BROCADE
A Broadcom Inc. Company

USER GUIDE

Brocade Fabric OS Extension User Guide, 8.2.1

Supporting Fabric OS 8.2.1

FOS-821-Ext-UG101
28 September 2018

Contents

Introduction

About This Guide

This document describes Brocade Extension, the platforms involved, and the configuration steps that you need to perform.

What's New in This Document for Fabric OS 8.2.1

Changes Made for the Initial Release

Changes to this publication, which support Brocade Fabric OS 8.2.1, include the following:

- Support for the Brocade 7810 Extension Switch, including the new sections Brocade 7840 Extension Switch Ports on page 47 and Brocade 7810 License Options on page 56.
- Completely revised the discussion of Extension Hot Code Load on page 22.
- Under Configuring Global LLDP Parameters on page 97:

  For the TLV port-desc, the string is now of the format: "Switch Model Name (in the case of a fixed-port switch)/Slot Model Name (in the case of a chassis): Mode + Speed + Slot/Port" (for example, Brocade SX6: WAN 10G 4/ge13).

  For the TLV sys-desc, the default system description string in the TLV advertised for the switch is now of the format: "Switch Model Name, Firmware Version" (for example, Brocade 7840, Fabric OS Version 8.2.1).
- Under Configuring Static and Dynamic LAGs Using LACP on page 99:

  A group of ports can now be expressed as a range rather than as a string of ports separated by commas. For example:

```
switch#admin> lldp --enable -port 3/40-56 -profile lldp_profile_1
```

Supported Hardware and Software

The following hardware platforms support Brocade Extension (Fibre Channel over IP features and IP Extension features) as described in this configuration guide:

- Brocade 7840 Extension Switch
- Brocade 7810 Extension Switch
- Brocade X6-4 Director and Brocade X6-8 Director with one or more Brocade SX6 Extension Blades
- Brocade DCX 8510-4 Backbone and Brocade DCX 8510-8 Backbone with one or more Brocade FX8-24 Extension Blades

As described in this configuration guide, the software supported is Brocade Fabric OS 8.2.1.

Contacting Brocade Technical Support

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

Online	Telephone
For nonurgent issues, the preferred method is to go to MyBrocade (my.brocade.com) and then go to one of the following sites: • My Cases • Software Downloads • Licensing tools • Knowledge Base	Required for Severity 1-Critical and Severity 2-High issues: • North America: 1-800-752-8061 (Toll-free) • International: 1-669-234-1001 (Not toll-free) Toll-free numbers are available in many countries and are listed at https://www.broadcom.com/support/fibre-channel-networking/.

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

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

Document Feedback

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

Extension Concepts and Features

Brocade Extension Concepts

Brocade extension switches and extension blades for Brocade director families provide a fast and reliable network infrastructure to address the requirements of storage area networks (SANs) that are extended beyond the traditional reaches of Fibre Channel (FC) communications. Without some form of distance extension, the distance between the source and the destination in a SAN is limited to a few kilometers.

Brocade 7840 and Brocade SX6 extension products support both FC/FICON-based data flows and IP-based storage data flows. Brocade extension solutions maximize replication and backup throughput over distance, using data compression, disk and tape protocol acceleration, and WAN-optimized TCP. Brocade extension supports applications such as remote data replication (RDR), centralized backup, and data migration.

Brocade extension uses the existing IP wide area network (WAN) infrastructure to connect Fibre Channel and IP fabrics between distant endpoints that are either impractical or costly using native Fibre Channel or IP connections. The basis of the connection is the extension tunnel, built on a physical connection between two extension switches or blades. Extension tunnels allow Fibre Channel and IP traffic to pass through the IP WAN. The extension tunnel connections ensure lossless transmission and that FC and IP frames are delivered in the correct order. The Fibre Channel fabric and all targets and initiators, whether FC or IP, are unaware of the presence of the IP WAN.

The extension tunnel provides load balancing across separate network paths, optimization for extended links, rate limiting to ensure optimal performance, and lossless link loss (LLL) recovery.

The two Brocade protocol technologies implemented in Brocade extension products are FCIP and IP Extension.

FCIP, or Fiber Channel over IP, is a tunneling protocol to link Fibre Channel over distance on standard IP networks. Used primarily for remote replication, backup, and storage access, FCIP provides Fibre Channel connectivity over IP networks between Fibre Channel devices or fabrics. The FCIP link is an inter-switch link (ISL) that transports FC control and data frames between switches. The following table outlines FCIP.

TABLE 1 FCIP Protocol

Network	WAN/MAN
Transport	FCIP/TCP/IP/Ethernet
Encapsulation	Brocade encapsulates Fibre Channel data sequences into compressed batches. Those batches fill TCP segments to their maximum size and then form IP datagrams.
IP-routable	Yes

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 uses the same VE_Ports and circuits that FCIP uses, or it can use its own. The following table outlines IP Extension.

TABLE 2 IP Extension Protocol

Network	WAN/MAN
Transport	IP Extension/TCP/IP/Ethernet
Encapsulation	Brocade encapsulates IP data flows (also called “streams”) into compressed batches. Those batches fill TCP segments to their maximum size and then form IP datagrams.
IP-routable	Yes

For additional information about implementation and technical details of Brocade extension technology, refer to the Brocade white paper Extension Trunking.

Extension Trunks, Tunnels, Circuits, and Interfaces

An extension tunnel is a conduit that contains one or more circuits. When a tunnel contains multiple circuits, it is also called an extension trunk because multiple circuits are trunked together. An extension tunnel, or extension trunk, is a single inter-switch link (ISL). Circuits are individual extension connections within the trunk, each with its own unique source and destination IP interface (IPIF) address.

To understand tunnels, you must understand the relationship between tunnels and VE_Ports. Because an extension tunnel is an ISL, each end of the tunnel requires its own Virtual E_Port (VE_Port). For example, the Brocade SX6 Extension Blade supports a number of VE_Ports. The available VE_Ports on the Brocade SX6 are numbered 16 to 35. Tunnels are more complicated than this, but the point is that each end of the tunnel is identified by number, and that number is directly associated with a VE_Port on the extension platform. Tunnels are frequently created between different VE_Ports, so from a configuration point of view, the tunnel number can be different at each end. Each extension platform, such as the Brocade 7810, the Brocade SX6 Extension Blade, Brocade 7840 Extension Switch, and Brocade FX8-24 Extension Blade, has a different numbering scheme for its VE_Ports.

Circuits exist within tunnels. A circuit is a connection between a pair of IP addresses that is defined within an extension tunnel. Circuits provide the links for traffic flow between source and destination interfaces that are located on either end of the tunnel. Each tunnel can contain a single circuit or multiple circuits.

NOTE

In this publication, the “source” or “local” is the switch that you are configuring, whereas the “destination” or “remote” is the switch on the other end of the tunnel or circuit. Local switch and remote switch will depend on which side of the of the tunnel you are configuring.

You must configure unique IPIFs as the local and remote endpoints of each circuit. On the local side, a circuit is configured with a source IPIF address and a destination address. On the remote side of the circuit, its source IPIF address is the same as the local-side destination address. Similarly, on the remote side of the circuit, its IPIF destination address points to the local-side source address. Multiple IPIFs can be configured on each physical Ethernet interface.

NOTE

On the Brocade FX8-24, the IP address (or IPIF) for each local and remote address must be individually unique. However, on the Brocade 7810, Brocade 7840 and Brocade SX6, each address pair must be unique. For example, the following address pairs use the same source address in each pair, but the destination addresses are different. For the Brocade FX8-24, these addresses are not unique. For the Brocade 7810, Brocade SX6, and Brocade 7840, the address pairs are 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

The circuit configuration parameters on the local and remote sides of the circuits and tunnel must match, in addition to the source and destination IPIF addresses pointing to each other. For example, if you configure IPsec on a tunnel, each end of the tunnel must be configured to use the same IPsec parameters. Other parameters for each circuit must match, such as MTU size, bandwidth allocation, QoS, VLAN ID, and keepalive values.

You must configure a gateway IP route for the circuit to the destination network when the remote IPIF is not on the same subnet as the local IPIF. You can define a specific number of routes per IPIF based on the extension platform. See Tunnel and Circuit Requirements for Brocade Extension Platforms on page 62 for specifications.

ATTENTION

When using Brocade IP Extension, the local and remote LAN subnet addresses must be different.

The following figure shows an example of two extension tunnels. The first tunnel is a trunk of four circuits, and the second tunnel is a trunk of two circuits. Each circuit is assigned a unique IPIF address. Those IPIFs are, in turn, assigned to Ethernet interfaces. In the figure, each IPIF is assigned to a different Ethernet interface. This is not required. Ethernet interface assignment is flexible depending on the environment’s needs, and assignments can be made as desired. For instance, multiple IP interfaces can be assigned to a single Ethernet interface. The circuit flows from IP interface to IP interface through the assigned Ethernet interfaces.

FIGURE 1 Extension Tunnel and Circuits

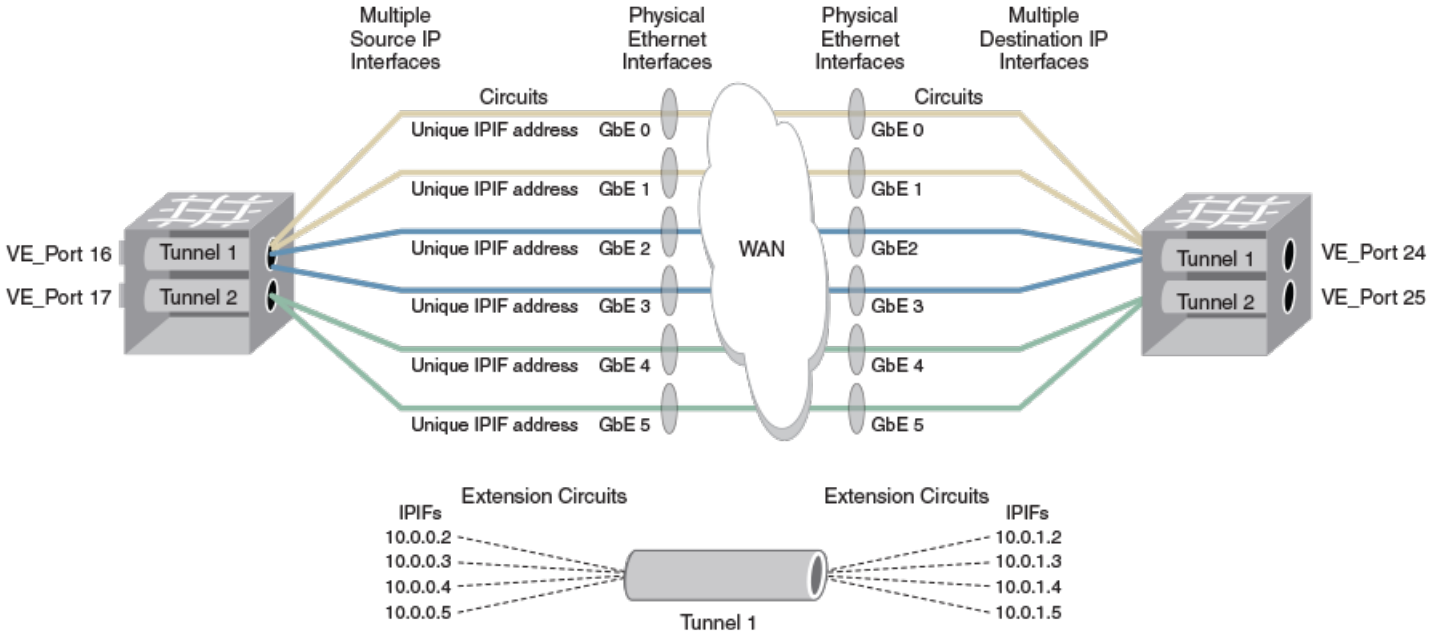

For specifications and restrictions on tunnels, circuits, and trunks for the Brocade 7810 Extension Switch, Brocade 7840 Extension Switch, SX6 Extension Blade, and FX8-24 Extension Blade, see the "Extension Platforms and Features" chapter in this guide.

VE_Ports and VEX_Ports

When Fibre Channel communications between switches includes fabric services, they must occur by means of an inter-switch link (ISL). ISLs are supported on E_Ports. Because an extension tunnel is an ISL, each end of that tunnel requires its own E_Port. There are various types of E_Ports:

- E_Port
- EX_Port
- VE_Port (virtual E_Port)
- VEX_Port (virtual EX_Port)

VE_Ports and VEX_Ports are virtual because they face the extension tunnel and thus enable communication across an extension tunnel. Extension trunking—multiple circuits in a tunnel—operates the same whether the virtual port is a VE_Port or a VEX_Port.

> **NOTE**
> VEX ports are supported only on the Brocade FX8-24 Extension Blade. The Brocade 7840, Brocade 7810, and Brocade SX6 do not support VEX_Ports.

EX and VEX ports are FC-routed (FCR) ports. Routed ports are the demarcation point of fabric services for a fabric. Fabric services do not extend beyond an EX or VEX port. Remote edge fabrics are edge fabrics connected through a WAN connection and VEX_Port.

TABLE 3 VE_Ports and VEX_Ports

E_Port Types	No FCR	FCR
Native FC	E_Port	EX_Port

TABLE 3 VE_Ports and VEX_Ports (continued)

E_Port Types	No FCR	FCR
Extended over tunnel	VE_Port	VEX_Port (on FX8-24 only)

Because an extension trunk is logically a *single* tunnel, only a single VE_Port or VEX_Port is used for each end of the tunnel, even though more than one circuit can be contained within the tunnel.

