x00000000
  IP-Extension          : Disabled
  Compression           : Deflate
  QoS BW Ratio          : 50% / 30% / 20%
  Fastwrite             : Enabled
  Tape Pipelining       : Disabled
  IPSec                 : Enabled
  IPSec-Policy          : MyIPsec
  Legacy QOS Mode       : Disabled
  Load-Level (Cfg/Peer): Failover (Failover / Failover)
  Local WWN             : 10:00:88:94:71:60:99:7e
  Peer WWN              : 10:00:88:94:71:60:97:5e
```

```
  RemWWN (config)      : 00:00:00:00:00:00:00:00
  Peer Platform        : 7810
  cfgmask              : 0x0000001f 0x4000020c
  Uncomp/Comp Bytes    : 8005081498552 / 4258266539141 / 1.88 : 1
  Uncomp/Comp Byte(30s): 0 / 0 / 1.00 : 1
  Flow Status          : 0
  ConCount/Duration    : 2 / 1d1h53m
  Uptime               : 1d1h48m
  Stats Duration       : 1d1h48m
  Receiver Stats       : 19210738534 bytes / 109884059 pkts /   73.00 Bps Avg
  Sender Stats         : 4258272630789 bytes / 293658171 pkts /   66.00 Bps Avg
  TCP Bytes In/Out     : 44407260010 / 4299975164240
  ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
  RTT (min/avg/max)    : 20 / 20 / 21 ms
  Wan Util             : 0.0%
  TxQ Util             : 0.0%
```

Displaying TCP Statistics

Use the `portshow fcipcircuit` command with the `--tcp` option to view TCP connection statistics for circuits. The following example displays TCP connection statistics for circuits associated with VE12:

```
switch:admin> portshow fcipcircuit 12 --tcp

 Circuit 12.0 (DP0)
 ==================================================
  Admin/Oper State     : Enabled / Online
  Flags                : 0x00000000
  IP Addr (L/R)        : 192.168.1.10 ge2 <-> 192.168.1.20
  HA IP Addr (L/R)     : 0.0.0.0  <-> 0.0.0.0
  Configured Comm Rates: 1250000 / 1250000 kbps
  Peer Comm Rates      : 1250000 / 1250000 kbps
  Actual Comm Rates    : 1250000 / 1250000 kbps
  Keepalive (Cfg/Peer) : 1000 (1000 / 1000) ms
  Metric               : 0
  Connection Type      : Default
  ARL-Type             : Auto
  PMTU                 : Disabled
  HA PMTU              : Disabled
  SLA                  : (none)
  Failover Group       : 0
  VLAN-ID              : NONE
  L2Cos (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
  L2Cos (IP:h/m/l)     : 0 / 0 / 0
  DSCP (FC:h/m/l)      : 0 / 0 / 0 (Ctrl:0)
  DSCP (IP:h/m/l)      : 0 / 0 / 0
  cfgmask              : 0x40000000 0x00010caf
  Flow Status          : 0
  ConCount/Duration    : 2 / 1d2h4m
  Uptime               : 1d1h59m
  Stats Duration       : 1d1h59m
  Receiver Stats       : 9607869778 bytes / 54957723 pkts /   47.00 Bps Avg
  Sender Stats         : 2129142294610 bytes / 146817971 pkts /   10.00 Bps Avg
```

```
TCP Bytes In/Out     : 22388309318 / 2195232509032
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)    : 20 / 20 / 21 ms
Wan Util             : 0.0%

TCP Connection 12.0 HA-Type:Main Pri:Low Conn:0x010a2b2e
================================================
 Local / Remote Port      : 3225 / 55992
 Duration                 : 1d1h59m
 MSS                      : 9074 bytes
 ARL Min / Cur / Max      : 125000 / 125000 / 625000
 ARL Reset Algo           : StepDown
 Send Window
 Size / Scale             : 15624192 / 9
 Slow Start Threshold     : 16777216
 Congestion Window        : 18201520
 Pkts InFlight            : 0
 Recv Window
 Size / Scale             : 15624704 (Max:15624704) / 9
 SendQ Nxt / Min / Max    : 0xc9e0436b / 0xc9e0436b / 0xc9e0436b
 RecvQ Nxt / Min / Max    : 0xa272d3bd / 0xa272d3bd / 0xa3613d6d
 RecvQ Pkts               : 0
 Sender Stats
 Sent Bytes / Pkts        : 93737308 / 711031
 Unacked Data             : 0
 Retransmits Slow / Fast  : 0 / 0 (High:0)
 SlowStart                : 0
 Receiver Stats
 Recv Bytes / Pkts        : 73905712 / 711034
 Out-of-Order             : 0 (High:0)
 Duplicate ACKs           : 0
 RTT / Variance (High)    : 19 ms (20 ms) / 0 ms (0 ms)

  (Output truncated…)
```

The counters can be reset to zero with the `--tcp --reset` arguments.

```
switch:admin> portshow fciptunnel 12 --tcp --reset
```

You can display the entire lifetime of statistics for the tunnel using the following command. The time basis for the statistics display in the output.

```
switch:admin> portshow fciptunnel 12 --tcp --lifetime
```

Displaying Circuits

Use the `portshow fcipcircuit --all-ls` command to view circuit information for all VFs. The following example displays circuit information for all VFs:

```
switch:admin> portshow fcipcircuit --all-ls
```

Displaying a Single Circuit

Use the `portshow fcipcircuit` command with the `<slot> <ve>` option to view information for a specific circuit. The following example displays information for circuit-1 on VE12:

```
switch:admin> portshow fcipcircuit 12 1
```

Displaying Circuit Performance

Use the `portshow fcipcircuit` command with the `--perf` option to view the circuit performance statistics. The following example displays circuit performance information for circuit 1 on slot 3, tunnel 12:

```
switch:admin> portshow fcipcircuit 3/12 1 --perf
```

Displaying GE Interface Performance

Use the `geportperfshow` command to view throughput information for Gigabit Ethernet (GbE) ports. The following example displays a GbE throughput port performance. The display updates until you press the **Enter** key. Use the `--help` option to show available command options.

```
switch:admin> geportperfshow

Throughput of GE port
ge0     ge1     ge2     ge3     ge4     ge5     ge6     ge7     Total
================================================================================
  0       0  11.36k  21.98k       0       0       0       0     33.34k
ge0     ge1     ge2     ge3     ge4     ge5     ge6     ge7     Total
================================================================================
  0       0  17.50k  25.94k       0       0       0       0     43.44k
ge0     ge1     ge2     ge3     ge4     ge5     ge6     ge7     Total
================================================================================
  0       0  10.22k  21.25k       0       0       0       0     31.47k
ge0     ge1     ge2     ge3     ge4     ge5     ge6     ge7     Total
================================================================================
  0       0  10.63k  20.84k       0       0       0       0     31.47k
ge0     ge1     ge2     ge3     ge4     ge5     ge6     ge7     Total
================================================================================
  0       0  18.27k  21.72k       0       0       0       0     40.00k
```

Displaying QoS Prioritization for a Circuit

Use the `portshow fcipcircuit` command with the `--perf --qos` option to view QoS prioritization for the default circuit. The following example displays QoS prioritization for slot-7, tunnel-24, circuit-0. The QoS performance statistics are shown for all QoS tunnels. The following example shows a FCIP-only tunnel, no IP Extension and the Control, High, Medium, and Low statistics:

```
switch:admin> portshow fcipcircuit 7/24 0 --perf --qos

Circuit 7/24.0 (DP0)
==================================================
Admin/Oper State       : Enabled / Online
Priority               : Control
Flags                  : 0x00000000
IP Addr (L/R)          : 192.168.5.2  7/ge2 <-> 192.168.1.2
HA IP Addr (L/R)       : 192.168.5.12 7/ge2 <-> 192.168.1.12
Configured Comm Rates  : 0 / 5000000 kbps
Peer Comm Rates        : 0 / 5000000 kbps
Actual Comm Rates      : 0 / 5000000 kbps
Keepalive (Cfg/Peer)   : 6000 (6000 / 6000) ms
```

```
Metric                  : 0
Connection Type         : Default
ARL-Type                : Auto
PMTU                    : Disabled
HA PMTU                 : Disabled
SLA                     : (none)
IPSEC                   : Enabled (0)
Failover Group          : 0
VLAN-ID                 : NONE
L2Cos (Ctrl)            : 0
 DSCP (Ctrl)            : 0
Cfgmask                 : 0x40000000 0x00010c2f
Flow Status             : 1
ConCount/Duration       : 1 / 5d12m
Uptime                  : 5d12m
Stats Duration          : 5d12m
Receiver Stats          : 0 bytes / 0 pkts /  0.00 Bps Avg
Sender Stats            : 0 bytes / 0 pkts /  0.00 Bps Avg
TCP Bytes In/Out        : 188332 / 158164
ReTx/OOO/SloSt/DupAck   : 0 / 0 / 0 / 0
RTT (min/avg/max)       : 1 / 1 / 4 ms
Wan Util (low/high)     : 0.0% / 0.0%