Once the tunnels are configured and the WAN-optimized TCP (WO-TCP) connections are made for a circuit, a logical inter-switch link is established between the switches. VE_Ports operate like E_Ports for all fabric services and Fabric OS operations, except that VE_Ports use TCP/IP and Ethernet as the transport instead of FC.

From the point of view of a Remote Edge Fabric, a VEX_Port appears as a normal E_Port. It follows the same Fibre Channel protocol as other E_Ports. However, VEX_Ports terminate the attached fabric at the port and do not allow fabrics to merge by propagating fabric services or routing topology information across the WAN to the Remote Edge Fabric. This provides edge-fabric or remote-edge-fabric isolation outward from the EX_Port or VEX_Port, respectively.

> NOTE
> VE_Ports or VEX_Ports cannot connect in parallel to the same domain at the same time as Fibre Channel E_Ports or EX_Ports.

An extension tunnel is assigned to a VE_Port or VEX_Port on the switch or blade at each end of the tunnel. Because multiple VE_Ports are supported on the extension switch or blade, you can create multiple tunnels on a switch or blade. There is no requirement for VE_Port numbers to be identical on each end of a tunnel.

Fibre Channel frames enter an extension tunnel through the VE_Ports (or VEX_Ports) and are encapsulated and passed to TCP layer connections. A data processing (DP) complex on the switch or blade handles the FC frame encapsulation, de-encapsulation, and transmission to the TCP link.

Gigabit Ethernet Interfaces

On the Brocade extension platforms, the Ethernet interfaces are abstracted from the VE/VEX_Ports, IP interfaces, and extension tunnels and trunks. This means that the Ethernet interfaces are not fixed to the VE/VEX_Ports, IP interfaces, or extension tunnels and trunks. Depending on the platform, you have access to a combination of 1-GbE interfaces, 10-GbE interfaces, and 40-GbE interfaces.

A GbE interface can be configured for WAN operations or LAN operations on platforms that support Brocade IP Extension. A GbE interface is used for WAN operation when it functions as the endpoint of a circuit. The interface is WAN-facing, and you create the tunnels and circuits, assigning ipifs (the circuit endpoints) to the GbE interface you desire to use. WAN is the only allowed GbE interface mode when a platform is configured for FCIP-only operation.

A GbE interface is used for WAN operation when it functions to pass circuits through it.

When a platform is configured for Hybrid mode, which supports both FCIP and IP Extension, you configure certain LAN-facing GbE ports for LAN mode operation. When GbE ports are in LAN mode, they can be members of a static link aggregation group (LAG) or dynamic LAG.

Ethernet Interfaces

Ethernet interfaces are assigned by associating them with IP interfaces (IPIFs). IP interfaces on the extension platforms are virtual entities within the platform. The IPIFs, via their IP addresses, are assigned to circuits by designating the source IP address when creating the circuit. The circuit is grouped into an extension trunk by designating the VE_Port (or VEX_Port) to which it belongs. Tunnels/trunks are identified by their VE_Port number. If an extension trunk already exists, you can create an additional circuit and include it in the tunnel.

Each Ethernet interface on the Brocade SX6, Brocade 7840, Brocade 7810, and Brocade FX8-24 can be used by circuits from multiple VE/VEX_Ports. Multiple VEs in different FIDs can share a Ge port provided it is in the default switch. If a Ge port is part of a non-default switch, however, then it can only be shared by VEs in that FID.

Each VE/VEX_Port can use multiple Ethernet ports. Each VE/VEX_Port can have multiple IP interfaces, or circuits. Each IP interface or circuit in a VE/VEX_Port is associated with an IP address. In turn, those IP addresses are assigned to specific Ethernet interfaces. However, a single circuit cannot span multiple Ethernet interfaces.

For Brocade Extension Trunking, the data processor (DP) that owns the VE_Port controls all member circuits. There is no distributed processing, load sharing, or LLL across DPs. Failover between DPs is done at the FC level by the Brocade FC switching ASIC, provided that the configuration permits it. The only components that can be shared are the Ethernet interfaces.

The IP network routes the circuits over separate network paths and WAN links based on the destination IP addresses, VLAN tagging, and possibly other Layer 2 and Layer 3 header attributes. Ethernet interfaces on the Brocade 7840, Brocade 7810, Brocade SX6, and the Brocade FX8-24 provide a single convenient connection to the data center LAN for one to many circuits.

Virtual Circuits and Tunnels

Upon creation, an extension tunnel consists of one or more circuits configured within the same tunnel. A VE_Port is a tunnel endpoint and a virtual tunnel is the actual object for transporting data of a particular traffic class. When an extension platform is configured to operate in Hybrid mode, which supports, both FCIP and IPEX, high, medium, and low data virtual-tunnels are created for both protocols. Created within each circuit member of the tunnel are four (FCIP only) or seven (FCIP and IPEX) virtual-tunnels (the F-Class tunnel, three IPEX QoS data virtual-tunnels, and three FCIP QoS data virtual-tunnels).

A tunnel consists of these virtual-tunnels:

- **Control virtual-tunnel**, which includes F-Class traffic. The control tunnel is "higher" priority and has no guaranteed minimum bandwidth. As needed, it can use all available bandwidth up to the maximum configured bandwidth rate although this traffic utilizes very little bandwidth.
- **High priority data virtual-tunnel**. The high data tunnel is for high-priority QoS data, and via Per-Priority-TCP-QoS has its own WO-TCP for this traffic class.
- **Medium priority data virtual-tunnel**. The medium data tunnel is for medium-priority QoS data, and via Per-Priority-TCP-QoS has its own WO-TCP for this traffic class.
- **Low priority data virtual-tunnel**. The low data tunnel is for low-priority QoS data, and via Per-Priority-TCP-QoS has its own WO-TCP for this traffic class.

QoS settings define the allotment of bandwidth during times of contention. The default settings are 50/50% between FCIP/IPEX, and within each of those there are the default allocation settings of 50/30/20% (the minimum configuration) respectively. You can change these settings as needed. If contention is absent, bandwidth to the virtual-tunnels is unrestricted; the available bandwidth can be used by any virtual-tunnel requesting it. These configurable bandwidth allotments apply only when the data waiting to be sent exceeds the available bandwidth to send it (termed "contention in the egress queue").

Data virtual-tunnels adhere to the minimum bandwidth configurations and will always receive those amounts. For example, a 10 Gb/s circuit used for both FCIP & IPEX and experiencing contention would by default reserve 1.5 Gb/s for a medium FCIP virtual-tunnel (10

Gb/s x 50% x 30% = 1.5 Gb/s). The lesson is that bandwidth allocations are not oversubscribed. Virtual-tunnels can consume at most the maximum communication rate configured on the circuit provided other virtual-tunnels are not currently using their allotted bandwidth.

NOTE
The Brocade FX8-24 does not support IPEX; it creates only four FCIP virtual-tunnels (1 control and 3 data).

An additional tunnel group is created for high availability (HA) when eHCL is configured on the Brocade SX6 or Brocade 7840. That tunnel group contains a main tunnel, a local backup tunnel (LBT), and a remote backup tunnel (RBT). When eHCL is active; the high, medium, and low QoS data tunnels collapse into a single QoS data tunnel for the duration of the eHCL update. For additional information, refer to Extension Hot Code Load on page 22.

FCIP Extension

FCIP, or Fiber Channel over IP, is a tunneling protocol to link Fibre Channel over distance on standard IP networks. Used primarily for remote replication, backup, and storage access, FCIP extension provides Fibre Channel connectivity over IP networks between Fibre Channel devices or fabrics. The FCIP link is an inter-switch link (ISL) that transports FC control and data frames between switches. FCIP extension is supported on all Brocade extension platforms.

FCIP extension provides the following:

- Adds protocol optimizations for extended distance performance for both FICON and FCP/SCSI flows that can be enabled on an extension tunnel.
- Network resiliency using Extension Trunking, which is a single logical tunnel comprised of one or more individual circuits
- Efficient use of VE_Ports
- High performance for high-speed WAN links (one or more 10-Gb/s and 40-Gb/s links)
- TCP Acceleration with WAN Optimized TCP (WO-TCP)
- Aggregation of circuit bandwidth
- Failover/failback
- Failover groups and metrics
- Spillover
- Use of disparate characteristic WAN paths
- Non disruptive link loss
- Lossless link loss (LLL)
- Adaptive Rate Limiting (ARL)
- In-Order Delivery (IOD)
- Deterministic path for protocol acceleration
- WAN bandwidth pooling—pool bandwidth from multiple links/providers
- High-speed compression using deflate
- High-speed IPsec (AES 256)
- Diagnostic and troubleshooting tools: SAN Health and WAN Tool
- QoS for FCIP with DSCP and/or 802.1P marking and enforcement
- 9216-byte jumbo frame support on the Brocade 7840, Brocade 7810, and the Brocade SX6 blade

Within the architecture of the Brocade 7840, Brocade 78100, Brocade SX6, and Brocade FX8-24, you find Brocade FC application-specific integrated circuits (ASICs) that know only the FC protocol. The VE/VEX_Ports are logical representations of actual FC ports. VE/

VEX_Ports are not part of the ASICs. Multiple FC ports feed a VE/VEX_Port, permitting high data rates well above 8 gigabits per second (Gb/s), 16 Gb/s, and 32 Gb/s, which is necessary for feeding the compression engine at high data rates and for high-speed trunks.

On the WAN side

- The Brocade 7840 and Brocade SX6 support 10-Gigabit Ethernet (GbE) and 40-GbE interfaces.
- The Brocade 7810 support up to six 1/10GbE optical interfaces or two 1GbE Copper interfaces and four 1/10GbE optical interfaces.
- The Brocade FX8-24 supports 10-GbE and 1-GbE interfaces.

Think of VE and VEX_Ports as the transition point from the FC world to the TCP/IP world inside the extension devices.

IP Extension

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 7840 Extension Switch, Brocade 7810 Extension Switch, and the Brocade SX6 Extension Blade platforms.

The Brocade 7840 Extension Switch and the Brocade SX6 Extension Blade support IP-based storage data flows as well as FCP/SCSI and FICON-based data flows. The Brocade 7810 Extension Switch supports IP-based storage data flows as well as FC-based data flows but does not support FICON-based data flows.

IP data flows across Brocade tunnels are referred to as IP Extension (termed *IPEX*), which enables you to use existing IP WAN infrastructure to connect IP storage. Additionally, IP Extension gives you visibility into and control of flows using various diagnostic tools, IPsec, compression, QoS, extension trunking, and lossless tunnel resiliency. IP Extension supports applications such as array-native IP remote data replication (RDR), IP-based centralized backup, and VM replication. In addition, IP Extension supports host-based and database replication over IP, NAS head replication between data centers, and data migration between data centers.

Brocade WAN Optimized TCP (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. This accelerates end device native TCP. IP Extension data across the IP WAN uses WO-TCP, a highly efficient and aggressive TCP stack for moving data between data centers.

IP Extension provides the following advantages:

- Data Center Interconnect (DCI): Unified support and management of both FC/FICON and IP
- Storage administrators: Provision once and over time connect many devices
- High performance for high-speed WAN links (one or more 10-Gb/s and 40-Gb/s links)
- WAN bandwidth pooling: pool bandwidth from multiple links/providers
- Lossless link loss (LLL)
- Adaptive Rate Limiting (ARL)
- Network resiliency using Extension Trunking
- Efficient protocol transport: negligible added overhead
- TCP Acceleration with WAN Optimized TCP
- Streams: Virtual windows on WAN Optimized TCP to eliminate head of line blocking (HoLB)
- High-speed compression using deflate
- High-speed IPsec (AES 256)
- Diagnostic and troubleshooting tools: SAN Health and WAN Tool
- Separate QoS for both FCIP and IP Extension with DSCP and/or 802.1P marking and enforcement

- Internal prioritization of FC traffic versus IP traffic with QoS distribution
- 9216-byte jumbo frame support for LAN and WAN networks

IP Extension and IP Networking

The following are key points to understand about IPEX and IP networking:

- IPEX extends only Layer 3 networking; it does not extend Layer 2 networking.
- IPEX requires different broadcast domains on each end of the extension tunnel. This means that a single subnet cannot span an IPEX tunnel.
- The Brocade 7840 Extension Switch, Brocade 7810 Extension Switch, and Brocade SX6 Extension Blade act as a next-hop gateway for IP data flows traversing two data centers. This is managed by creating the IPEX LAN interfaces plus configuring (on the end-devices) specific routes for the remote subnet that points to the IPEX gateway interface (ipif).
- The LAN interfaces are on the local side of the network in which the end-devices are attached, typically via a local Ethernet data center switch.
- In contrast, WAN interfaces are on the side of the network that sends traffic outside of the data center potentially communicating with multiple routers before arriving at the final destination. In the following illustration, Data Center A and Data Center B are examples of LANs that communicate through a WAN.
- IP Extension supports dynamic and static link aggregation groups (LAGs) on the LAN interface. A LAG is a group of physical Ethernet interfaces that are treated as a single logical interface.
- Beginning with Fabric OS 8.2.0, IPEX supports neighbor discovery through Link Level Discovery Protocol (LLDP) on the GbE interfaces. This applies to both the WAN and LAN interfaces.
- Beginning with Fabric OS 8.0.1, IPEX supports policy-based routing (PBR), which allows a router to be connected to the LAN side of an IPEX platform.
- IPEX uses traffic control lists (TCLs) to filter traffic to the correct destination. The default action (no TCL is configured) is to deny all traffic. Configuring TCLs on each extension switch or blade allows traffic to pass from endpoint to endpoint through a designated tunnel.

  NOTE
  Creating a TCL is critical, even if only a single tunnel (VE_Port) is being implemented. IPEX will not function without a configured TCL.

NOTE
IPEX does not allow FIPS mode, and when a platform is operating in FIPS mode, IPEX is disallowed. This restriction does not apply to the Brocade 7810 switch because FIPS is supported on the CP side and this switch supports single-applicable Hybrid mode (includes IP Extension).

FIGURE 2 Extension Tunnel Concept and TCP/IP Layers for FCIP and IP Extension

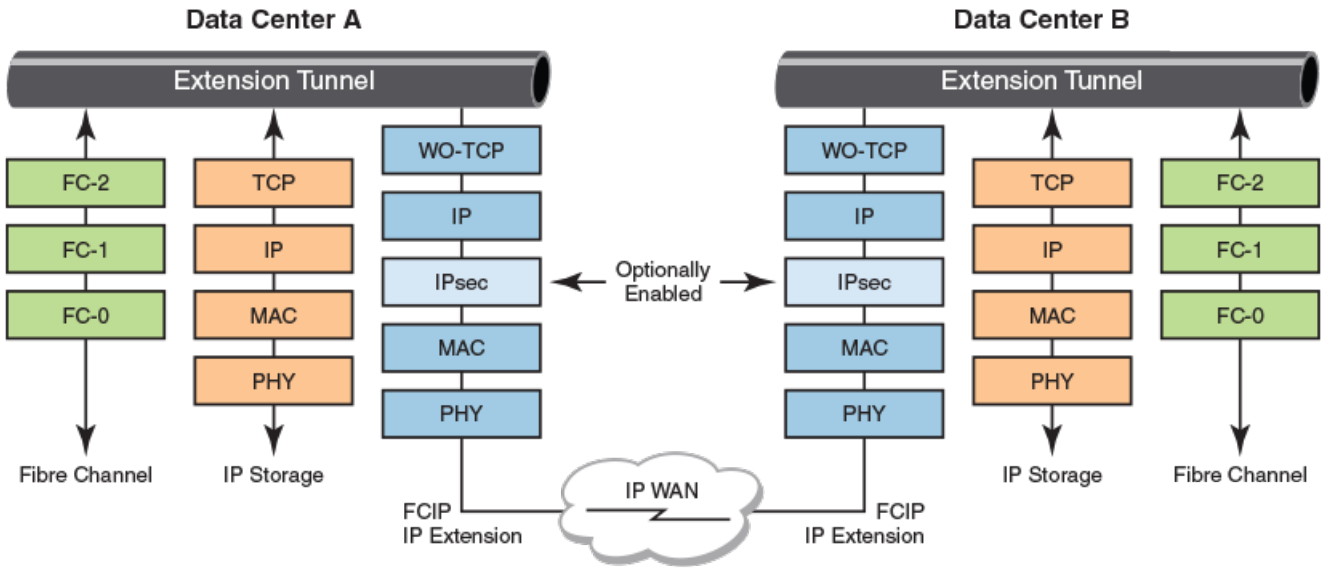

IPEX forwarding and processing are done within the extension platform, the field-programmable gate arrays (FPGAs), and the data processors (DPs). The VE_Port is a logical endpoint for a tunnel, and IPEX uses them as well, even if no FC or FICON traffic is being transported. IPEX uses switch virtual interfaces (SVIs) on each DP as a communication port with the IP storage end devices. The SVI becomes the gateway IP interface for data flows that are headed toward the remote data center. At the SVI, end-device TCP sessions are locally terminated and reformed on the remote side (termed "TCP proxying"). 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.

Which side is considered the local or remote depends on where the TCP session initiates. On the Brocade 7840 switch and Brocade SX6, eight of the ten GbE interfaces available for IPEX LAN side connectivity are connected to end devices or data center LAN switches. The Brocade 7810 Extension Switch supports four LAN ports only. The Brocade 7840 switch and Brocade SX6 have eight 10-GbE interfaces reserved for WAN-side (tunnel) connectivity whereas the Brocade 7810 Extension Switch has six GbE interfaces.

Extension Trunking

Extension Trunking is an advanced feature of the Brocade Extension platforms for both FCIP and IPEX that enables bandwidth aggregation and lossless failover for increased resiliency over IP WANs. It is a way of managing the use of WAN bandwidth and provides redundant paths over the WAN that can protect against transmission loss due to WAN failure. Furthermore, Extension Trunking provides granular load balancing per batch weighted round-robin.

Trunking is enabled by creating multiple circuits within a tunnel, members of a single VE_Port. The tunnel will utilize multiple circuits to carry data between a source and destination data center. For circuit capacities on Brocade Extension switches and blades, see Tunnel and Circuit Requirements for Brocade Extension Platforms on page 62.

Redundancy and Fault Tolerance

Multiple extension tunnels can be defined between pairs of extension switches or blades, but doing so defeats the benefits of a single multiple-circuit extension tunnel. Defining two tunnels between a pair of switches or blades is not as redundant or fault-tolerant as having multiple circuits in one tunnel.

Extension Trunking provides lossless link loss (LLL). LLL ensures that all data lost in flight is retransmitted and reordered before being delivered to upper layer protocols. This is an essential feature to prevent interface control checks (IFCCs) on mainframes that use FICON and SCSI timeouts for open-system-based replication. For more information about LLL on specific Brocade extension switches and blades, see Circuit Failover on page 78.

NOTE
When you create multiple parallel tunnels between the same switch domains, you must enable lossless dynamic load sharing (DLS). This is because there can be routing updates that occur when tunnels come up or go down. Each routing update can cause dropped or unrouteable frames if the destination is to a domain via a switch connected through a peer tunnel.

Failover Circuits and Groups

Failover circuits and groups are a feature supported on the Brocade 7840 Extension Switch, Brocade 7810 Extension Switch, Brocade FX8-24 Extension Blade, and the Brocade SX6 Extension Blade. Failover groups allow you to define a set of metric 0 and metric 1 circuits that are part of a failover group. When all metric 0 circuits in the group fail, metric 1 circuits take over operation, even if there are metric 0 circuits still active in other failover groups. Typically, you would configure only one metric 0 circuit in a failover group. All metric 0 circuits in a group must fail before a metric 1 circuit is used.

For additional information about circuit failover configuration, see Circuit Failover on page 78.

Spillover Circuits

The spillover feature is supported on the Brocade 7840 Extension Switch, Brocade 7810 Extension Switch, and the Brocade SX6 Extension Blade. Spillover lets you configure a set of primary circuits to use all the time and a set of secondary circuits to use only during high-utilization periods. When a tunnel is configured for spillover, it will run over the lower metric circuits (metric 0) until the bandwidth utilization is reached on those circuits. When the lower metric bandwidth is exceeded, the remaining traffic is sent over the higher metric 1 circuits.

For example, consider three 1Gb/s circuits that are configured on a tunnel.

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

For more information about spillover circuit configuration, refer to Configuring Spillover on page 138.

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

Service-Level Agreement

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 the Brocade 7840 Extension Switch, Brocade 7810 Extension Switch, and the Brocade SX6 Extension Blade. When using an SLA, you can define a set of network service-level parameters, such as packet loss, and make certain that the network meets those parameters before bringing the circuit into service.

When configured, the SLA automatically invokes WAN Tool and runs the test on the circuit. 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 enabled and placed into service. If the circuit goes down (bounces), it reverts to test mode and the SLA parameters must be met before the circuit is enabled again and placed in service.

For more information about SLA configuration, see Configuring a Service-Level Agreement on page 108.

Considerations for Multiple Tunnel Use with Protocol Optimization

If you want to use multiple extension tunnels within the same switch with protocol optimization features, you can use either Virtual Fabric Logical Switch/Logical Fabric (LS/LF) configurations (best practice) or traffic isolation (TI) zones (not recommended)

Protocol optimization features include the following:

- FastWrite
- Open Systems Tape Pipelining (OSTP)
- FICON Acceleration

These features require deterministic FC frame routing between initiators and targets when multiple tunnels or VE_Ports exist. TI zones and LS/LF configurations provide deterministic flows between the switches. Noncontrolled, parallel (equal-cost multipath) tunnels are not supported between domains when protocol optimization is enabled on one or more tunnels unless you limit the routing of source ID and destination ID (SID/DID) pairs to a specific tunnel by using TI zones or Virtual Fabric (VF) LS/LF configurations.

Note the following additional restrictions:

- The recommended best practice is to have identical Fabric OS versions at both ends of an extension tunnel.
- When planning Fabric OS upgrades or downgrades, it is recommended that you upgrade or downgrade both endpoints of an extension tunnel with the same Fabric OS version.
- When configuring tunnels to support large numbers of devices, consider the memory limitations of the Brocade extension switch or blade if you are enabling any protocol optimization feature. If too many devices are present or activated at one time, protocol optimization, such as FICON Acceleration, can be negatively impacted. Refer to Memory Use Limitations for Large-Device Tunnel Configurations on page 37.
- We strongly discourage use of Traffic Isolation Zones when they include VE ports with advanced features enabled (FICON XRC, FICON Tape, OSTP or FW).

IP WAN Network Considerations

Because Brocade extension tunnels use TCP connections over an existing wide area network (WAN), consult with the WAN carrier and IP network administrator to ensure that the network hardware and software equipment operating in the data path can properly support the TCP connections. Keep the following considerations in mind:

- Routers and firewalls that are in the data path must be configured to pass traffic through a specific TCP port on the switch.
  - On all platforms, if IPsec is used, the network must allow both Encapsulating Security Payload (ESP) traffic and UDP port 500 (IKE) traffic to pass through.
  - On the Brocade 7840 switch, Brocade 7810 switch, and the Brocade SX6 blade, the Brocade WO-TCP implementation selects a port between 49152 and 65535 as the ephemeral (or initiating) port to open up to port 3225 and 3226.
- On the Brocade 7840 switch, Brocade 7810 switch, and the Brocade SX6 blade, the TCP URG flag is frequently set. Ensure that these flags are not dropped or modified from ports 3225 and 3226.
- The Brocade FX8-24 blade uses TCP port 3225.
- The Brocade 7810 switch Brocade 7840 switch, and the Brocade SX6 use TCP ports 3225 and 3226.
- To enable recovery from a WAN failure or outage, be sure that diverse, redundant network paths are available across the WAN.

- Be sure that the underlying WAN infrastructure can support the redundancy and performance expected in your implementation.
- If you use jumbo frames on the Brocade 7810 switch, Brocade 7840 switch, or the Brocade SX6 blade, ensure that the WAN network will support it.

  NOTE
  The IP DF (Do Not Fragment) bit is set on Brocade Extension Tunnels due to the inefficiency of fragmentation.

- Use AUTO MTU to avoid any unexpected tunnel behavior, which might occur happen when the WAN network does not support hard set MTU values for the IPIF.

IP LAN Network Considerations

The following considerations apply when configuring the extension platform for IP Extension support:

- The LAN interface functions as a virtual switch, or gateway. That is why it is referred to as a switch virtual interface (SVI).
- IP extension forms an extension tunnel through the WAN to allow the local LAN to communicate with the remote LAN.
- The SVI on the local extension platform and remote extension platform must be on different subnets. Otherwise, the IP traffic will not be forwarded.
- When configuring the IP route on the extension platform, the SVI is the next-hop gateway.
- IP extension optimizes the TCP traffic. However, if you are doing UDP forwarding or using IPsec in your LAN, such traffic will not be optimized.
- Double VLAN tagging (sometimes referred to as nested VLAN tagging or QinQ is not supported.

Extension Hot Code Load

Extension Hot Code Load (eHCL) tunnels are supported on the Brocade 7840 Switch and the Brocade SX6 Blade. An eHCL tunnel provides high-availability (HA) support and allows non disruptive firmware updates. eHCL tunnels require four IP addresses per circuit, two at each endpoint. Each IP address has its own IP Interface (IPIF). The four addresses are the traditional local and remote IP addresses for the production circuit and the local and remote HA IP addresses used by eHCL during firmware updates. Typically, the local pair of IP addresses is included in the same subnet, and the remote pair of IP addresses is included in the same subnet, although this is not a requirement. All IP addresses must be able to communicate across the IP infrastructure.

As the Brocade 7810 Switch has only one DP, it does not support eHCL on itself. However, if a Brocade 7810 switch is connected to a Brocade 7840 switch or Brocade SX6 blade, the Brocade SX6 blade will support eHCL although the Brocade 7840 switch will not.

NOTE
Consequently, the `--local-ha-ip` option for `portcfg fciptunnel|fcipcircuit` will not be configurable on a Brocade 7810 Switch. However, you can still use the `--remote-ha-ip` option to allow a Brocade 7810 Switch to connect to a Brocade 7840 Switch or Brocade SX6 Blade for remote eHCL. When the Brocade 7810 Switch is configured for an attached Brocade 7840 Switch or Brocade SX6 Blade, only the `--local-ha-ip` option should be used. If you use the `--remote-ha-ip` option on the Brocade 7840 Switch or Brocade SX6 Blade, the tunnel will never reach an ONLINE state, remaining in a DEGRADED state. Only the main tunnel (MT) group and remote-backup tunnel (RBT) group will appear on the Brocade 7810 Switch, and only the main tunnel (MT) group and local-backup tunnel (LBT) group will be found on the connected Brocade 7840 Switch or Brocade SX6 Blade.
Please be aware that the `portshow fciptunnel --hcl-status` command functions on a Brocade 7810 Switch even though it does not support eHCL. This command on the Brocade 7810 Switch provides DP connectivity information for the MT and RBT groups.