Circuit 7/24.0 (DP0)
==================================================
Admin/Oper State        : Enabled / Online
Priority                : FC-High
Flags                   : 0x00000000
IP Addr (L/R)           : 192.168.5.2  7/ge2 <-> 192.168.1.2
HA IP Addr (L/R)        : 192.168.5.12 7/ge2 <-> 192.168.1.12
Configured Comm Rates   : 2500000 / 5000000 kbps
Peer Comm Rates         : 2500000 / 5000000 kbps
Actual Comm Rates       : 2500000 / 5000000 kbps
Keepalive (Cfg/Peer)    : 6000 (6000 / 6000) ms
Metric                  : 0
Connection Type         : Default
ARL-Type                : Auto
PMTU                    : Disabled
SLA                     : (none)
IPSEC                   : Disabled (0)
Failover Group          : 0
VLAN-ID                 : NONE
L2Cos (FC-High)         : 0
 DSCP (FC-High)         : 0
cfgmask                 : 0x40000000 0x00213def
Flow Status             : 0
ConCount/Duration       : 1 / 5d12m
Uptime                  : 5d12m
Stats Duration          : 5d12m
Receiver Stats          : 0 bytes / 0 pkts /  0.00 Bps Avg
Sender Stats            : 0 bytes / 0 pkts /  0.00 Bps Avg
TCP Bytes In/Out        : 1083588332 / 1083589164
ReTx/OOO/SloSt/DupAck   : 0 / 0 / 0 / 0
```

```
RTT (min/avg/max)       : 1 / 1 / 6 ms
Wan Util (low/high)     : 0.0% / 0.0%

Circuit 7/24.0 (DP0)
==================================================
Admin/Oper State        : Enabled / Online Warning
Priority                : FC-Medium
Flags                   : 0x00000000
IP Addr (L/R)           : 192.168.5.2  7/ge2 <-> 192.168.1.2
HA IP Addr (L/R)        : 192.168.5.12 7/ge2 <-> 192.168.1.12
Configured Comm Rates   : 1500000 / 5000000 kbps
Peer Comm Rates         : 2000000 / 5000000 kbps
Actual Comm Rates       : 1500000 / 5000000 kbps
Keepalive (Cfg/Peer)    : 6000 (6000 / 6000) ms
Metric                  : 0
Connection Type         : Default
ARL-Type                : Auto
PMTU                    : Disabled
SLA                     : (none)
IPSEC                   : Disabled (0)
Failover Group          : 0
VLAN-ID                 : NONE
L2Cos (FC-Medium)       : 0
 DSCP (FC-Medium)       : 0
cfgmask                 : 0x40000000 0x00213def
Configuration Warnings :
  Min-comm-rate / QoS-Ratio / Dist-Ratio
Flow Status             : 0
ConCount/Duration       : 1 / 5d12m
Uptime                  : 5d12m
Stats Duration          : 5d12m
Receiver Stats          : 0 bytes / 0 pkts /  0.00 Bps Avg
Sender Stats            : 0 bytes / 0 pkts /  0.00 Bps Avg
TCP Bytes In/Out        : 1083589204 / 1083588460
ReTx/OOO/SloSt/DupAck   : 0 / 0 / 0 / 0
RTT (min/avg/max)       : 1 / 1 / 4 ms
Wan Util (low/high)     : 0.0% / 0.0%

Circuit 24.0 (DP0)
==================================================
Admin/Oper State        : Enabled / Online Warning
Priority                : FC-Low
Flags                   : 0x00000000
IP Addr (L/R)           : 192.168.5.2  7/ge2 <-> 192.168.1.2
HA IP Addr (L/R)        : 192.168.5.12 7/ge2 <-> 192.168.1.12
Configured Comm Rates   : 1000000 / 5000000 kbps
Peer Comm Rates         : 500000 / 5000000 kbps
Actual Comm Rates       : 500000 / 5000000 kbps
Keepalive (Cfg/Peer)    : 6000 (6000 / 6000) ms
Metric                  : 0
Connection Type         : Default
ARL-Type                : Auto
PMTU                    : Disabled
```

```
SLA                    : (none)
IPSEC                  : Disabled (0)
Failover Group         : 0
VLAN-ID                : NONE
L2Cos (FC-Low)         : 0
 DSCP (FC-Low)         : 0
cfgmask                : 0x40000000 0x00213def
Configuration Warnings:
  Min-comm-rate / QoS-Ratio / Dist-Ratio
Flow Status            : 0
ConCount/Duration      : 1 / 5d12m
Uptime                 : 5d12m
Stats Duration         : 5d12m
Receiver Stats         : 0 bytes / 0 pkts /  0.00 Bps Avg
Sender Stats           : 0 bytes / 0 pkts /  0.00 Bps Avg
TCP Bytes In/Out       : 1083583264 / 1083583984
ReTx/OOO/SloSt/DupAck  : 0 / 0 / 0 / 0
RTT (min/avg/max)      : 1 / 1 / 1 ms
Wan Util (low/high)    : 0.0% / 0.0%
```

Troubleshooting

This section describes WAN analysis tools that are designed to test connections, trace routes, and estimate the end-to-end IP path performance characteristics between a pair of Brocade Extension port endpoints.

These tools are available as options for the `portCmd` command.

- `portCmd --ping` – Tests connections between a local IP address and a destination IP address.
- `portCmd --traceroute` – Traces the IP network path from a local IP address to a destination IP address.
- `portCmd --pmtu` – Helps in identifying the maximum MTU possible in the WAN network.
- `portCmd --wtool` – Generates traffic over a pair of IP addresses to test link characteristics.
- `portShow fcipTunnel --perf` – Displays performance statistics for existing and established VE_Ports.

Portcmd --ping

The `portCmd --ping` command tests the connection between an IPIF IP address of a local GE interface and a destination IP address.

The `portCmd --ping` command has the following syntax:

```
portCmd --ping [slot/ge_port] dp0|1 -s <Source IP Address> -d <Destination IP Address>
```

As DP complexes share GE interfaces, the interface is identified as ge[X].dp[Y]. For example, `ge[X].dp[Y]` . This identification directs the command to a specific DP. For more information of command syntax and output examples, refer to the *Brocade Fabric OS Command Reference Manual*.

Portcmd --traceroute

The `portCmd traceroute` command traces routes from a local Ethernet port to a destination IP address. The `portCmd --traceroute` command has the following syntax:

```
portCmd --traceroute [slot/ge_port] dp0|1 -s <Source IP Address> -d <Destination IP Address>
```

DPs share GE interfaces, and the port identification is `ge[X].dp[Y]` . This port identifier directs the command to a specific DP. For example, `ge2.dp0` . For more information about command syntax and output examples, refer to the *Brocade Fabric OS Command Reference Manual*.

Portcmd --wtool

The WAN Tool allows you to generate traffic at a specified rate in kb/s over a pair of IP addresses. Use this tool to test the network link for the following issues:

- Maximum throughput
- Congestion
- Packet loss percentage
- Out-of-order delivery
- Other network conditions

The main purpose of this tool is to determine the health of a link before deployment as a circuit.

The following table shows the supported features on the two extension platforms:

Table 44: WAN Tool Supported Features

WAN Tool Supported Features	Brocade 7810 Extension Switch	Brocade 7850 Extension Switch	Brocade SX6 Extension Blade
Max SLA sessions per DP	12	20	20
Max WAN Tool sessions per DP	6	10	10
Max Aggregated BW	2.5Gb/s	50Gb/s (6VE mode) 25Gb/s (18VE mode)	20Gb/s (10VE mode) 10Gb/s (20VE mode)
IPsec	Yes	Yes	Yes
L2cos	Yes	Yes	Yes
Bidirectional traffic	Yes	Yes	Yes
DSCP Marking	Yes	Yes	Yes
L2cos Marking	Yes	Yes	Yes
Connection-type	Yes	Yes	Yes
PMTU	Yes	Yes	Yes
Jumbo Frames (1280 to 9216)	Yes	Yes	Yes

A circuit and WAN Tool cannot operate at the same time since they both cannot utilize the same TCP ports 3225 and 3226. Therefore, you must disable the circuit that you are testing at the local and remote switch before you can configure a WAN Tool connection. When you configure WAN Tool on both switches with the necessary parameters, non-guaranteed TCP connections are established between the switches. Issuing the WAN Tool start command starts traffic flow on these connections.

Multiple non-guaranteed TCP connections are established for the WAN Tool session to ensure that the traffic being generated between the IP pair is as balanced as possible. The configured rate is split equally among various 500Mb/s connections. For example, if you configure a 10Gb/s rate for the test session, 20 500Mb/s connections are created. Another example, if you configure a 1Gb/s rate, two 500Mb/s connections are created. If the rate cannot be split equally into 500Mb/s connections, connections with different rates are created. For example, if you configure a 1.5Gb/s rate, four 375Mb/s connections are created. You can verify that these connections are created after configuring WAN Tool on both switches using the `portcmd--wtool` *`wt-id`* `show -c` command.

For an example of the `Portcmd --wtool` command, see Configuring WAN Tool and Displaying Results.

The following requirements should be considered when using the WAN Tool:

- Brocade SX6 Extension Blade – Each DP supports a maximum of 10 WAN Tool sessions, which reflects the number of enabled WAN Tools and SLA sessions. One blade supports a maximum of 20 sessions.
- Brocade 7850 Extension Switch – Each DP supports a maximum of 10 WAN Tool sessions, which reflects the number of enabled WAN Tools and SLA sessions. The platform maximum is 20 sessions.
- Brocade 7810 Extension Switch – The platform maximum is six sessions.
- A test session can run over an IP path being used by an existing circuit between two switches; however, the circuit must be disabled at each end before configuring the session (applicable to WAN Tool only).
- You must configure the WAN Tool session at each end of the circuit.
- Before a test can start, you must enable each WAN Tool session at each end of the circuit. When the testing is complete, disable and retain the WAN Tool session, which allows you to enable it again as needed.
- After configuration, you can start a test from one switch to solely examine the unidirectional traffic to the opposite switch, or you can test the bidirectional traffic between both switches using the bidirectional option. If you select the latter, you can start the session at either switch.
- You can configure multiple test sessions (one per circuit) for a single port, but the total rate that is configured for all sessions must be equal to or less than the physical speed of the port (100Gb/s, 40Gb/s, 25Gb/s, 10Gb/s, or 1Gb/s).

For example, on a 10Gb/s interface, you can configure four 2.5Gb/s sessions. Alternatively, on a 40Gb/s interface, you can configure four 10Gb/s sessions and cannot cross the maximum supported DP limit.

- The MTU size that is used in the test session is obtained from the IPIF configured value. You can change the MTU size for the IP address pair being tested using the `portcfg ipif [slot/]ge#.dp# modify ip_addr mtu` command. For more information about this command, refer to the *Brocade Fabric OS Command Reference Manual*.
- SLA uses WAN Tool for testing and is referred to as AWT. SLA testing is triggered in the following ways:
- The circuit was created specifying the SLA configuration option, SLA testing runs before bringing a circuit online.
- After a circuit has come online, SLA testing runs when a circuit bounces for any reason.
- SLA testing can be triggered by disabling and enabling a circuit, disabling and enabling a switch, or disabling and enabling a GE interface on which the circuit is configured.

During an AWT session, you might see the following messages:

- When the AWT is enabled: [XTUN-3000], 1161/308, CHASSIS, INFO, SKY76, WAN Tool session 24.6 ENABLED.
- When the AWT session establishes and starts the traffic: [XTUN-3002], 1183/313, CHASSIS, INFO, SKY76, WAN Tool session 24.6 STARTED.
- When the AWT successfully meets SLA requirements: [XTUN-3007], 1250/330, CHASSIS, INFO, SKY76, WAN Tool session 24.6 SLA requirements meet.
- When the AWT fails to meet the SLA loss requirements: [XTUN-3006], 691, CHASSIS, WARNING, SKY76, WAN Tool session 24.6 SLA Failed to meet SLA requirements Reason SLA Drop.
- When the AWT fails to meet the SLA run-time requirements: [XTUN-3006], 691, CHASSIS, WARNING, SKY76, WAN Tool session 24.6 SLA Failed to meet SLA requirements Reason SLA Timeout.
- After completion, the AWT stops traffic and disables the session: [XTUN-3003], 280, CHASSIS, INFO, SKY76, WAN Tool session 24.6 STOP. and [XTUN-3001], 288, CHASSIS, INFO, SKY76, WAN Tool session 24.6 DISABLED.

WAN Tool Commands

Use the `portcmd --wtool` to configure a WAN Tool session. The general syntax for creating a test session including all command options is as follows:

```
portcmd --wtool wt-id create --admin-status [enable|disable] --src src_ip --dst dst_ip - time test_time --
rate link_rate --ipsec policy_name [--bi-directional|-uni-directional] -- dscp dscp --l2cos L2Cos --
connection-type type .
```

You must configure the following parameters on each switch:

- WAN Tool session test ID (*wt-id*): The ID does not have to match on each switch, but this is recommended for an easier comparison of test results on both ends of the circuit when multiple test sessions are created. (However, the rate must be the same.) Valid IDs are 0 through 9. You can configure 10 sessions for each DP and a maximum of 20 sessions per slot.
- Administrative status (`create --admin-status [enable|disable]`): This parameter must be enabled before a test can be run.
- Link rate (*link_rate*) in kb/s: Configure the same link rate on the switch at each end of the circuit. The WAN Tool connections do not fully establish until the same rate is specified for each switch.
- IPsec policy name (*policy_name*): The policy name can be different on each switch; the IPsec policy configuration parameters must be the same on each switch.
- Source IP (*src_ip*) and destination IP (*dst_ip*) address: The source address is the destination address and the destination address is the source address on the opposite switch.
- Bidirectional or unidirectional (`[--bi-directional|-uni-directional]`). This parameter is optional. If you use this parameter, you must configure it on both switches.
- Test session time (*test_time*): The test duration time in minutes must be configured on both switches, but it does not need to match on both switches. The test session uses the time that is configured on the switch where the test started. If bidirectional is specified, the session runs for the time that is configured on the switch where the test started, and then it runs for the time (if configured) configured on the opposite switch.

NOTE
You can create a WAN Tool session without all the required parameters. However, all the required parameters must be configured before a test session can be enabled. For more details on the WAN Tool command and parameters, refer to the *Brocade Fabric OS Command Reference Manual*.

Modify the link rate, test time, test direction (`--bi-directional`) parameters, and clear statistics for a WAN Tool test session after creating a test session, using the `portcmd --wtool wt-id modify` command.

NOTE

You must stop the WAN Tool session before modifying parameters using `portcmd --wtool wt-id stop`

The following examples show the usage of the `modify` parameter:

- To modify the rate, use the `portcmd --wtool wt-id modify --rate link_rate` command.
- To clear test results, use the `portcmd --wtool wt-id modify --clear` command.

Start and stop a configured test session on a specific switch using the following commands:

- `portcmd --wtool wt-id start` . The test duration must be specified with the create or modify parameters. The time can be modified while traffic is not running. The next start command will use the updated time value.
- `portcmd --wtool wt-id stop`

Clear test statistics using the `portcmd --wtool wt-id modify clear` command.

Display historical statistics using the `portcmd --wtool wt-id show --history` command.

NOTE
WAN Tool session historical stats are collected when the `start` command is issued. Until a session has been completed at least once, statistics cannot be displayed. Statistics are also stored when a session is administratively disabled.

For SLA WAN Tool (AWT) sessions, historical statistics are stored automatically when the session runs to completion and it is administratively disabled.

Disable the test sessions using the `portcmd --wtool wt-id modify -a disable` command.

Alternatively, you can delete test sessions using the `portcmd --wtool wt-id delete` command. Delete all configured test sessions using `portcmd --wtool delete -all` .

At this point, after disabling or deleting the configured test sessions, you can re-enable the circuit for operation in a tunnel using the `portCfg fcipcircuit create` command. Display statistics from a WAN Tool session using `portcmd --wtool wt-id show` , where *wt-id* is the ID (0-15) you used to create the test session. Display all test sessions (if multiple test sessions are configured) using `portcmd -- wtool all show` .

For more information about WAN Tool command and parameters, refer to the *Brocade Fabric OS Command Reference Manual*.

Configuring a WAN Tool Session and Displaying Results

Perform the following steps to configure a WAN tool session and display results. When you configure a WAN tool session, you can optionally set the connection type to be a listener or an initiator. Use the `portcmd --wtool` command with no arguments to display all available command options.

1. Connect to the switch and log on as admin.
2. Enter the `Portcmd --wtool` command with the `--help` option to view the options.

```
switch:admin> portcmd --wtool --help

Usage: portcmd --wtool [<slot>/]<session> <action>

  Session ID:
```

```
    [<slot>/]<session>      - WTool session ID Range: 0 - 19.

  Action:
    create          - Create a WTool session.
    modify          - Modify a WTool session.
    delete          - Modify a WTool session.
    delete-all      - Delete all WTool sessions.
    show            - Show WTool session information.
    start           - Start a WTool session.
    stop            - Stop a WTool session.
    stop-all        - Stop all Automated WTool Sessions.

  Create / Modify Args:
   -a,--admin-status enable|disable
                         - Enable or Disable this session.
   -s,--src <ipaddr>     - Set the source IP address.
   -d,--dst <ipaddr>     - Set the destination IP address.
   -t,--time <min>       - Set the session runtime in minutes.
   -r,--rate <kbps>      - Set the session committed rate.
   -i,--ipsec-policy <name>
                         - Set the IPSec Policy for this session.
   -b,--bi-directional   - Set the session for bi-directional traffic.
   -u,--uni-directional  - Set the session for uni-directional traffic.
      --dscp <dscp>      - Set the DSCP marking for this session.
      --l2cos <l2cos>    - Set the L2CoS for this session.
   -C,--connection-type  - Set the connection type. Allowable types are
                           [default] [listener] [initiator]

  Show Args:
   --summary             - Display summary wtool session info.
   --all-ls              - Display sessions from all logical switches.
   -d,--detail           - Display detailed wtool session info.
   -c,--connection       - Show connection information for sessions.
   -H,--historical       - Display historical stats.
   -p,--peer             - Display local and remote statistics.
   -S,--sla              - Display Running SLA configured sessions
   -R,--reset            - Reset WAN Tool session stats.
   -h,--help             - Show this help menu.
```

3. Disable the circuit for the IP pair that you wish to test at each switch using the `portcfg fcipcircuit modify --admin-status disable` command.
   The following example shows how to disable circuit 1:
   ```
   switch1:admin>portCfg fcipcircuit 24 modify 1 --admin-status disable
   ```
4. Verify that the circuit is disabled using the `portshow fciptunnel -c` command. The OpStatus for circuit 1 should be Down.
5. Establish a test connection on the circuit by configuring a WAN Tool session on the switch at one end of the circuit.
   The following example configures a test connection (WAN Tool session 0) on circuit-1 between source IP of 10.1.1.1 and destination IP of 10.1.1.2:
   ```
   Switch1:admin> portcmd --wtool 0 create \
   --src 10.1.1.1 \
   --dst 10.1.1.2 \
   --rate 10000000 \
   ```

```
--time 10 \
--admin-status enable
```

6. Configure the WAN Tool session on the switch at the other end of the circuit.