Extension HCL Operation

eHCL allows non disruptive firmware updates on the Brocade 7840 Switch and the Brocade SX6 Blade for FCIP and IPEX traffic over extension tunnels configured for HA. eHCL benefits mainframe environments by supporting nonstop connectivity for applications such as replication and tape grids. eHCL maintains extension connectivity across the WAN during firmware updates without disrupting active I/O, the loss of data, or allowing data to become out of order. eHCL is not supported on the Brocade 7810 switch or the Brocade FX8-24 blade.

eHCL supports serial upgrades. Recommended practice is to perform firmware upgrades on switches at the local site, followed by upgrades on the switches at the remote site. Beginning with Fabric OS 8.2.0, concurrent (or parallel) eHCL is supported on Gen 6 chassis. Refer to Extension HCL Enhancements in Fabric OS 8.2.0 on page 26 for additional information.

> NOTE
> When using Hybrid mode in Fabric OS releases prior to Fabric OS 8.1.0, eHCL is disruptive to IPEX traffic and non disruptive to FCIP traffic.

The Brocade 7840 switch and the Brocade SX6 blade each have two data processor (DP) complexes, referred to as DP0 and DP1. (See Network DP Components on page 51). An Extension HCL firmware update occurs on one DP complex at a time. When initiated, the update always starts on DP0. Before DP0 is updated to the new firmware, traffic fails over to DP1 to maintain communication between the local and remote switch. The failover process guarantees no loss and in-order delivery of data.

eHCL uses three tunnel groups contained within the extension tunnel (see the following figure) to perform the non disruptive update. Consider the main data center "local" and the DR site as "remote." The perspective in this section stays location consistent. The local backup tunnels (LBTs) and remote backup tunnels (RBTs) are automatically created when you configure an eHCL HA circuit. There is no requirement for all circuits within a tunnel to be eHCL-enabled. Some circuits can be eHCL-enabled and some might not.

To create the circuits for a tunnel, you must configure IP addresses on ipif (IP interfaces) for both the local DP0 and DP1, and for the remote DP0 and DP1. The ipifs you create are for the main tunnel (local DP0 to remote DP0), the local backup tunnel (local DP1 to remote DP0), and the remote backup tunnel (local DP0 to remote DP1). You create the production circuit and the ipifs not the backup circuits. When you configure these circuits, you designate which IP addresses to use with high availability (HA) via arguments in the `portcfg` command. This collection of circuits and tunnels form the *HA tunnel group*.

When eHCL is properly configured and a main tunnel exists between local DP0 and remote DP0.the following events occur during the firmware upgrade.

> NOTE
> All FCIP and IPEX traffic is maintained without interruption during the upgrade process.

- The main tunnel (MT) group provides connectivity during normal operations.
- A local backup tunnel (LBT) group maintains connectivity from the local switch DP1 to the remote switch DP0 while the former is upgrading. This tunnel, dormant during non-eHCL, is created automatically on the local DP1.
- A remote backup tunnel (RBT) maintains connectivity from remote switch DP1 to the local switch DP0 while the remote switch DP0 is upgrading. This tunnel group, dormant during non-eHCL operation, is created automatically on the remote DP1 when the corresponding local HA IP address is defined.

Creating the extension tunnel (with the`portcfg fciptunnel` command) creates the main tunnel (MT), which carries traffic through the extension tunnel to the remote switch. The LBT is created upon specifying the local HA IP address for the circuit, whereas the RBT is created upon specifying the remote HA IP address for the circuit. All three tunnel groups (MT, LBT, and RBT) are associated with the same VE_Port on the local and remote DP complexes. When an extension tunnel is configured to be eHCL capable, the LBT and RBT tunnel groups are always present. These connections are established when a tunnel and its circuit(s) are created or upon the switch booting if they already exist.

NOTE
The QoS traffic for high, medium, and low priorities is collapsed into a single QoS tunnel for the duration of eHCL operation. QoS priorities are restored when the eHCL operation completes. If the tunnel is configured for Differentiated Services Code Point (DSCP), all traffic is tagged with medium DSCP when it enters the WAN.

These tunnel groups are utilized in the following Extension HCL 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.
    a. Perform feature disable processing on the MT on DP0.
    b. Traffic from the MT on DP0 is rerouted to DP1 through the LBT to the remote switch. In-order data delivery is maintained.
    c. DP0 reboots with the new firmware, and the configuration is reloaded.
    d. 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.
    a. Perform feature disable processing on the MT on DP1.
    b. Traffic from the MT on DP1 is rerouted to DP0 through the LBT so that data traffic can continue between the switches. In-order data delivery is maintained.
    c. DP1 reboots with the new firmware, and the configuration is reloaded.
    d. Traffic from the LBT is failed-back to DP1 through the MT.
5. After the firmware is updated on DP1 and all MTs, the LBT, and the RBT are online, the Extension HCL firmware update is complete.

During the update process, tunnels and trunks change state (up or down). The MT provides connectivity during normal operations. It is up during normal operation and down only during the eHCL process. The RBT and LBT are normally up during normal operation, but do not handle non-HA production traffic. They operate to handle traffic only during the eHCL process. RBT handles traffic when the remote switch DP0 undergoes the eHCL process. The RBT is visible as a backup tunnel on local DP0.

To configure Extension HCL, refer to Configuring Extension Hot Code Load on page 144.

Extension HCL Limitations and Considerations

Following are limitations and considerations for eHCL:

- eHCL is exclusive to the Brocade 7840 Extension switch and the Brocade SX6 Extension blade. It is incompatible with the Brocade FX8-24 Extension blade.
- Although the Brocade 7810 Switch is compatible with eHCL, because it has a single DP, it does not support eHCL itself. However, if the Brocade 7810 Switch is connected to a Brocade 7840 Extension switch or Brocade SX6 Extension blade, it will support remote eHCL for the connected switch. Consequently, you cannot configure the `--local-ha-ip <ip-addr>` option for the `portcfg fciptunnel|fcipcircuit` command on the Brocade 7810 Switch. However, you can still use the `--remote-ha-ip <ip-addr>` option. When you intend to use such a configuration for an attached Brocade 7840 or SX6 Extension product, use only the `--local-ha-ip <ip-addr>` option. (Else, if you use `--remote-ha-ip <ip-addr>` on the Brocade 7840 Extension switch or SX6 Extension blade products, the tunnel will never reach an ONLINE state; at best, it will remain DEGRADED.) In this scenario, only the main tunnel group and remote-backup tunnel groups will appear on the Brocade 7810 Switch, and only the main tunnel group and local-backup tunnel groups will appear on the connected Brocade 7840 Extension Switch or Brocade SX6 Extension blade.

NOTE

The `portshow fciptunnel --hcl-status` command is available for the Brocade 7810 Switch even though it does not support eHCL. You can use the output of this command to view the DP connection information.

- No configuration changes are permitted during the eHCL process. This includes modifying tunnel or circuit parameters. New device connections that require "zone checking" may experience timeouts during the CP reboot phase of the firmware download. The CP performs all zone checking and so must be active to process new SID/DID connection requests like PLOGIs.
- eHCL supports Virtual Fabrics (VF) and FC Routing (FCR with the IR license) and all existing features.
- eHCL was designed for all environments including mainframe FICON XRC and tape and open systems disk replication (EMC SRDF, HDS Universal Replicator, IBM Global Mirror, HP Remote Copy, and others). eHCL supports asynchronous and synchronous environments as well as IPEX.
- The Brocade 7840 Extension switch, and the Brocade SX6 Extension blade have two data processor (DP) complexes: DP0 and DP1. (This discussion does not apply to the Brocade 7810 Extension switch, as it has only one data processor.) During the eHCL process, each DP reloads sequentially, while the other DP remains operational. Consider the following for planning and use of the switch during this process:
  - Because only one DP complex remains operational at a time, the total switch capacity is temporarily diminished by up to 50 percent.
  - FCIP and IPEX data is not lost and remains in-order. eHCL does not cause FICON interface control check (IFCC).
  - Implementation of eHCL requires proper planning. Considerable bandwidth may be available based on licensing and compression to the Fibre Channel (FC) and FICON side of the extension switches. Furthermore, there are typically A and B paths in redundant replication networks, which provides additional bandwidth during firmware updates. Deployment planning may include apportioning one of the two DP complexes to HA or limiting maximum use to 50 percent of the licensed capacity across both DP complexes to reserve adequate bandwidth for high-availability operations.
- Although most firmware updates for Fabric OS 7.4.0 and later will support eHCL, not every Fabric OS release will guarantee firmware capable of using this feature. Refer to the Fabric OS release notes for details.
- The firmware on the switch at each end of the tunnel must be compatible. Start with the same version of Fabric OS and upgrade to the same version on both ends. If not, you might either introduce instability and aberrant behavior or prevent successful tunnel formation when the main tunnel attempts to come back online.
- eHCL does not require any additional communication paths. For the normal operation of data replication and/or tape backup, you will have an existing extension tunnel (MT) connected across the WAN. The two backup tunnels will exist alongside the main tunnel.
- eHCL leverages RASlog warnings and error messages (WARN/ERROR).
- If parallel tunnels (VE_Ports) are configured between local and remote sites, if each tunnel has multiple circuits, you must set the VE link cost to static. Else, FC traffic might be disrupted during eHCL activity.
- When Teradata Emulation is enabled on an Extension tunnel, eHCL is not supported. You must perform a disruptive firmware update.
- Do not configure certain pairs of VE ports in different logical switches with different traffic routing policies. For example, on a Brocade 7840, one logical switch has Exchange-Based Routing (EBR) and the other has Port-Based Routing (PBR), which may be common in mainframe environments. VE_Ports 24 and 34 are used in the different logical switches. During eHCL, these VEs share backend virtual channels (VCs) when the VE is 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 4 VE_Port Pairs and Differing LS Traffic Policies

Brocade 7840 VE_Port		Brocade SX6 VE_Port	
DP0	DP1	DP0	DP1
24	34	16	26
25	35	17	27
26	36	18	28
27	37	19	29
28	38	20	30
29	39	21	31
30	40	22	32
31	41	23	33
32	42	24	34
33	43	25	35

Extension HCL Enhancements in Fabric OS 8.2.0

Brocade Fabric OS 8.2.0 includes enhancements to the Extension HCL feature.

Non-terminated TCP (NTTCP) Traffic

In Fabric OS 8.2.0, NTTCP traffic is queued during the failover/failback cycles of the firmware upgrade instead of being dropped. Some traffic loss can occur if the amount of NTTCP data received during the failover/failback exceeds the available queue size. In addition, a WAN flush is performed on the non-terminated streams that carry NTTCP, UDP, or ICMP traffic. This action reduces the probability of application I/O timeouts during the failover/failback cycle.

The NTTCP traffic is queued and non-terminated WAN streams are flushed only when all participating DPs on a VE_Port are running Fabric OS 8.2.0 or a subsequent release.

Independent Resumption of LAN TCP Connections

Prior to Fabric OS 8.2.0, traffic resumed on an LAN TCP connection after all LAN connections on a VE_Port failed over. In instances of a slow-draining LAN connection, all other LAN connections on that VE_Port were delayed until the slow-draining device completed its failover. In Fabric OS 8.2.0 and subsequent releases, the failover sequence of a LAN TCP connection consists of the following phases:

- WAN flush complete, which indicates the data for the LAN connection is flushed on the WAN in both directions.
- LAN Tx complete, which indicates the data for a LAN connection is flushed on the LAN in both directions.

After receiving the WAN flush complete indication for all LAN connections on a VE_Port, each LAN TCP connection can be independently resumed as soon as its LAN Tx is complete. This action decouples having to wait on a slow-draining LAN connection before all LAN connections can be resumed on a given VE_Port. All participating DPs on a VE_Port must be running Fabric OS 8.2.0 or a subsequent release.

Concurrent eHCL Support on Gen 6 Chassis

In Fabric OS 8.2.0 or a subsequent release, concurrent (or parallel) HCL is supported on a Gen 6 chassis, for example the Brocade X6 platform. Concurrent eHCL means that between two Gen 6 chassis running Fabric OS 8.2.0 and configured for eHCL support, you can initiate firmware upgrades at the local site and remote site at the same time.

ARL Protocol Rollover Management

From Fabric OS 8.2.0 onwards, the management of ARL protocol rollover during eHCL is now 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 protocol rollover is enabled for that VE_Port.

Fibre Channel SAN Considerations

There are many considerations in the implementation and design of a Fibre Channel storage area network (FC SAN). For detailed information, refer to the Brocade white paper, SAN Design and Best Practices.

Adaptive Rate Limiting

Adaptive Rate Limiting (ARL) is performed on circuits to change the rate at which the tunnel transmits data through the IP network. ARL uses information from the TCP connections to determine and adjust the rate limit for the circuit dynamically. This allows connections to utilize the maximum available bandwidth while providing a minimum bandwidth guarantee. ARL is configured on a per-circuit basis because each circuit can have available different amounts of bandwidth. Any single circuit is limited to 10 Gb/s, unless the hardware imposes a lower bandwidth. The following list identifies ports and available bandwidth:

- XGE (10GbE) ports on the Brocade FX8-24 Extension blade (xge0 and xge1)
- 10G (1/10GbE) ports on the Brocade 7840 Switch (ge2-ge17)
- 10G (1/10GbE) ports on the Brocade SX6 Extension blade (ge2-ge17)
- 40GbE ports on the Brocade 7840 Extension switch (ge0 and ge1)
- 1GbE ports on the Brocade 7810 Extension switch (ge0-ge1)
- 10G (1/10G-GbE) ports on the Brocade 7810 Extension switch (ge2-ge7)
- 40GbE ports on the Brocade SX6 Extension blade (ge0 and ge1)
- 1GbE ports on the Brocade FX8-24 Extension blade (ge0-ge9)

Brocade ARL always maintains at least the minimum configured bandwidth, and it never tries to exceed the maximum level. Everything in between is adaptive. Brocade ARL tries to increase the rate limit up to the maximum until it 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 climb toward the maximum. On the other hand, if congestion events are encountered, Brocade 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.

For additional information and details of ARL, refer to the Brocade white paper, Brocade Adaptive Rate Limiting.

Brocade 7840 Switch, Brocade 7810 Switch, and Brocade SX6 Blade Support for ARL

ARL accommodates shared bandwidth; however, the amount of storage data using Extension connections continues to grow and consume larger and faster links. On supported Gen 5 and Gen 6 platforms (Brocade 7840 Switch, Brocade 7810 Switch, and the Brocade SX6 Blade), 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 implemented by ARL is optimized to increase overall throughput. ARL dynamically preserves bandwidth and evaluates network conditions to see whether additional back-offs are required.

ARL maintains round-trip-time (RTT) stateful information to better predict network conditions and to allow intelligent and granular decisions about proper adaptive rate limiting. When ARL encounters a network error, it looks back at prior stateful information, which will be different relative to the current state. Rate-limit decisions are then modified using the ARL algorithm. When configured for automatic selection, ARL will dynamically determine which algorithm to use based on the changing network conditions.

On the Brocade SX6 Blade, the Brocade 7810 Switch, and the Brocade 7840 Switch, you can configure the type of ARL algorithm that is used for backing off the traffic. The default is automatic, and the ARL logic determines the best approach. The ARL choices for these platforms are as follow:

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

Brocade FX8-24 Extension Blade Support for the ARL Backoff Algorithm

ARL on the Brocade FX8-24 Extension Blade uses a single backoff algorithm, which is “reset to floor.” This means that when congestion is detected at the current bandwidth, it resets to the minimum configured bandwidth for the circuit and restarts the climb toward maximum bandwidth.

FSPF Link Cost Calculation When ARL Is Used

Fabric Shortest Path First (FSPF) is a link state path selection protocol that directs traffic along the shortest path between the source and the destination based upon the link cost. When ARL is used, the link cost is calculated based on the sum of the maximum traffic rates of all established, currently active low-metric 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 are used, configure Lossless DLS. This action avoids FC frame loss during routing updates that are made because of bandwidth updates.

NOTE
Multiple parallel tunnels are supported, but their use is not recommended.

ARL Considerations

Consider the following limitations when configuring ARL:

- As a best practice, the aggregate of the maximum-rate bandwidth setting through a tunnel should not exceed the bandwidth of the WAN link. For example, given a 2Gb/s WAN link, the aggregate of the ARL maximum rates connected to that WAN link can be no more than 2Gb/s. For ingress rates, there is no limit because the FC flow control (BBC) rate-limits the incoming data.
- The aggregate of the minimum configured values cannot exceed the speed of the Ethernet interface, which is 1Gb/s for GbE ports or 10Gb/s for 10GbE ports, and 40Gb/s for 40GbE ports.
- Configure minimum rates of all tunnels so that the combined rate does not exceed the specifications listed for the extension product in the Tunnel and Circuit Requirements for Brocade Extension Platforms on page 62.

- For 1GbE, 10GbE, and 40GbE ports, the ratio between the minimum committed rate and the maximum committed rate for a single circuit cannot exceed five times the minimum. For example, if the minimum is set to 1 Gb/s, the maximum for that circuit cannot exceed 5Gb/s. This limit is enforced in software.
- The ratio between any two circuits on the same tunnel should not exceed four times the lower circuit. For example, if one circuit is configured to 1Gb/s, any other circuit in that same tunnel should not exceed 4Gb/s. This limit is *not* enforced in software, but is strongly recommended.
- For any circuit, the minimum bandwidth values must match on both the local side and the remote side, and so must the maximum bandwidth values.
- ARL is invoked only when the minimum and maximum bandwidth values on a circuit differ. In other words, if you configure both the minimum and maximum bandwidth values on a circuit, say, to 1,000,000 (1 Gb/s), ARL is not used.

Compression Options

Brocade Fabric OS software provides an advanced compression architecture that supports multiple modes to optimize compression ratios for various throughput requirements. The available modes include hardware-based compression and software-based compression. The compression mode available depends on the Brocade extension platform and the protocol (FCIP or IPEX).

NOTE
Throughput for all compression modes depends on the compression ratio achievable for the data pattern. Brocade makes no promises, guarantees, or assumptions about the compression ratio that any application may achieve.

Compression Options for the Brocade 7840 Extension Switch, the Brocade 7810 Extension Switch, and the Brocade SX6 Extension Blade

The following compression options are available for the Brocade 7840 Extension switch and the Brocade SX6 Extension blade.

- None: No compression.
- Fast deflate: Hardware-based compression. This mode initiates a deflate-based algorithm to compress data before it enters the DP and to decompress the data after it leaves the DP. It provides the highest throughput at 40Gb/s per DP before compression, but it provides the least amount of compression. Because fast deflate is hardware-based, this mode is appropriate for synchronous applications. The Brocade 7810 Extension switch does not support fast deflate.
- Deflate: Processor-based compression. This mode initiates the processor compression engine in deflate mode with a preference for 16Gb/s total speed per DP before compression. It provides a lower speed than fast deflate, but a faster speed than aggressive deflate. Deflate compression provides more compression than fast deflate, but is typically not as much compression as aggressive deflate.
- Aggressive deflate: Processor-based compression. This mode initiates the processor engine in deflate mode with a preference for compression. This mode is the slowest at 10Gb/s before compression, but it typically provides the highest level of compression.

Follow the guidelines for assigning explicit compression levels for tunnels in the following table.

TABLE 5 Assigning Compression Levels

Total tunnel bandwidth on a DP	Compression level
More than 4Gb/s	Fast deflate
2Gb/s to 4Gb/s	Deflate
Less than 2Gb/s	Aggressive deflate

The enhancements for IP Extension allow you to configure compression on the tunnel at a protocol level. The compression options override the main tunnel compression level and set the compression for the specified protocol to the desired mode. The available modes depend on the protocol, FC or IP.

TABLE 6 Extension Hybrid Mode Protocol Compression Choices

Compression level	FC protocol support	IP protocol support
Deflate	Yes	Yes
Fast deflate	Yes	No
Aggressive deflate	Yes	Yes

Compression Options for the Brocade FX8-24 Extension Blade

The following compression options are available for the Brocade FX8-24 Extension Blade:

- None: No compression.
- Standard: Hardware compression mode.
- Moderate: A combination of hardware and software compression that provides more compression than hardware compression alone. This option supports up to 8 Gb/s of FC traffic.
- Aggressive: Processor-based compression that provides a more aggressive algorithm than that used for the standard and moderate options. This option supports up to 2.5 Gb/s of FC traffic.
- Auto: Allows the system to set the best compression mode based on the tunnel's configured bandwidth and the aggregate bandwidth of all tunnels in the extension blade.

Follow the guidelines for assigning explicit compression levels for tunnels, as shown in the following table.

TABLE 7 Assigning Compression Levels

Total Tunnel Bandwidth on a DP	Compression Level
More than 2 Gb/s	Standard
More than 512 Mb/s and less than or equal to 2 Gb/s	Moderate
Equal to or less than 512 Mb/s	Aggressive

FastWrite and Open Systems Tape Pipelining

Brocade FastWrite is an algorithm that reduces the number of round trips required to complete a SCSI write operation. FastWrite can maintain throughput levels over links that have significant latency. The RDR (Remote Data Replication) application still experiences latency; however, reduced throughput due to that latency is minimized for asynchronous applications, and response time is reduced by up to 50 percent for synchronous applications.

Open Systems Tape Pipelining (OSTP) enhances open systems SCSI tape read and write I/O performance. When the extension link is the part of the network with the longest latency, OSTP can provide accelerated speeds for tape read and write I/O over tunnels. To use OSTP, you must also enable FastWrite.

OSTP accelerates SCSI read and write I/O to sequential devices (such as tape drives) over extension links, which reduces the number of round-trip times needed to complete the I/O over the IP network and speeds up the process.

For FastWrite and OSTP to work, both ends of a tunnel must have matching FastWrite and OSTP configurations. You enable FastWrite and OSTP during the tunnel configuration process. They are enabled on a per-tunnel basis.

FastWrite and OSTP Network Configurations

The Fibre Channel Protocol (FCP) features used in FastWrite and OSTP require a single deterministic path between initiators and targets when multiple tunnels exist. If there are Equal-Cost Multi-Path (ECMP) tunnels between the same SID/DID pairs, protocol optimization will fail when a command is routed over one tunnel and the response is returned over a different tunnel.

The following figures show network configurations supported by FastWrite and OSTP. Neither configuration contains multiple ECMP paths. The first figure shows a single tunnel with FastWrite and OSTP enabled. The second figure shows multiple tunnels, but none of them creates a multiple ECMP path.

NOTE

Only one emulating tunnel is supported between an initiator port and a peer device port.

Brocade extension devices can distinguish between storage flows that use protocol optimization and those that do not use protocol optimization. For example, SAN Volume Controller (SVC) from IBM Corporation does not use FastWrite, but Asynchronous Symmetrix Remote Data Facility (SRDF/A)from EMC Corporation does use FastWrite. Both applications functioning over the connection are fully supported for FastWrite because FastWrite will not engage with the IBM SVC flows yet will engage with the SRDF/A flows across the same VE_Port. This is also true when using OSTP with IBM SVC. Both flows can utilize the same VE_Port with FastWrite and OSTP enabled.

FIGURE 3 Single Tunnel with FastWrite and OSTP Enabled

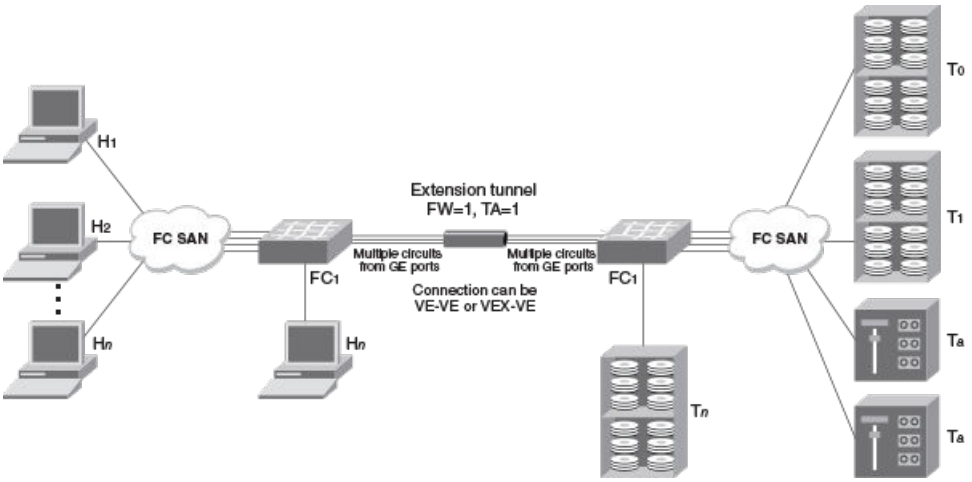

FIGURE 4 Multiple Tunnels to Multiple Ports with FastWrite and OSTP Enabled on a Per-Tunnel, Per-Port Basis

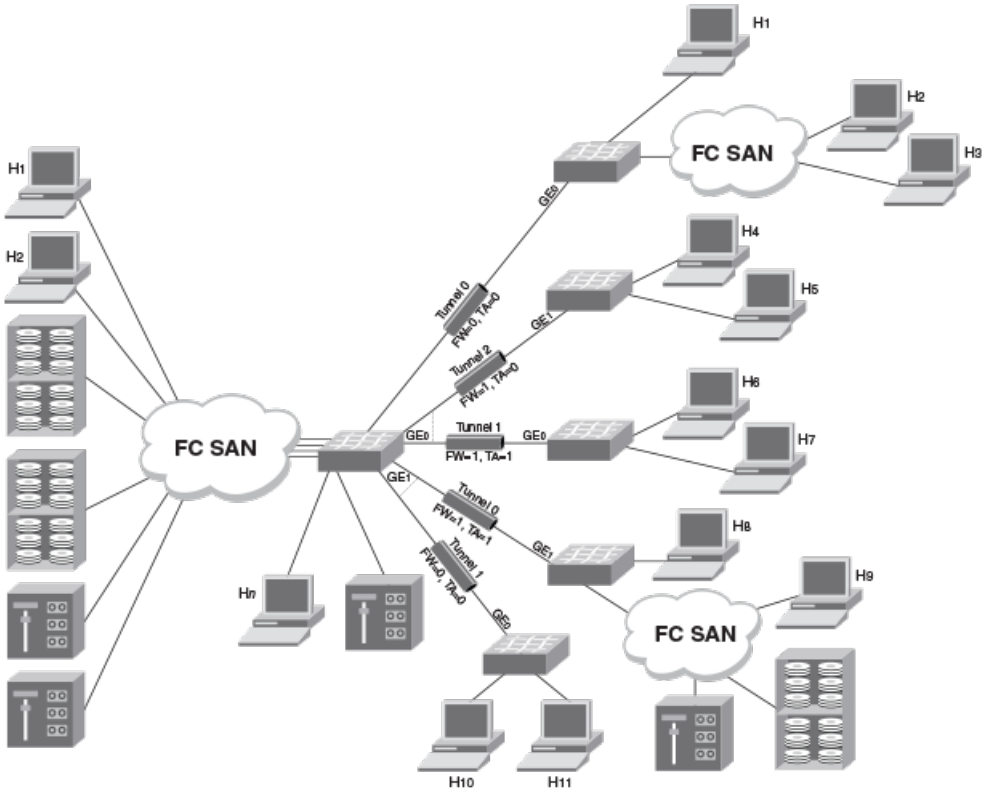

In some cases, VF LS/LF or Traffic Isolation (TI) Zoning configurations can be used to control the routing of SID/DID pairs to individual tunnels.