```
Switch2:admin>portcmd --wtool 0 create \
--src 10.1.1.2 \
--dst 10.1.1.2 \
--rate 10000000 \
--time 10 \
--admin-status enable
```

The `wt-id (0)` option does not need to match the configuration on Switch1, but this is recommended for an easier comparison of test results on both ends of the circuit when multiple test sessions are created. The rate must be the same for both switches. The source address of Switch1 becomes the destination address for Switch2 and the destination address becomes the source address. See WAN Tool Commands for a list of WAN Tool command parameter values that must be identical for both switches in the circuit.

NOTE

If the rate is not configured the same on both sides, the rate is negotiated. The detailed output of the `portcmd --wtool wt-id show --detail` command displays the configured, current, and peer rates.

```
Switch1:admin> portcmd --wtool 2 show --detail

WTool Session: 2 (DP1)
==============================================
Admin / Oper State    : Enabled / Online
Up Time               : 9s
Run Time              : 0s
Time Remaining        : 1m0s
IP Addr (L/R)         : 2002:141::65 ge11 <-> 2002:140::65
IP-Sec Policy         : FID_99Presh
PMTU Discovery (MTU)  : disabled (6666)
Bi-Directional        : enabled
L2CoS / DSCP          : (none) / (none)
Configured Comm Rate  : 15000 kbps    <=== rate mismatch negotiation
Peer Comm Rate        : 20000 kbps    <=== rate mismatch negotiation
Actual Comm Rate      : 15000 kbps    <=== rate mismatch negotiation
Tx rate               : 2171.87 Kbps (  0.27 MB/s)
Rx rate               : 2171.87 Kbps (  0.27 MB/s)
Tx Utilization        : 14.48%
Rx Utilization        : 14.48%
RTT (Min/Max)         : 1 ms/6 ms
RTT VAR (Min/Max)     : 1 ms/3 ms
Local Session Statistics
Tx pkts               : 0
Peer Session Statistics
Rx pkts               : 0
Ooo pkts              : 0
Drop pkts  : 0
Drop% (Overall/5s)    : 0.00% / 0.00%
```

7. Verify that the WAN Tool test connection has established using the `portcmd --wtool wt-id show` command, the `portcmd --wtool wt-id show -c` command, and the `portcmd --wtool wt-id show -d` command.

```
Switch1:admin> portcmd --wtool 0 show

wantool-id: (0)
========================================
```

```
State              : Established
Up Time            : 7m37s
Run Time           : 0s
Time remaining     : 0s
IP Addr (L/R)      : 10.1.1.2 <-> 10.1.1.1
PMTUD              : Disabled
Comm Rate          : 10000000 Kbps (1220.70 MB/s)
Tx rate            : 4562.50 Kbps (0.56 MB/s)
Rx rate            : 4539.69 Kbps (0.55 MB/s)
Tx Utilization     : 0.05%
Rx Utilization     : 0.05%
RTT (Min/Max)      : 0.10ms/0.28ms
RTT VAR (Min/Max)  : 0.09ms/0.34ms
Local Session Statistics
Tx pkts            : 0
Peer Session Statistics
Rx pkts            : 0
Ooo pkts           : 0
Drop pkts          : 0 (0.00%)

Switch1:admin> portcmd --wtool 0 show --connection

Id  Port(L/R)     Rate(Tx/Rx)  UpTime  RunTime
=========================================================================
6   63494 / 3225  0.03  /  0.03  8m8s  0s
17  63490 / 3225  0.03  /  0.03  8m8s  0s
14  63498 / 3225  0.03  /  0.03  8m8s  0s
3   61443 / 3226  0.03  /  0.03  8m8s  0s
11  61447 / 3226  0.03  /  0.03  8m8s  0s
9   61446 / 3226  0.03  /  0.03  8m8s  0s
1   61442 / 3226  0.03  /  0.03  8m8s  0s
20  63491 / 3225  0.03  /  0.03  8m8s  0s
8   63495 / 3225  0.03  /  0.03  8m8s  0s
12  63497 / 3225  0.03  /  0.03  8m8s  0s
4   63493 / 3225  0.03  /  0.03  8m8s  0s
16  63489 / 3225  0.03  /  0.03  8m8s  0s
13  61448 / 3226  0.03  /  0.03  8m8s  0s
19  61440 / 3226  0.03  /  0.03  8m8s  0s
5   61444 / 3226  0.03  /  0.03  8m8s  0s
15  61449 / 3226  0.03  /  0.03  8m8s  0s
7   61445 / 3226  0.03  /  0.03  8m8s  0s
18  61441 / 3226  0.03  /  0.03  8m8s  0s
10  63496 / 3225  0.03  /  0.03  8m8s  0s
2   63492 / 3225  0.03  /  0.03  8m8s  0s
=========================================================================
Number of Connections:20

Switch1:admin> portcmd --wtool 0 show -d

WTool Session: 0 (DP0)
================================================
Admin / Oper State  : Enabled / Online
Up Time             : 55m54s
```

```
Run Time            : 0s
Time Remaining      : 1m0s
IP Addr (L/R)       : 10.1.9.76 ge9 <-> 10.1.9.77
IP-Sec Policy       : (none)
PMTU Discovery (MTU): disabled (1500) Bi-Directional  : disabled
L2CoS / DSCP        : (none) / (none) Configured Comm Rate  : 20000 kbps
Peer Comm Rate      : 20000 kbps
Actual Comm Rate    : 20000 kbps
Tx rate             : 517.04 Kbps (  0.06 MB/s)
Rx rate             : 517.04 Kbps (  0.06 MB/s)
Tx Utilization      : 2.59%
Rx Utilization      : 2.59%
RTT (Min/Max)       : 1 ms/1 ms
RTT VAR (Min/Max)   : 1 ms/1 ms
Local Session Statistics
Tx pkts             : 0
Peer Session Statistics
Rx pkts             : 0
Ooo pkts            : 0
Drop pkts           : 0
Drop% (Overall/5s)  : 0.00% / 0.00%

Switch1:admin> portcmd --wtool show

Session OperSt Flags   LocalIP   RemoteIp  TxMBps RxMBps Drop%
--------------------------------------------------------------------------------
0       Up     ----4-- 10.1.1.1  10.1.1.2  0.64   0.64   0.00
--------------------------------------------------------------------------------
Flags (wtool): S=SLA v=VLAN i=IPsec 4=IPv4 6=IPv6 L=Listener I=Initiator
```

The example output from the `--wtool 0 show` command indicates that the connection has an established state. The example output from the `--wtool 0 show -c` command displays the non-guaranteed TCP connections that are created between TCP ports 3225 and 3226 to balance the traffic. For the 10Gb/s test connection, 20 non-guaranteed TCP connections are created. The output from the `--wtool 0 show -d` command shows additional details.

8. Display peer information using the `--wtool show --peer` command.

```
switch:admin> portcmd --wtool 0 show --peer

WTool Session (0)          (Local)                    (Remote)
--------------------------------------------------------------------------------
Admin / Oper State         : Up                       : Up
Up Time                    : 6m49s                    : 6m49s
Run Time                   : 0s                       : 0s
Time Remaining             : 10m0s                    : -
Port                       : ge9                      : -
IP Addr                    : 10.1.1.1                 : 10.1.1.2
IP-Sec Policy              : (none)                   : (none)
Configured Comm Rate       : 10000000 kbps            : 10000000 kbps
Actual Comm Rate           : 10000000 kbps            : 10000000 kbps
PMTU Discovery (MTU)       : disabled (1500)          : disabled (1500)
L2cos /DSCP                : (none) / (none)          : (none) / (none)
Tx Rate                    : 5156.03 Kbps             : 0.00 Kbps
Rx Rate                    : 5144.34 Kbps             : 0.00 Kbps
Tx Utilization             : 0.05%                    : 0.00%
```

```
Rx Utilization             : 0.05%                    : 0.00%
RTT (Min/Max)              : 1 ms/1 ms                : 1 ms/1 ms
RTT VAR (Min/Max)          : 1 ms/1 ms                : 1 ms/1 ms
Tx pkts                    : 0                        : 0
Rx pkts                    : 0                        : 0
Ooo pkts                   : 0                        : 0
Drop pkts                  : 0                        : 0
Drop% (Overall / 5s)       : 0.00% / 0.00%            : 0.00% / -
```

9. SLA for automated WAN tool sessions, view the configured and negotiated SLA configuration for the session. This makes it easier to see the SLA requirements and whether the requirements were negotiated.