This provides deterministic flows between the switches and allows the use of ECMP. Refer to the *Brocade Fabric OS Administration Guide* for more information about Traffic Isolation Zoning.

FICON Acceleration

FICON Acceleration provides specialized data management techniques and automated intelligence to accelerate FICON tape read and write and IBM Global Mirror data replication operations over distance, while maintaining the integrity of command and acknowledgment sequences.

To use the features of FICON Acceleration, you must obtain the Advanced FICON Acceleration license, which allows interoperability for the following features:

- Write and read tape pipelining
- IBM z/OS Global Mirror emulation
- Teradata emulation

FICON Acceleration licenses are available on the Brocade 7840 switch and the Brocade DCX 8510 family for the FX8-24 on an individual slot basis.

NOTE

This license is not required on Brocade Gen 6 platforms with a Brocade SX6 blade.

NOTE
The Brocade 7810 Extension Switch supports neither FICON nor mainframe applications.

As part of FOS 8.2.1, support for Brocade 7840 and SX6 configurations, FICON Emulation will support Exchange Based Routing policy.

This support comprises two major changes.

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

This change should improve FICON over FCIP performance improvements as FICON flows will not all be mapped to the same VT or VC (therefore head of line blocking will be avoided in both the WAN path and the FC egress path processing).

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 provider) indicate support for Exchange Based Routing.

For additional information about FICON configuring and administration, refer to the *Brocade Fabric OS FICON Configuration Guide*. In addition, you can refer to the eBook, *Brocade Mainframe Connectivity Solutions*.

VM Insight

VM Insight is a feature implemented in Fabric OS 8.1.0 that allows flows between virtual machines (VMs) to be tracked and identified with 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. If the server added the optional VMID header, the header will be used between the FCIP complex and the target device and will be included in FCP sequences back to the initiator.

Support of this feature requires that the drivers for HBA and storage vendors support the VMID frame-tagging method that identifies each subflow through the fabric.

The following table shows VM Insight support.

TABLE 8 VM Insight Supported Platforms

	FX8-24 (all Fabric OS)	7840/SX6 (before Fabric OS 8.1.0)	7840/SX6 (Fabric OS 8.1.0 or later)	7810 (Fabric OS 8.2.1)
Non-emulated tunnels	Supported	Supported	Supported	Supported
FCP emulated tunnels (FastWrite, OSTP)	Not supported	Not supported	Supported	Supported
FICON-emulated tunnels	Not supported	Not supported	Not supported	Not Supported

For more information, refer to the *Brocade Flow Vision Configuration Guide*, the *Brocade Monitoring and Alerting Policy Suite Configuration Guide*, and the *Brocade Network Advisor SAN + IP User Manual*.

NVMe Support over Extension

NVMe is supported over FCIP tunnels on the Brocade 7810 switch, the Brocade 7840 switch, the Brocade SX6 blade and the Brocade FX8-24 blade.

IP Security Encryption

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

The following sequence of events invokes the IPsec protocol.

- IPsec and Internet Key Exchange (IKE) policies are created and assigned on peer switches or blades on both ends of the tunnel.

  NOTE
  On the Brocade FX8-24 Blade, you enable IPsec and provide a preshared key inline on the tunnel. There is no specific policy creation.

- 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 association (SA) parameters include encryption and authentication algorithms, Diffie-Hellman key exchange, and SAs.
- Data is transferred between IPsec peers based on the IPsec parameters and keys stored in the SA database.
- When authentication and IKE negotiation are complete, the IPsec SA is ready for data traffic.
- SA lifetimes terminate through deletion or by timing out. An SA lifetime equates to approximately two billion frames of traffic passed through the SA or after a set time interval has passed.
- When the SA is about to expire, an IKE rekey will occur to create a new set of SAs on both sides and data will start using the new SA. IKE and SA re-keys are non disruptive.

On the Brocade 7840 Switch, the Brocade 7840 Switch, and the Brocade SX6 Blade, IPsec support in Fabric OS 8.1.0 and subsequent releases support Suite B Cryptographic Suites for IPsec. This enables configuration of the Elliptic Curve Diffie-Hellman (ECDH) for key agreement and Elliptic Curve Digital Signature Algorithm (ECDSA) for peer authentication. Suite B configuration requires the generation and importation of Suite B-compliant X.509 end-entity certificates as well as 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 prior Fabric OS releases.

TABLE 9 IPsec Suite B Comparison on the Brocade 7810, the Brocade 7840, and the Brocade SX6

Attribute	Prior to Fabric OS 8.1.0 releases	Suite B in Fabric OS 8.1.0 and later
Authentication	Shared key	ECDSA-P384
Diffie-Hellman	MODP-2048	ECDH-P384
PRF	PRF-HMAC-512	PRF-HMAC-384
Integrity	None	HMAC-384-192
Encryption	AES-256-GCM	AES-256-CBC (IKE connection) AES-256-GCM (data connection)

NOTE
Brocade 7810 Switch was introduced with Fabric OS 8.2.1, so please interpret the table accordingly.

As part of Suite B support, you can use digital certificates and third-party certificate authority (CA) to verify the identity of the certificates. This requires each end of the secure connection to have access or a means to lookup 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 preshared key profile allows you to specify the key when IPsec is configured. The public key infrastructure (PKI) profile uses CA certificates.

IPsec for the Extension Switches and Blades

Advanced Encryption Standard, Galois/Counter Mode, Encapsulating Security Payload (AES-GCM-ESP) 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 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.
- The IKEv2 key exchange protocol is used by peer switches and blades for key agreement.
- IKEv2 uses UDP port 500 to communicate between the peer switches or blades.
- All IKEv2 traffic is protected using AES-256 encryption.
- An 128-bit GCM or 192-bit hash message authentication code (HMAC) is used to check data integrity and detect third-party tampering.
- Pseudo-random function (PRF) is used to generate multiple security keys, using 384-bit or 512-bit HMAC.
- A 2048-bit Modular Exponential (MODP) or 384-bit Elliptic Curve Diffie-Hellman group is used for both IKEv2 and IPsec key generation.
- No encryption key is used for more than 4 hours or 2 billion frames, whichever comes first. When the SA lifetime expires, a new key is generated, limiting the amount of time 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 lifetime expires.

  After the key lifetime of 4 hours or 2 billion frames has been reached, new keys are generated and the old keys are discarded, This process of rekeying is non disruptive.
- Encapsulating Security Payload (ESP) is used as the transport mode. ESP uses a hash algorithm to calculate and verify an authentication value, and it encrypts only the TCP header and TCP payload.
- A circuit in a nonsecure tunnel can use the same Ethernet interface as a circuit in 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.

Limitations Using IPsec over Tunnels

The following limitations apply to using IPsec:

- Network Address Translation (NAT) is not supported.
- Authentication Header (AH) is not supported.
- IPsec-specific statistics are not supported on the Brocade FX8-24 Blade; however, the Brocade 7810 Switch, the Brocade 7840 Switch, and the Brocade SX6 Blade do provide statistics.
- On the FX8-24 blade, IPsec is configurable only on IPv4-based tunnels; however, the Brocade 7810 Switch, the Brocade 7840 Switch and the Brocade SX6 Blade support IPsec on IPv6 as well as IPv4-based tunnels.

- Older versions of the Brocade FX8-24 Blade do not support IPsec on VE_Ports 22 through 31. For these blades, a RASlog warning message displays saying that the blade is not at the correct version to support IPsec-enabled tunnels on VE_Ports 22 through 31.
- On the Brocade FX8-24 Blade, a tunnel using IPsec in Fabric OS 7.0.0 and later must use the **--legacy** tunnel option if connecting to a peer running a version beforeFabric OS 7.0.0.
- IPsec is not allowed with the **--connection-type** tunnel option set to anything other than the default.

IPv6 Addressing

This implementation of IPv6 uses unicast addresses for the interfaces with circuits. The link-local unicast address is automatically configured on the interface, but using the link-local address space for circuit endpoints is not allowed. Site-local unicast addresses are not allowed as circuit endpoints.

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

Note the following IPv6 addressing points:

- Anycast addresses are not used. Each IPv6 interface has a unique unicast address, and addresses configured are assumed to be unicast.
- Multicast addresses cannot be configured for an IPv6 interface with circuits.
- The IPv6 8-bit Traffic Class field is defined by the configured Differentiated Services field for IPv6 (RFC 2474). This configuration is done on the circuit using the Differentiated Services Code Point (DSCP) parameters to fill the 6-bit DSCP field.
- The IPv6 optional Extension Headers are not supported.
- Parts of the Neighbor Discovery protocol (RFC 4861) are used in this implementation.
  - Hop limits (such as Time to Live (TTL)) are learned from the Neighbor Advertisement packet.
  - The link-local addresses of neighbors are learned from Neighbor Advertisement.
  - The IPv6 link-local address for each GE interface is configured at startup and advertised to neighbors. The user does not configure the interface link-local address.
- IPv6 addresses and routes must be statically configured by the user. Router advertisements and IPv6 stateless address autoconfiguration (RFC 2462) are not supported.
- ICMPv6 message types in RFC 4443 and ICMPv6 message types used for Neighbor Discovery (ND) are supported.
- Path MTU discovery
  - For the Brocade 7810 Switch, the Brocade 7840 Switch and the Brocade SX6 Blade, PMTU discovery is supported. See Path MTU Discovery on page 77.
  - For the Brocade FX8-24 Blade, path MTU (PMTU) discovery is not supported. The MTU option in the `portcfg ipif` command is optional. If not configured, an MTU of 1500 bytes is used. The maximum IP MTU supported is 1500 bytes (including the 40-byte fixed IPv6 header), the same as for IPv4. The minimum IP MTU allowed is 1280 bytes, including the 40-byte fixed IPv6 header. Any network used for IPv6 circuits must support an IP MTU of 1280 bytes or larger. IPv6 fragmentation is not supported. The Layer 4 protocol ensures that the PDU is less than the IP MTU (including headers).
- IPv6 addressing with IPsec
  - For the Brocade 7810 Switch, the Brocade 7840 Switch and the Brocade SX6 Blade, IPv6 addressing can be used when implementing IPsec.
  - For the Brocade FX8-24 Blade, IPv6 addressing cannot be used when implementing IPsec.

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

Memory Use Limitations for Large-Device Tunnel Configurations

The data processing layer on the Brocade Extension switch and blade data processing (DP) complex has access to reserved memory used for control block structure allocations. Following are related specifications for the Brocade Extension switches and blades.

TABLE 10 VE_Ports and DRAM Pool Sizes for Brocade Extension Products

Product	DP VE_Ports	DP DRAM Pool Size
Brocade 7840 Extension Switch	DP0: 24 through 33	2.4 GB
	DP1: 34 through 43	2.4 GB
Brocade 7810 Extension Switch	DP0: 12-15	256 MB
Brocade SX6 Extension Blade	DP0:16 through 25	2.4 GB
	DP1: 26 through 35	2.4 GB
Brocade FX8-24 Extension Blade	DP0: 22 through 31	536 MB
	DP1: 12 through 21	536 MB

Tunnel processing will create more control blocks when any type of emulation feature is enabled, such as FCP/SCSI Fast Write, Open Systems Tape Pipelining, or FICON. In those cases, be sure to not include too many devices running over the tunnel. If too many devices are present or activated at one time, emulation operations can be negatively impacted. Even without emulation enabled, too many devices running over the tunnel may impact operations at some point because of memory consumption.

NOTE

A configuration that works without an emulation feature, such as FICON Acceleration, FastWrite, or Open Systems Tape Pipelining (OSTP), may not work when emulation features are 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 7840 Switch, VE_Port 28. (The Brocade 7840 Switch has no blade slots.)

```
switch:admin> portshow xtun 28 -dram2
Dram2 Pool Info:

Dram2 Cacheable: 0x8000000500d57300 - 0x80000005a0d572ff (2684354560)
pMaster=0x012517be00 heap_free=2676601216 heap_alloc=7753344 defrags=0 failedAllocs=0

Fast Free Usage=3070 Fast Free Allocates=15763317
Merged & Defrag'd Queues:
               User       MAX
Q SegSize     Alloc     UserA     Alloc'd  Returnd       Max   Current  Bytes
== =======    =====     =====     =======     =======       ===   =======  =====
1 0x0080      3213      3213        3213           0       275       275  0x00008980
2 0x0100       603       603         603           0        40         6  0x00000600
3 0x0180       306       306         306           0        43        43  0x00004080
4 0x0200       105       105         105           0        99        65  0x00008200
5 0x0280      1681      1681        1681           0         4         3  0x00000780
6 0x0300       396       396         396           0         1         0  0x00000000
7 0x0380        12        12          12           0         5         0  0x00000000
8 0x0400       267       267         267           0        33         0  0x00000000
9 0x0480       335       335         335           0        21         0  0x00000000
11 0x0580       50        50          50           0        18         0  0x00000000
12 0x0600      419       419         419           0         1         0  0x00000000
13 0x0680       37        37          37           0         9         0  0x00000000
15 0x0780      294       294         294           0         1         0  0x00000000
18 0x0900        2         2           2           0         1         0  0x00000000
19 0x0980        1         1           1           0         2         0  0x00000000
22 0x0b00       16        16          16           0         3         0  0x00000000
24 0x0c00       50        50          50           0        35         0  0x00000000
25 0x0c80        5         5           5           0         1         0  0x00000000
28 0x0e00        3         3           3           0         1         0  0x00000000
29 0x0e80       16        16          16           0         1         0  0x00000000
30 0x0f00      534       534         534           0         1         0  0x00000000
32 0x1000       18        18          18           0         2         0  0x00000000
33 0x1080        7         7           7           0         1         0  0x00000000
34 0x1100        1         1           1           0         5         0  0x00000000
40 0x1400       66        66          66           0         1         0  0x00000000
42 0x1500        3         3           3           0         1         0  0x00000000
43 0x1580        2         2           2           0         1         0  0x00000000
44 0x1600       34        34          34           0         2         0  0x00000000
48 0x1800       48        48          48           0         2         1  0x00001800
51 0x1980        9         9           9           0         1         0  0x00000000
53 0x1a80        5         5           5           0         1         0  0x00000000
56 0x1c00        4         4           4           0         4         0  0x00000000
59 0x1d80        2         2           2           0         2         0  0x00000000
64 0x2000       23        42          45          22    327680    326722  0x9f884000

Fast Free Queues:
  Q  SegSz/dec.   Allocs      Frees     Max    Current   Bytes
=== ======/==== ========== ========== ======== ======== ==========
  1 0x0080/128      19681      21610     3058     1929 0x0003c480
  2 0x0100/256     132687     132692      155        5 0x00000500
  3 0x0180/384        124        124       62        0 0x00000000
  4 0x0200/512        351        361       62       10 0x00001400
  5 0x0280/640        131        466      335      335 0x00034580
  6 0x0300/768        569        571       27        2 0x00000600
  8 0x0400/1024    9785262    9785460      263      198 0x00031800
  9 0x0480/1152          0        335      335      335 0x0005e380
11 0x0580/1408    1265810    1265860       50       50 0x00011300
13 0x0680/1664          0         15       15       15 0x00006180
15 0x0780/1920        131        185       54       54 0x00019500
24 0x0c00/3072    1265810    1265860       50       50 0x00025800
25 0x0c80/3200     126581     126586        5        5 0x00003e80
30 0x0f00/3840    3039599    3039653       54       54 0x00032a00
33 0x1080/4224     126581     126586        5        5 0x00005280
```

```
48 0x1800/6144          0         14       14       14 0x00015000
51 0x1980/6528          0          9        9        9 0x0000e580
                ---------- ---------- -------- -------- ----------
                  15763317                          3070 0x001B7680
Total Bytes in DRAM2 Pool: 2678401024 (free)    1799808 (fastfreed)
Total DRAM Bytes Allocated: 7753344 (in use)

switch:admin>
```

Control Blocks Created during FCP Traffic Flows

For FCP/SCSI traffic flows, tunnel processing creates control block structures based upon the SID/DID pairs used between initiators and devices. When either FastWrite or OSTP (read or write) is enabled, additional structures and control blocks are created for each logical unit number (LUN) on a SID/DID-pair basis. FCP processing in an emulated tunnel configuration will create multiple control blocks for each LUN if there are multiple SID/DID pairs that can be used to access that same LUN. The following structures and control blocks are created:

- Initiator, target, nexus (ITN) structure: Each FCP-identified SID/DID flow will be recorded in an ITN structure.
- Initiator, target, LUN (ITL) control block: Each specific LUN on a SID/DID flow will have an ITL control block created for the flow.
- Turbo write block (TWB) structure: FCP emulation processing also 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 upon the SID/DID pairs called a FICON device path block (FDPB). If any FICON emulation feature is enabled, additional control blocks are created for each SID/DID pair, logical partition (LPAR) number (FICON channel block structure), LCU number (FICON control unit block structure), and for each individual FICON device address on those LCUs (FICON device control block structure). FICON Exchange control blocks are also created if the switch is operating in EBR mode at FOS 8.2.1 or higher.

> NOTE
> The Brocade 7810 Extension switch supports neither FICON Emulation nor FICON SANs.

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

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

This number grows quickly in extended direct-attached storage device (DASD) configurations, such as those used in IBM z/OS Global Mirror, also known as Extended Remote Copy (XRC).

FDCBs Example

The following example assumes that the tunnel is used to extend two channel paths (CHPIDs) from a System Data Mover (SDM) site to a production site. It also assumes that there are two SDM-extended LPARs, that the IBM DS8000 production controllers have 32 LCUs per chassis, and that each LCU has 256 device addresses.

Using the preceding equation, the number of extended FICON device control block images created would be the following:

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

Considerations for Tunnel Control Block Memory and Device Configuration

The `portshow xtun` *`slot/ve_port`* `-fcp -port -stats` command displays current usage and control block sizes per tunnel 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 created for FCP and FICON traffic flows during tunnel processing.

The following example shows 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 FCP/SCSI VE_Port output.

```
switch:admin> portshow xtun 24 -fcp -port -stats
Slot(0.dp0) VePort(24) Port Stats:
  Global Queue Stats:
   Name,cnt,max,usage,size,total size
   Data,10,12,4794,8192,81920
   Message,1,2,2374,7448,7448
   Stat,0,0,0,1152,0
   Stat Cache,0,0,0,0,0
   Global stats,0,0,0,0,113200
  Port Queue Stats:
   Name,cnt,max,usage,size,total size
   Image,4,4,4,0,0
   SRB,2,4,7781217,0,0
   TWB,3,102,26746084,0,0
  Port Struct Allocation Stats:
   Name,cnt,max,usage,size
   IMAGE,4,4,4,6104
   ITN,4,0,0,5608
   OXTBL,11,0,0,8192
   ITL,32,0,0,5048
   TWB,4,102,26746086,1008
  Global Buffer Stats:
   Name,current,min,max
   Write Data Storage,18544,0,3063052
```

```
Read Data Storage,0,0,0
XBAR % avail,100,100,100
WIRE % avail,99,99,100
SWCOMP % avail,100,100,100
```

Use output from `portshow xtun slot/ve_port -fcp -port -stats` command in conjunction 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 rule of thumb, no more than 80 percent of the tunnel DP complex 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 percent of the pool is allocated, consider redesigning the tunnel configuration to ensure continuous operation. When you redesign the tunnel configuration, examine the existing number of SID/DID pairs in the configuration and determine whether new switches, chassis, or blades are needed to reduce the impact on DRAM2.

For Fabric OS releases before 7.4, RASlog message XTUN-1008 provides notification of DRAM2 memory usage. The message is generated by the DP complex when significant memory thresholds are reached. The following thresholds are shown for the Brocade FX8-24 blade:

- 66%
- 33%
- 17%
- 8%
- 0.07%

For Fabric OS 8.0.1 and later, each Brocade FX8-24 Blade, Brocade SX6 Blade DP complex, and Brocade 7840 Switch DP complex generates the XTUN-1008 RASlog message when the following percentages of the DRAM memory pool are available:

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

RASlog messages include the amount of allocated memory from the pool, the amount of free memory in the pool, and the total pool size. Refer to RASlog messages to determine if you need to reduce the size of the extended configuration or to plan for additional switch resources.

Brocade switches and blade DPs are expected to support no more than the number of FICON device control blocks (FDCBs) and extended LUNs (ITLs) noted in the following table.

TABLE 11 FDCBs and ITLs Per Product DP

Product	FDCBs	ITLs
Brocade 7840 Extension Switch	512,000	200,000
Brocade 7810 Extension Switch	0	30,000
Brocade SX6 Extension Blade	512,000	200,000
Brocade FX8-24 Extension Blade	160,000	65,000

The following applies to the Brocade 7840 Switch and the Brocade SX6 Blade.

During Extension Hot Code Load (eHCL) operations, duplicated emulation and non-emulation control blocks are created on the same DP for the high-availability portion of the tunnel. That means that at one point in time during the eHCL process, twice the normal memory requirements are consumed. This duplication process occurs on the remote non-eHCL DP when the primary local DP is undergoing feature disable processing.

The amount of DRAM2 memory on the Brocade 7840 Switch and Brocade SX6 Blade should be able to support eHCL operations with approximately 512K FICON devices active through the VE_Ports on that DP.

Because each customer configuration is unique, the supported number and types of devices will be different. In large configurations, the administrator should review memory usage periodically to ensure continued, reliable operations of the tunnel and emulation features.

Firmware Downloads

For the Brocade FX8-24 Blade, if Fibre Channel traffic or FCIP traffic is active on Fibre Channel ports, the traffic will be disrupted during a firmware download. Similarly, firmware download on a Brocade 7810 Switch will be disruptive during upgrade or downgrade (FCIP and IPEX datapath traffic is disrupted), but native FC-FC traffic will not be disrupted during firmware download.

The Brocade 7840 Switch and the Brocade SX6 Blade support the Extension Hot Code Load (eHCL) feature. During an eHCL action, traffic is failed over to one DP complex while the firmware upgrades in the other DP complex. With eHCL, active FC traffic on Fibre Channel ports and VE_Ports is not disrupted during a firmware download. For more information on this process, see Extension Hot Code Load on page 22.

You must configure eHCL if you want to perform non disruptive firmware downloads. Otherwise, all traffic is disrupted during the download.

> ATTENTION
> When Teradata Emulation is enabled on an extension tunnel, eHCL is not supported. Any Teradata connections are disrupted by a firmware download even when eHCL is configured.

The best practice is to update the switch or blade at both ends of the tunnel with the same maintenance release of Fabric OS software.

For details on downloading firmware, refer to the chapter on installing and maintaining firmware in the *Brocade Fabric OS Administration Guide*.

Extension Platforms and Features

Extension Platforms and Features Overview

The following platforms support Brocade Extension features in Fabric OS software:

- Brocade 7810 Extension Switch
- Brocade 7840 Extension Switch
- Brocade SX6 Extension Blade in the following chassis:
  - Brocade X6-4 Director
  - Brocade X6-8 Director
- Brocade FX8-24 Extension Blade in the following chassis:
  - Brocade DCX 8510-4
  - Brocade DCX 8510-8

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

Tunnel compatibility across extension platforms is outlined below.

TABLE 12 Compatibility Across Platforms

Platforms	Brocade 7840	Brocade 7810	Brocade SX6	Brocade FX8-24
7840	Yes	Yes	Yes	No
7810	Yes	Yes	Yes	No
SX6	Yes	Yes	Yes	No
FX8-24	No	No	No	Yes

NOTE

Extension connections are not supported between the Brocade 7810 switches, Brocade 7840 switches, and Brocade SX6 blades and previous generation products like Brocade 7800/7500 switches and Brocade FX8-24/FR4-18i blades.

The following table provides details about extension platform capabilities.

TABLE 13 Extension Capabilities by Platform

Capability	Brocade 7840	Brocade 7810	Brocade SX6	Brocade FX8-24
Extension Trunking	Yes	Yes	Yes	Yes
Adaptive Rate Limiting	Yes	Yes	Yes	Yes
1Gb/E and 10Gb/E ports	Yes (1/10Gb/s)	Yes	Yes (1/10Gb/s)	Yes (1Gb/s and 10Gb/s optional)
40GbE ports	Yes Enabled using the 7840 WAN Rate Upgrade 2 license.	No	Yes No additional license is required.	No
FC ports	Yes (2, 4, 8, 16Gb/s)	Yes (4, 8, 16, 32Gb/s)	Yes (4, 8, 16, 32Gb/s)	Yes (1, 2, 4, 8Gb/s)
Compression	Yes • Deflate • Aggressive deflate • Fast deflate	Yes • Deflate • Aggressive deflate	Yes • Deflate • Aggressive deflate • Fast deflate	Yes LZ and deflate
Protocol acceleration • FastWrite • Open Systems Tape Pipelining – OSTP read – OSTP write	Yes	Yes	Yes	Yes
QoS • Marking DSCP • Marking 802.1P - VLAN tagging	Yes	Yes	Yes	Yes
FICON extension • IBM z/OS Global Mirror acceleration • Tape read acceleration • Tape write acceleration • Teradata emulation	Yes	No	Yes	Yes
IPsec • AES-256-GCM • SHA-512 HMAC • IKEv2	Yes Transport mode encrypted data transfer (ESP) method	Yes Transport mode encrypted data transfer (ESP) method	Yes Transport mode encrypted data transfer (ESP) method	Yes Transport mode encrypted data transfer (ESP) method
VEX_Ports	No	No	No	Yes
Support for third-party WAN optimization hardware	No	No	No	Yes Support is limited to Silver Peak for Fabric OS 7.1.0b and to Riverbed for Fabric OS 6.4.x.