```
switch:admin> portcmd --wtool 24.0 show -d

WTool Session: 24.0 (DP0)
=========================================
Admin / Oper State      : Enabled / Running
Up Time                 : 25s
Run Time                : 25s
Time Out                : 1h29m35s
Time Remaining          : 4m59s
IP Addr (L/R)           : 170.195.7.10 ge7 <-> 171.196.7.10
IP-Sec Policy           : (none)
PMTU Discovery (MTU)    : disabled (1500)
Bi-Directional          : disabled
L2CoS / DSCP            : (none) / (none)
SLA (Run Time / Timeout / Loss)
Configured              : (5m / 1h30m / .5%)    <=====
Actual                  : (5m / 1h30m / .2%)    <=====
Configured Comm Rate    : 1000000 kbps
Peer Comm Rate          : 1000000 kbps
Actual Comm Rate        : 1000000 kbps
Tx rate                 : 999565.18 Kbps ( 124.95 MB/s)
Rx rate                 : 750737.63 Kbps ( 93.84 MB/s)
Tx Utilization          : 99.96%
Rx Utilization          : 75.07%
RTT (Min/Max)           : 50 ms/51 ms
RTT VAR (Min/Max)       : 1 ms/9 ms Local Session Statistics
Tx pkts                 : 153112464
Peer Session Statistics
Rx pkts                 : 114832047
Ooo pkts                : 0
Drop pkts               : 38278953
Drop %    (overall/5s)  : 25.00%/.5%
```

10. Multiple WAN Tool sessions, verify basic connection information using the `--wtool all show` command.

```
Switch1:admin> portcmd --wtool all show

Session OperSt Flags    LocalIP      RemoteIp     TxMBps  RxMBps  Drop%
--------------------------------------------------------------------------------
1       Up     -i4pv    10.1.9.77    10.1.9.76    0.06    0.06    0.00
25.0    Up     S-4--    10.1.8.77    10.1.8.76    0.06    0.06    0.00
--------------------------------------------------------------------------------
Flags (wtool): S=SLA v=VLAN i=IPsec 4=IPv4 6=IPv6 L=Listener I=Initiator
```

Output for this example shows that WAN Tool session 1 was created to test the circuit with IP address pair 10.1.9.77 and 10.1.9.76 and session 25.0 was created testing the circuit with IP address pair 10.1.8.77 and 10.1.8.76.

11. Start traffic on the test connection by entering the `portcmd --wtool wt-id start` command.

```
Switch1:admin> portcmd --wtool 0 start
```

12. Verify that the test session started by entering the `portcmd --wtool show wt-id --detail` command.

```
Switch1:admin> portcmd --wtool show --detail

WTool Session: 24.0 (DP0)
=========================================
Admin / Oper State      : Enabled / Running
Up Time                 : 10s
Run Time                : 9s
Time Out                : 3m50s
Time Remaining          : 1m51s
IP Addr (L/R)           : 10.1.1.2 ge9 <-> 10.1.1.1
IP-Sec Policy           : (none)
PMTU Discovery (MTU)    : disabled (1500)
Bi-Directional          : disabled
L2CoS / DSCP            : (none) / (none)
Configured Comm Rate    : 1000000 kbps
Peer Comm Rate          : 1000000 kbps
Actual Comm Rate        : 1000000 kbps
Tx rate                 : 999624.45 Kbps ( 124.95 MB/s)
Rx rate                 : 1000000.00 Kbps ( 125.00 MB/s)
Tx Utilization          : 99.96%
Rx Utilization          : 100.00%
RTT (Min/Max)           : 1 ms/1 ms
RTT VAR (Min/Max)       : 1 ms/1 ms
Local Session Statistics
Tx pkts                 : 810024
Peer Session Statistics
Rx pkts                 : 792029
Ooo pkts                : 0
Drop pkts               : 0 (0.00%)
```

The State shows that the test is running and other statistics display as well, such as test Run Time and Time Remaining.

13. Start the test from the other switch by entering the `portcmd --wtool wt-id` command.

NOTE
If you used the `bi-directional` option when creating the session, you can start the session on either switch.

```
Switch2:admin> portcmd --wtool 0 start
```

14. Verify that the test session started on the other switch by entering the `portcmd --wtool wt-id show` command.

```
Switch2:admin> portcmd --wtool 0 show
```

15. Delete the WAN Tool session on both switches, or disable the WAN Tool session on both switches.
    - To delete the WAN Tool session, use the `portcmd --wtool wt-id delete` command. This deletes the session and allows you to enable the circuit for a normal traffic.
    - To disable the WAN Tool session, use the `portcmd --wtool wt-id modify --admin -status disable` command. This disables the session without deleting it, so you can use the same session again.
    - To verify that the WAN Tool session is disabled, enter the `portcmd --wtool wt-id show` command or the `portcmd --wtool wt-id show -d` command.

16. Enable the circuit from each switch using the `portcfg fcipcircuit port modify` command. The following example enables the circuit from Switch1.

```
Switch1:admin> portCfg fcipcircuit 24 modify 1 --admin-status enable
```

Resolving Test Session Problems

If the output from the `portcmd --wtool wt-id show` command shows that the state is down, constantly in progress, or the connection times out (changes from an up to down state), the WAN Tool test connection is not being established. Verify that you have properly configured the session on both ends with the appropriate parameters. For the list of parameters required for each switch, see WAN Tool Commands.

The following WAN Tool configuration issues can cause problems establishing a connection:

- The test rate does not match on each switch.
- The test rate on a single circuit or multiple circuits on a port is greater than the rate allowed for the port. This generates a warning that the bandwidth has been exceeded and blocks you from creating a session.
- The IPsec policy does not match on each switch. The IPsec names need not to match, but the names must refer to the same policy.
- The configured source and destination IP addresses are not correct on one or both switches.

This section contains information for troubleshooting extension links.

- When deleting extension links, you must delete them in the exact reverse order in which they were created. That is, first delete the tunnels, followed by any IP route configurations, then the IP interfaces, and finally the port configuration.
- The `portcmd --ping` command verifies network connectivity; it does not verify that everything is correctly configured for a tunnel.
- For the tunnel to function correctly, the following configurations are made at both ends of the tunnel:
  - GE interfaces
  - IPIF
  - IP routes
  - ARL
  - IPsec (optional)
  - Compression (optional)
  - FastWrite (optional)
  - OSTP (optional)

Because of an IPsec RASLog limitation, you might not be able to determine if an incorrect configuration causes an IPsec tunnel to remain inactive. This misconfiguration can occur on either end of the tunnel. As a result, you must correctly match the encryption method, authentication algorithm, and other configurations on each end of the tunnel.

Gathering Information

The following commands must be issued and their data that are collected before the `supportSave` command is run. Using the `supportSave` command can take 10 minutes or more to run, and some of the information is time-critical.

1. Use the `tracedump -n` command to generate a local trace dump locally.
2. Run and collect the data for issues specific to tunnel ports using the following commands:
   - `slotshow`
   - `portshow slot/GE_port`
3. Use the following commands to gather information on the current GE port configurations and extension mode configurations:

- `extncfg --show`
- `portcfgge --show`

4. Use the following commands to gather information on various parameters of port channel and configuration parameters for ports and GbE ports:
   - `portchannel --show`
   - `portchannel --show -detail`
   - `portshow tcl`
   - `portshow tcl --detail`
   - `portshow lan-stats --global`
   - `portshow lan-stats --per-flow`
   - `portshow lan-stats --hist-stats`
5. Run and collect the data using the following commands if you use multiple logical switches:
   - `portshow ipif --all-ls`
   - `portshow arp --all-ls`
   - `portcfgshow iproute --all-ls`
   - `portshow fciptunnel --circuits --all-ls`
6. Finally, gather the data using the `supportSave` command.

For more information about the above commands, refer to the *Brocade Fabric OS Administration Guide* and *Brocade Fabric OS Command Reference Manual*.

A Tunnel Does Not Come Online

Perform the following steps to troubleshoot a tunnel that fails to come online:

1. Confirm that the Ethernet port is online.

```
switch:admin> portshow ge0
Eth Mac Address: 00.05.1e.37.93.06
Port State: 1    Online
Port Phys: 6    In_Sync
Port Flags: 0x3 PRESENT ACTIVE
Port Speed: 10G
```

2. Confirm that the IP configuration is correct on both tunnel endpoints using the following command:

```
switch:admin> portshow ipif ge0.dp0

 Port       IP Address                   / Pfx MTU   VLAN  Flags
-------------------------------------------------------------------------------
ge0.dp0     192.168.5.2                  / 24  1500  0     U R M I
ge0.dp1     192.168.5.12                 / 24  1500  0     U R M I
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running I=InUse
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport

switch:admin> portshow fciptunnel 24 --circuit --config

Tunnel: VE-Port:24 (idx:0, DP0)
================================================
Oper State              : Enabled
TID                     : 24
Flags                   : 0x00000000
IP-Extension            : Disabled
```

```
Compression             : None
QoS BW Ratio            : 50% / 30% / 20%
Fastwrite               : Disabled
Tape Pipelining         : Disabled
IPSec                   : Disabled
Load-Level (Cfg/Peer)   : Failover (Failover / Failover)
Local WWN               : 10:00:50:eb:1a:14:a9:46
Peer WWN                : 10:00:50:eb:1a:13:ad:16
RemWWN (config)         : 00:00:00:00:00:00:00:00
cfgmask                 : 0x001007ff 0x40000208
Flow Status             : 0
ConCount/Duration       : 1 / 5d2h
Uptime                  : 5d2h
Stats Duration          : 5d2h
Receiver Stats          : 7704928 bytes / 51813 pkts /  8.00 Bps Avg
Sender Stats            : 6303500 bytes / 51810 pkts /  8.00 Bps Avg
TCP Bytes In/Out        : 4465088992 / 4463631336
ReTx/OOO/SloSt/DupAck   : 0 / 0 / 0 / 0
RTT (min/avg/max)       : 1 / 1 / 6 ms Wan Util  : 0.0%
TxQ Util                : 0.0%

Circuit 24.0 (DP0)
==================================================
Admin/Oper State        : Enabled / Online Warning  <======
Flags                   : 0x00000000
IP Addr (L/R)           : 192.168.5.2 ge0 <-> 192.168.1.2
HA IP Addr (L/R)        : 192.168.5.12 ge0 <-> 192.168.1.12
Configured Comm Rates   : 5000000 / 5000000 kbps
Peer Comm Rates         : 5000000 / 5000000 kbps
Actual Comm Rates       : 4500000 / 5000000 kbps
Keepalive (Cfg/Peer)    : 6000 (6000 / 6000) ms
Metric                  : 0
Connection Type         : Default
ARL-Type                : Auto
PMTU                    : Disabled
SLA                     : (none)
Failover Group          : 0
VLAN-ID                 : NONE
L2Cos (FC:h/m/l)        : 0 / 0 / 0 (Ctrl:0)
L2Cos (IP:h/m/l)        : 0 / 0 / 0
DSCP (FC:h/m/l)         : 0 / 0 / 0 (Ctrl:32)
DSCP (IP:h/m/l)         : 0 / 0 / 0
cfgmask                 : 0x40000000 0x00213def
Configuration Warnings:                          <======
  Min-comm-rate / QoS-Ratio / Dist-Ratio
Flow Status             : 0
ConCount/Duration       : 1 / 5d2h
Uptime                  : 5d2h
Stats Duration          : 5d2h
Receiver Stats          : 7704928 bytes / 51813 pkts /  8.00 Bps Avg
Sender Stats            : 6303500 bytes / 51810 pkts /  8.00 Bps Avg
TCP Bytes In/Out        : 4465092320 / 4463634664
ReTx/OOO/SloSt/DupAck   : 0 / 0 / 0 / 0
```

```
RTT (min/avg/max)       : 1 / 1 / 6 ms
Wan Util                : 0.0%
```

In this example, the Online Warning indicates a problem and the information that is shown in Configuration Warnings implies where the problem might be. Both ends of the circuit should be configured with the same parameters.

3. Use the `portcmd --ping` command with the local and remote circuit IP addresses.

   If the ping is successful, there is an end to end connectivity. If the ping is not successful, try to ping from the local IP address to the next hop router gateway. If that is successful, try to ping to the remote router gateway.

   The `-s` value is the ping source IP address, and the `-d` value is the destination IP address.

   ```
   portcmd --ping ge0.dp0 -s 10.10.10.1 -d 10.10.10.2
   ```

   If the command is successful, then you have IP connectivity and your tunnel should come up. If not, continue to the next step.

4. Use the `portcmd --traceroute` command with the local and remote circuit IP addresses.

   ```
   portcmd --traceroute ge0.dp0 -s 10.10.10.1 -d 10.10.10.2
   ```

5. If there are routed IP connections that provide for the tunnel, confirm that both ends of the tunnel have defined IP routes and that the route gateways are correct. The tunnel or route lookup might fail to come online because of a missing or incorrect IP route.
6. Confirm that the tunnel is correctly configured using the following command:

   ```
   portshow fciptunnel <slot/VE>
   ```

   Confirm that the IPsec, compression, FastWrite, and OSTP settings match at each end or the tunnel might not come up. Confirm the local and remote IP addresses are accurate.

7. Generate an Ethernet sniffer trace.

Rule out all possible blocking factors. Routers and firewalls that are in the data path must be configured to pass traffic (TCP ports 3225 and 3226) and for the IPsec traffic (UDP port 500). If possible blocking factors have been ruled out, simulate a connection attempt using the `portCmd --ping` command, from source to destination, and then generate an Ethernet trace between the two endpoints. The Ethernet trace can be examined to further troubleshoot the connectivity.

A Tunnel Goes Online and Offline

A bouncing tunnel that goes on and offline is a common problem; this usually occurs because of oversubscription of the available bandwidth, which can result in the following behaviors:

- More data is sent to the link than it can accommodate
- Management data gets lost or is queued for too long
- Timeouts are encountered

Perform the following steps to gather information:

1. Verify the available link bandwidth.
2. Confirm that the IP path is exclusive to the extension traffic.
3. Confirm that traffic shaping is configured to limit the available bandwidth by using the following command:

   ```
   portShow fciptunnel all --tcp
   ```

   Examine data from both routers. This data shows retransmissions, indicating input and output rates on the tunnels.

4. Use the WAN Tool to gather WAN performance data.

Extension Performance Monitor (EPM)

Extension Performance Monitor (EPM) is designed to collect data for issues with the extension tunnel performance.

The EPM provides a ready-made command available on the switch to run while they are experiencing performance issues. The main function of EPM is to run both user-level and engineering-level debug commands repeatedly over a user-specified duration and then log the information for a retrieval through supportSave. With this feature, the support and engineering team can better understand and troubleshoot customer performance issues.

You can monitor the following four main areas for performance with an extension switch or blade:

- System (SYS)
- LAN
- WAN
- Fiber Channel (FC)

Monitoring the system involves overall DP performance including memory, SSO, and pool information. LAN monitoring deals with the IP Extension connection health and the basic extension tunnel performance. WAN monitoring captures information about the overall health and performance of the extension tunnels. FC monitoring captures performance information that is related to the emulation of FC frames over an extension tunnel.

NOTE
LAN monitoring is supported only in IP Extension mode.

Only one instance of EPM can be running at any one time per chassis. If EPM is already running, a new instance of the command is not allowed and an error message is displayed. You can manually stop the EPM session before the scheduled termination time.

In multi-slot or multi-DP environments, you can collect EPM data for any one slot/dp pair. You must specify the slot/dp pair for the data collection. The log files are saved on the switch in `/var/log/` . The default file name uses the `epm_<date><slot_dp>_extnperfmon.log` format. You can generate up to five log files on an active CP at any one time. If you try to generate a sixth log file, you are prompted to run a supportSave or clean up the old files.

supportSave captures the log files from all CPs and all partitions, and it removes them from the switch after successfully adding them to the `esmd.tar` file. By default, EPM monitors the system for 15 minutes. The monitoring duration ranges from 5 minutes to 60 minutes.

NOTE

- EPM does not support monitoring a system during the supportSave process. If a supportSave is issued during the monitoring period, the monitor gracefully shuts down with the reason code supportSave detected.
- EPM does not support monitoring a system during a firmware upgrade or downgrade.

Using the extnperfmon Command

Use the `extnperfmon` command to view information on the Brocade Extension Switches and Blades. The following CLI output shows an example of the `extnperfmon` command with the `--help` option:

```
switch:admin> extnperfmon --help

Usage:
extnPerfMon <cmd> [<args>]

Commands:
```

```
--status    Display the current status of the performance monitor and show any current log files that are
 found.
--monitor <args>    Set the options for the performance monitor.
--start [<args>]    Start the performance monitor with the current settings and optionally include arguments
 to set for the session.
--stop    Stop the currently running monitoring session.
--cleanup [-force]    Cleanup log files from any performance monitor sessions. (optionally use -force option
 to bypass confirmation)
--help    Display command syntax help menu for performance monitor.

Arguments for 'monitor' and 'start' commands:
-dp [<slot>/]dp<#>    Specify the DP to monitor (required).
[-time <mins>]    Specify the time in minutes to monitor (default:15) [-filename <string>] Specify a string
 name to call this monitoring session
(31 characters max, default:"extnperfmon").
[-sys]    Enable SYSTEM monitoring flags (default:enabled).
[-lan]    Enable LAN/IPEXT monitoring flags (default:disabled).
[-wan]    Enable WAN monitoring flags (default:disabled).
[-fc]     Enable FC monitoring flags (default:disabled).
[-all]    Enable ALL monitoring flags (default:disabled).
```

Extension Performance Monitoring --status Command

Use the `extnperfmon --status` command to view the current status of an EPM and any current log files from prior sessions. The following example shows the `--status` output with all default values and when EPM has never been run:

```
switch:admin> extnperfmon --status
Nothing currently configured.
Existing performance monitor files: 0
```

The following example shows the `--status` output when there is an active session monitoring slot 3, DP 0 for the `-sys` , `-lan` , and `-wan` options:

```
switch:admin> extnperfmon --status

Session File Name:            extnperfmon
Session Status:               Monitoring
Modules Monitoring:           SYS,LAN,WAN
DP Monitored:                 3/DP0
Duration/Time remaining:      60/30 minutes
Start Time:                   8/02/2023 18:59:22
End Time:                     8/02/2023 19:59:22 (expected)
```

Extension Performance Monitoring --monitor Command

Use the `extnperfmon --monitor` command to set up the configuration of the EPM session for its various options. You can set the `-dp` , `-time` , `-filename` , `-lan` , `-wan` , `-fc` , or `-all` options. The following monitor command is used to set the `-filename` to SLOT3_PERFORMANCE_LAN_WAN, set the `-dp` option to slot 3/DP 1, set the `-time` to 20 minutes, and add `-lan` and `-wan` modules to be monitored.

```
switch:admin> extnperfmon --monitor -filename SLOT3_PERFORMANCE_LAN_WAN -dp 3/dp1 -time 20 -lan -wan
Operation Succeeded.
switch:admin> extnperfmon --status
```

```
Session File Name:        SLOT3_PERFORMANCE_LAN_WAN
Session Status:           Not Run
Modules Monitoring:       SYS,LAN,WAN
DP Monitored:             3/DP1
Duration/Time remaining: 20 (minutes) / --
Start Time:               --
End Time:                 --
```

Extension Performance Monitoring --status Command

Use the `extnperfmon --status` command to view the current status of an EPM and any current log files from prior sessions. The following example shows the `--status` output with all default values and when EPM has never been run:

```
switch:admin> extnperfmon --status
Nothing currently configured.
Existing performance monitor files: 0
```

The following example shows the `--status` output when there is an active session monitoring slot 3, DP 0 for the `-sys` , `-lan` , and `-wan` options:

```
switch:admin> extnperfmon --status

Session File Name:            extnperfmon
Session Status:               Monitoring
Modules Monitoring:           SYS,LAN,WAN
DP Monitored:                 3/DP0
Duration/Time remaining:      60/30 minutes
Start Time:                   8/02/2023 18:59:22
End Time:                     8/02/2023 19:59:22 (expected)
```

Extension Performance Monitoring --stop Command

Use the `extnperfmon --stop` command to stop the active monitoring session. The following is an example of the `--stop` command while no session is running:

```
switch:admin> extnperfmon --stop

Stopping the Extension Performance Monitor session. This process may take up to 2 minutes.
extnperfmon: Object does not exist.
No extnperfmon instance found. Unable to stop.
```

The following is an example of the `--stop` command while a session is running.

```
switch:admin> extnperfmon --stop Stopping the Extension Performance Monitor session. This process may take up
 to 2 minutes.
extnperfmon: Operation Succeeded.
```

Extension Performance Monitoring --cleanup Command

Use the `extnperfmon --cleanup` command to clean up the log files stored in /var/log and /mnt/var/log in both the active and backup CP. The following is an example of the `--cleanup` option:

```
switch:admin> extnperfmon --cleanup
Existing performance monitor files:
```

```
    /mnt/var/log/epm_210320_004327_DP0_extnperfmon.log
    /mnt/var/log/epm_210320_004614_DP0_extnperfmon.log

Continue with file deletion (y/n)? y

Operation Completed.
```

Optionally, use the `-force` command to clean up log files without being prompted for confirmation.

```
switch:admin> extnperfmon --status
-----------------------------------------------------

Session File Name:         extnperfmon
Session Status:            Stopped - User Request
Modules Monitoring:        SYS (0x7f)
DP Monitored:              DP0
Duration/Time remaining:   5 (minutes) / --
Start Time:                04/20/2023-18:27:31
End Time:                  04/20/2023-18:29:11
-----------------------------------------------------


Existing performance monitor files:
  /var/log/epm_210420_143217_DP0_extnperfmon.log.gz
  /var/log/epm_210420_143515_DP0_extnperfmon.log.gz
  /var/log/epm_210420_182731_DP0_extnperfmon.log.gz
switch:admin> extnperfmon --cleanup -force

Existing performance monitor files:
  /var/log/epm_210420_143217_DP0_extnperfmon.log.gz
  /var/log/epm_210420_143515_DP0_extnperfmon.log.gz
  /var/log/epm_210420_182731_DP0_extnperfmon.log.gz

extnperfmon: Operation Succeeded.
switch:admin> extnperfmon --status
-----------------------------------------------------

Session File Name:          extnperfmon
Session Status:             Stopped - User Request
Modules Monitoring:         SYS (0x7f)
DP Monitored:               DP0
Duration/Time remaining:    5 (minutes) / --
Start Time:                 04/20/2023-18:27:32
End Time:                   04/20/2023-18:29:12
-----------------------------------------------------

Existing performance monitor files: (none)
```

FTRACE

FTRACE is a support tool that is used primarily by your support provider to communicate an issue with Broadcom Support.

FTRACE can be used in a manner similar to that of a channel protocol analyzer. You can use FTRACE to troubleshoot problems through an SSH session rather than using an analyzer or sending technical support personnel to the installation site. FTRACE is always enabled on Brocade Extension platforms, and trace data is automatically captured.

FTRACE Configuration

A default configuration for FTRACE is provided for each DP on each platform, which allows tracing events specific to each DP.

The `portcfg ftrace slot/ve_port cfg` command is interactive. Use this command only under the direction of an authorized support representative.

Each DP has an FTRACE instance. The default FTRACE configuration defines eight buffers on each DP for trace events, and 300,000 trace entries per buffer. The default FTRACE configuration enables auto checkout (ACO) for the first four buffers and disables ACO for the last four. The Brocade Extension platform has SSD disk on each DP, which is used to save copies of triggered FTRACE buffers. Saving FTRACE buffers to SSD is enabled by default, and by the Save to Flash `portcfg ftrace ve_port cfg` command.

ACO can be enabled for each defined FTRACE buffer. FTRACE processing varies when the FTRACE buffer is defined as ACO enabled or ACO disabled.

ACO enabled: If the FTRACE buffer is defined with ACO enabled, when that buffer is the current FTRACE buffer and a trigger event occurs, FTRACE post fills that buffer to the end, or add the post fill percentage of more trace entries. During the post filling process, the FTRACE buffer state is reported as post fill. When the post filling process has completed, the buffer state is reported as checked out, and the next sequential available buffer is assigned to the current buffer (state current). If all FTRACE buffers are marked as checked out, FTRACE no longer records trace entries. The default configuration, therefore, captures at least the first four error traces, and permanently checks out those buffers, and moves them to the ACO-off buffers. FTRACE buffers that have been checked-out are included in a supportSave. When the supportSave is complete, the buffers are returned to the unused state and available for new traces. You can use the `portshow ftrace ve_port cmd` command to check in or check out buffers.

ACO disabled: If the FTRACE buffer is defined with ACO disabled, when that buffer is the current FTRACE buffer and a trigger event occurs, FTRACE processing completes the same post filling process as described for ACO enabled. When completed, if the `Save to Flash` configuration option is enabled, the buffer moves to a saving state, and the next available buffer will be made the current trace buffer. FTRACE saves as many as eight buffers in the DP SSD file system. If there are already eight saved FTRACE buffers in the file system, the oldest trace buffer is replaced by the current buffer being saved. When the `Save to Flash` processing completes, the buffer is marked as triggered. If the `Save to Flash` option is not enabled, the buffer is immediately marked as triggered and the next sequentially available FTRACE buffer is marked as the current buffer.

In the default configuration, FTRACE captures at least the first four error events in buffers 0, 1, 2, and 3. It captures the last three error events in triggered buffers 4-7, and always has the current buffer. Also, buffers 4-7 potentially have as many as 10 saved prior trigger events that are reported and saved in the DP SSD file system.

FTRACE data contents are included in a switch SupportSave capture. After the SupportSave has been captured, the FTRACE buffers are reset and all buffers that were previously either Checked Out or Triggered return to an unused state.

Change the FTRACE ACO configuration using the `portcfg ftrace`*`slot/ve_port`*`cfg` command. For more information, see Changing Configuration Settings.

Changing Configuration Settings

Use the `portcfg ftrace` *`slot/ve_port`* command to change the FTRACE configuration settings. The configuration for FTRACE is defined using the first VE_Port on the DP as follows:

- Brocade 7810 Extension Switch: VE12 on DP0
- Brocade 7850 Extension Switch: VE24 on DP0 and VE33
- Brocade SX6 Extension Blade: VE<slot>/16 on DP0 and VE<slot>/26 on DP1

To change FTRACE configuration settings on the first DP complex (DP0) on a Brocade SX6 Extension Blade, if applicable, set the context where the VE_Port 7/16 is defined, and then enter the following command:

```
portcfg ftrace 7/16 cfg
```

To change FTRACE configuration settings on the second DP complex (DP1) on a Brocade SX6 Extension Blade, if applicable, set the context to where VE_port 7/26 is defined, and then enter the following command:

```
portcfg ftrace 7/26 cfg
```

The `portcfg` command is an interactive command sequence and prompts for configuration items.

Common Acronyms

This section provides information on common acronyms and their expansions.

Table 45: Brocade Extension Platforms – Common Acronyms

Acronym	Expansion
AAA	Authentication Authorization Accounting (Security)
ACC	BSL Automated Case Creation
ACL	Access Control List (IP)
AES	Advanced Encryption Standard (Encryption)
AH	Authentication Header (Encryption)
ARL	Adaptive Rate Limiting (Fabric OS Extension feature)
ASIC	Application Specific Integrated Circuit
BET	Brocade Extension Trunking (Fabric OS Extension feature)
BPA	BSL Best Practice Assessment
BSL	Brocade Support Link
BT	Brocade Trunking (Fabric OS FC feature)
BUM	Broadcasts Unknowns Multicasts (IP)
CA	Certificate Authority (Certificates)
CBC	Cipher Block Chaining (Encryption)
CHPID	Channel Path ID (Mainframe)
CIDR	Classless Inter-Domain Routing (Subnet prefixes)
CIR	Committed Information Rate
CNSA	Commercial National Security Algorithm Suite
CPI	BSL Configuration, Performance, Inventory
CRL	Certificate Revocation List (CA and Certificates)
CSR	Certificate Signing Request (CA and Certificates)
CUP	Control Unit Port (Mainframe)
DASD	Direct-Attached Storage Device (Mainframe)
DBR	Device Based Routing (FC)
DCA	BSL Data Collection Assistant
DID	Destination ID (FC destination PID)
DID	Domain ID (FC)
DISL	Dedicated ISL (FC E_Port)
DLS	Dynamic Load Sharing (FC)
DNS	Domain Name System
DP	Data Processor (Extension HW)
DRAM	Dynamic Random Access Memory
DSCP	Differentiated Services Code Point (QoS)

Acronym	Expansion
DST	Destination
DWDM	Dense Wave Division Multiplexing (Networking)
E_Port	Expansion port (ISL)
EBR	Exchange Based Routing (FC)
ECDH	Elliptic Curve Diffie-Hellman (Encryption)
ECDSA	Elliptic Curve Digital Signature Algorithm (Encryption)
ECM	Extension Configuration Manager (Fabric OS Extension feature)
eHCL	Extension Hot Code Load (Fabric OS Extension feature)
EPM	Extension Performance Monitor (Fabric OS Extension feature)
ESP	Encapsulating Security Payload (Encryption)
EX_Port	SAN expansion port for FCR
F_Port	SAN fabric port
FA	BSL Fabric Analytics
FC	Fibre Channel
FCIP	Fibre Channel over Internet Protocol
FCP	Fibre Channel Protocol
FCR	Fibre Channel Routing
FDCB	FICON Device Control Block
FDPB	FICON Device Path Block
FEC	Forward Error Correction
FICON	Fiber Connection (Mainframe)
FID	Fabric ID (Fabric OS VF LS feature)
FMS	FICON Management Server (Fabric OS FICON feature)
FOS	Brocade Fabric Operating System
FPGA	Field Programmable Field Array
FPI	Fabric Performance Impact (Fabric OS FC feature)
FPIN	Fabric Performance Impact Notification (Fabric OS FC feature)
FSPF	Fabric Shortest Path First
FV	Fabric Vision
FW	Firmware
Gbps Gb/s	Gigabits per second
GCM	Galois Counter Mode (Encryption)
GE	Gigabit Ethernet (interfaces)
GW	Gateway (Networking)
HA	High Availability
HCL	Hot Code Load
HEE	High Efficiency Encapsulation (Fabric OS Extension feature)
HMAC	Hash-based Message Authentication Code
HoLB	Head of Line Blocking
HSRP	Hot Standby Router Protocol (Networking)

Acronym	Expansion
ICL	Inter-Chassis Link (Core Routing Blade Links)
IFCC	Interface Control Check (Mainframe)
IFL	Inter-Fabric Link (EX_Port)
IKE	Internet Key Exchange (Encryption, Certificates, IKEv2)
IOD	In-Order Delivery (FC and FICON)
IP	Internet Protocol (Networking)
IPEX	IP Extension (Fabric OS Extension feature)
IPIF	IP Interface (Extension)
IPsec	IP Security (Encryption)
IPv4	IP version 4 (Networking)
IPv6	IP version 6 (Networking)
ISL	Inter-Switch Link (E_Port)
ITL	Initiator-Target-LUN (Storage)
ITN	Initiator-Target Nexus (Storage)
KATOV	Keepalive Time Out Value (Fabric OS Extension feature)
L2	Layer 2 (Ethernet)
L2CoS	Layer 2 Class of Service (QoS)
L3	Layer 3 (IP)
LACP	Link Aggregation Control Protocol (Ethernet)
LAG	Link Aggregation Group (Ethernet)
LAN	Local Area Network (Ethernet)
LBT	Local Backup Tunnel (eHCL)
LCU	Logical Control Unit (Mainframe)
LD	Long Distance (Fabric OS ISL type)
LDAP	Lightweight Directory Access Protocol (Security)
LED	Light Emitting Diode
LLDP	Link Layer Discovery Protocol (Ethernet)
LLL	Lossless Link Loss (Fabric OS Extension feature)
LPAR	Logical Partition (Mainframe)
LS	Virtual Fabrics: Logical Switch (Fabric OS feature)
LWL	Long Wave Length (fiber)
MAC	Media Access Control address (Ethernet)
MAN	Metro Area Network
Mbps Mb/s	Megabits per second
MMD	Modified Multiplicative Decrease (ARL)
MMF	Multi-Mode Fiber (fiber)
ms	millisecond
MT	Main Tunnel (eHCL)
MTU	Maximum Transmission Unit (Networking)
N_Port	FC Node port (FC)

Acronym	Expansion
NAS	Network Attached Storage (Storage)
NAT	Network Address Translation (Networking)
ND	IPv6: Neighbor Discovery (Networking)
NTP	Network Time Protocol (Networking)
NT-TCP	Non-Terminated TCP (Fabric OS Extension feature: Wtool and SLA)
NVMe	Non-Volitile Memory Express (Storage)
ODSS	BSL On-Demand SupportSave
OOO	Out of Order
OSTP	Open Systems Tape Pipelining (Fabric OS Extension feature)
PBR	Policy Based Routing (Networking)
PBR	Port Based Routing (FC)
Pfx	Prefix (Networking)
PG	Performance Group (Fabric OS FC feature)
PHY	Physical layer (OSI model)
PKI	Public Key Infrastructure (CAs and Certificates)
PMTU	Path MTU (Networking)
PRF	Pseudo Random Function (Encryption)
PSK	Pre-Shared Key (Encryption)
pWDM	Passive Wave Division Multiplexing (Networking)
QoS	Quality of Service (Networking)
QSFP	Quad Small Form Pluggable (optics)
RADIUS	Remote Authentication Dial-In User Service (Security)
RAS	Reliability Availability Serviceability (RASlog)
RBT	Remote Backup Tunnel (eHCL)
RCA	BSL Remote Configuration Assistance
RDA	BSL Remote Diagnostics Assistance
RDR	Remote Data Replication
RFC	Request For Comments (IETF)
RJ-45	Standard for 8 pin copper based jack
RTT	Round Trip Time
Rx	Receive
SA	Security Association (Encryption)
SAN	Storage Area Network
SAN	Subjective Alternative Name (Certificates)
SCSI	Small Computer System Interface (ANSI)
SDDQ	Slow Drain Device Quarantine (Fabric OS FC feature)
SDM	System Data Mover (Mainframe)
Sec	second
SFP	Small Form Pluggable (1, 2, and 4 Gbps)
SFP+	Enhanced Small Form Pluggable (8, 10, 16, and 25 Gbps)

Acronym	Expansion
SID	FC source PID
SLA	Service Level Agreement (Fabric OS Extension feature)
SMF	Single Mode Fiber (fiber)
SNMP	Simple Network Management Protocol (Networking)
SRC	Source
SSH	Secure Shell
STP	Spanning Tree Protocol (Ethernet)
SWL	Short Wave Length (fiber)
Syslog	A message logging server
TACACS+	Terminal Access Controller Access Control System (Security)
TCL	Traffic Control List (Fabric OS Extension feature)
TCP	Transmission Control Protocol (Networking)
TLV	Type Length Value
TO	Traffic Optimizer (Fabric OS FC feature)
TTL	Time To Live (IP)
Tx	Transmit
UDP	User Datagram Protocol (Networking)
VE	Virtual E_Port (Extension tunnel)
VF	Virtual Fabrics (Fabric OS FC feature)
VLAN	Virtual LAN (Networking)
VM	Virtual Machine
VMID	A VM’s unique ID
VRRP	Virtual Router Redundancy Protocol (Networking)
WAN	Wide Area Network (Networking)
WDM	Wave Division Multiplexing (Networking)
WO-TCP	WAN Optimized TCP (Fabric OS Extension feature)
xISL	Base-switch tagged FC ISL carrying multiple LS traffic

Revision History

The revision history provides a list of the significant changes made in each version of the document.

FOS-92x-Ext-UG100; April 28, 2023
Initial document version.

Documentation Legal Notice

This notice provides copyright and trademark information as well as legal disclaimers.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom 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 or to view the licensing terms applicable to the open source software, please download the open source attribution disclosure document in the Broadcom Support Portal. If you do not have a support account or are unable to log in, please contact your support provider for this information.


BROADCOM®