TABLE 13 Extension Capabilities by Platform (continued)

Capability	Brocade 7840	Brocade 7810	Brocade SX6	Brocade FX8-24
IPv6 addresses for extension tunnels	Yes	Yes	Yes	Yes[1]
Support for jumbo frames	Yes IP MTU of 9216 is the maximum.	Yes IP MTU of 9216 is the maximum.	Yes IP MTU of 9216 is the maximum.	No IP MTU of 1500 is the maximum.
Path maximum transmission unit (PMTU) discovery	Yes Maximum discoverable size is 9100 bytes.	Yes Maximum discoverable size is 9100 bytes.	Yes Maximum discoverable size is 9100 bytes.	No
Extension Hot Code Load (eHCL)	Yes	No	Yes	No
WAN Tool service-level agreement (SLA) support	Yes	Yes	Yes	No
Link Level Discovery Protocol (LLDP)	Yes All Ethernet interface	Yes All Ethernet interface	Yes All Ethernet interface	No
Support Modes	Hybrid mode and FCIP mode	Hybrid	Hybrid mode and FCIP mode	FCIP mode The Brocade FX8-24 has 10G, 1G, and Dual modes.
Keep-alive packets (KAP) for LACP and LLDP	Yes	Yes	Yes	No

The following note apply to the preceding table:

1. IPv6 addressing is not supported with IPsec.

The following table shows IP Extension capabilities on supported platforms.

TABLE 14 IP Extension Capabilities by Platform

Capability	Brocade 7840	Brocade 7810	Brocade SX6
Hybrid mode (FCIP and IP Extension)	Yes	Yes	Yes
Link aggregation group (static and dynamic LAG)	Yes Only in Hybrid mode 1GbE and 10GbE ports	Yes Only in Hybrid mode 1GbE and 10GbE ports	Yes Only in Hybrid mode 1GbE and 10GbE ports
Switch virtual interface (SVI) IP interface (IPIF)	IP traffic through a Brocade extension tunnel	IP traffic through a Brocade extension tunnel	IP traffic through a Brocade extension tunnel
IPEX compression	Deflate, aggressive deflate, and fast deflate	Deflate and aggressive deflate	Deflate, aggressive deflate, and fast deflate
Traffic control list (TCL)	Yes	Yes	Yes
LAN-side jumbo frames	Yes	Yes	Yes
Policy-based routing (PBR)	Yes	Yes	Yes

Brocade 7840 Extension Switch, Brocade 7810 Extension Switch, and Brocade SX6 Extension Blade Overview

The Brocade 7840 Extension Switch, Brocade 7810 Extension Switch, and the Brocade SX6 Extension Blade share a number of design features. However, because the internal hardware architecture and design of the Brocade 7810 Switch is analogous to a "single DP" Brocade 7840 Switch, there are some key differences.

Brocade 7840 Switch and Brocade SX6 Blade

- Two data processor (DP) complexes, DP0 and DP1
- Up to 20 VE_Ports
- Two 40GbE ports
- Sixteen 1/10GbE ports

The following table shows the Ethernet interface properties on the Brocade 7840 Switch and Brocade SX6 Blade.

TABLE 15 Brocade 7840 Switch and Brocade SX6 Blade

Ethernet Interface Properties	Data Processor 0	Data Processor 1
10GbE Interfaces	Sixteen 1/10GbE interfaces shared by both DPs	Sixteen 1/10GbE interfaces shared by both DPs
40GbE Interfaces (WAN side only)	Two 40GbE interfaces shared by both DPs	Two 40GbE interfaces shared by both DPs
Maximum Bandwidth per VE Port	20Gb/s	20Gb/s
Maximum WAN Bandwidth per DP	20Gb/s	20Gb/s
Maximum IP Extension Bandwidth per DP	20/20Gb/s **NOTE** Requires 2:1 compression using fast deflate LAN/WAN	20/20Gb/s **NOTE** Requires 2:1 compression using fast deflate LAN/WAN
Maximum FCIP Bandwidth per DP	40Gb/s FC in (to the DP); 20Gb/s IP out (to the DP)	40Gb/s FC in (to the DP); 20Gb/s IP out (to the DP)
Maximum Number of VE_Ports	No VEX_Ports Default 5 per DP at 20Gb/s VE_Port maximum bandwidth 10 per DP at 10Gb/s VE_Port maximum bandwidth	No VEX_Ports Default 5 per DP at 20Gb/s VE_Port maximum bandwidth 10 per DP at 10Gb/s VE_Port maximum bandwidth
Maximum FCIP Bandwidth in FCIP Mode per DP without Fast Deflate Compression	20Gb/s	20Gb/s
Maximum FCIP Bandwidth in FCIP Mode per DP with Fast Deflate Compression and 2:1 Compression	40Gb/s	40Gb/s
Maximum FCIP Bandwidth in Hybrid Mode per DP without Fast Deflate Compression	10Gb/s	10Gb/s
Maximum FCIP Bandwidth in Hybrid Mode per DP with Fast Deflate Compression and 2:1 Compression	20Gb/s	20Gb/s

NOTE
Extension tunnels created on a Brocade 7840 Switch or Brocade SX6 Blade cannot connect to the FX8-24 Blade.

Brocade 7810 Switch

The Brocade 7810 Extension Switch has only one DP and the following table of Ethernet interface properties reflects this difference.

TABLE 16 Brocade 7810 Switch

Ethernet Interface Properties	Data Processor 0
1/10-GbE Interfaces (1G is the default; 10G with advanced license)	6
RJ-45 Ports	2
Maximum Bandwidth per VE	2.5Gb/s
Maximum WAN Bandwidth	2.5Gb/s
Maximum IP Extension Bandwidth	10Gb/s
Maximum Number of VE_Ports	4 VE_Ports (no VEX_Ports)
FCIP Bandwidth without Compression	2.5Gb/s
FCIP Bandwidth with 2:1 Compression	5Gb/s
FCIP Bandwidth with 4:1 Compression	10Gb/s

Brocade 7840 Extension Switch Ports

NOTE
You cannot connect extension tunnels created on a Brocade 7840 Switch to interfaces on an FX8-24 Blade. However, the Brocade 7840 Switch can connect with a Brocade 7840 Switch, a Brocade 7810 Switch, and a Brocade SX6 Blade in another Brocade X6 Director chassis.

The following figure illustrates the FC ports, 10/1GbE ports, and 40GbE ports on the Brocade 7840 Switch or Brocade 7810 Switch.

FIGURE 5 Brocade 7840 Switch Ports and Status Indicators

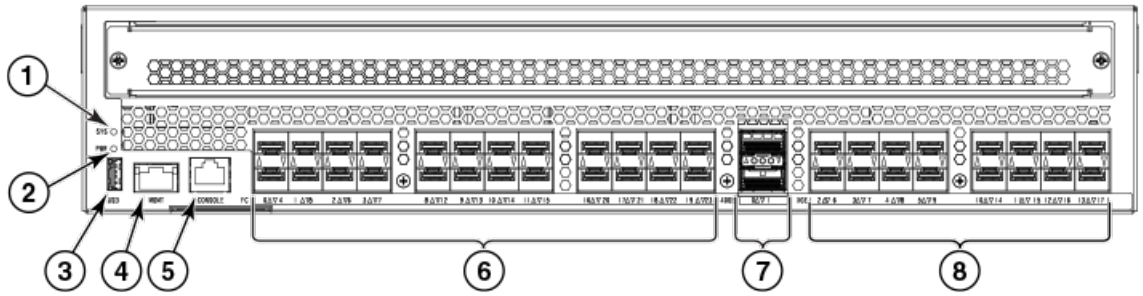

1. System (SYS) status LED
2. Power (PWR) LED
3. USB port
4. Ethernet management (mgmt) port
5. Serial console management port
6. FC ports 0 through 23
7. 40GbE ports 0 and 1
8. 1/10GbE ports 2 through 17

The Brocade 7840 Extension Switch provides 24 16-Gb/s FC ports (FC0–FC23) numbered 0 through 23 on the switch, two 40-GbE ports (ge0–ge1) numbered 0 through 1 on the switch, and 16 1/10-GbE ports (ge2–ge17) numbered 2 through 17 on the switch. Up to 20 VE_Ports are supported for tunnel configurations. Typically, only one VE_Port is needed per remote site.

NOTE
IPEX is not supported in 20VE mode, only in 10VE mode (the default).

NOTE
The 40 GbE ports are enabled with the Brocade 7840 Switch Upgrade 2 license.

Brocade 7810 Extension Switch Ports

The following illustration identifies system LEDs and ports on the port-side view of the Brocade 7810 Extension Switch.

FIGURE 6 Brocade 7810 Extension Switch Port Numbering

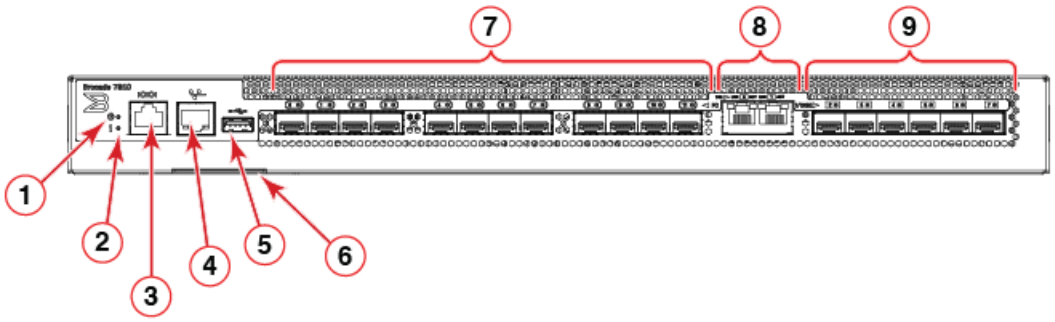

1. System power LED
2. System status LED
3. Serial console port (RJ-45)
4. Ethernet management port (RJ-45)
5. USB port
6. Serial number pull-out tab
7. 32 Gb/s SFP+ FC ports (0-11)
8. 1 GbE copper (RJ-45) ports (0-1)
9. 1/10 GbE SFP+ ports (2-7)

NOTE
All the ports are connected to a single Gen6 ASIC.

Brocade SX6 Extension Blade Ports

NOTE
You cannot connect extension tunnels created on a Brocade SX6 Blade to interfaces on a Brocade FX8-24 Blade. The Brocade SX6 Blade can connect with a Brocade SX6 Blade in another Brocade X6 Director chassis or with a Brocade 7840 or 7810 Extension Switch.

The following figure illustrates the FC ports, 10/1GbE ports, and 40GbE ports on the Brocade SX6 Extension Blade.

FIGURE 7 Brocade SX6 Extension Blade Port Numbering

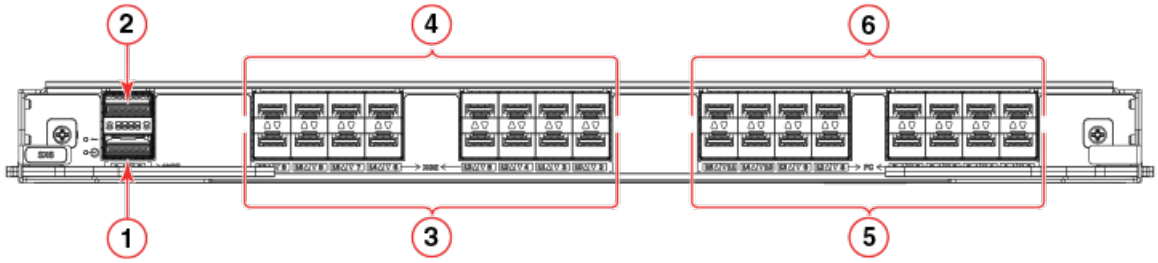

1. 40GbE port 0
2. 40GbE port 1
3. 10/1GbE ports 2 through 9 (right to left)
4. 10/1GbE ports 10 through 17 (right to left)
5. FC ports 0, 1, 2, 3, 8, 9, 10, 11 (right to left)
6. FC ports 4, 5, 6, 7, 12, 13, 14, 15 (right to left)

The Brocade SX6 Blade has the following external ports:

- Sixteen Fibre Channel (FC) SFP+ ports that support Fibre Channel Routing services and connection to FC devices for the Brocade Extension feature. These ports support 32Gb/s transceivers operating at 8, 16, or 32Gb/s or 16Gb/s transceivers operating at 4, 8, or 16Gb/s. The ports also support 10Gb/s transceivers. Supported port speed depends on the installed transceiver. FC ports can auto-negotiate speeds with connecting ports.
- Sixteen 10/1GbE SFP+ and two 40GbE QSFP ports. These ports allow connection of blades to IP WANs and allow Fibre Channel and IP I/O to pass through the IP WAN using extension tunnels. The 10/1GbE ports operate at 10Gb/s or 1Gb/s fixed speeds with appropriate 10Gb/s or 1Gb/s transceivers installed. The 40GbE QSFP ports operate at a 40Gb/s fixed speed.

Ethernet Port Groups

This section describes the numbered Ethernet ports on the Brocade 7810 Switch, the Brocade 7840 Switch, and the Brocade SX6 Extension platforms.

Brocade 7840 and Brocade SX6 Extension Switch

The Brocade 7840 Switch and the Brocade SX6 Blade support eight groups of Ethernet ports. Specific recommendations can be applied to ports within a group to help alleviate traffic congestion problems.

> **NOTE**
> For the Brocade 7810 Switch, each Ethernet port is its own group. So, the rest of this discussion does not apply.

Switch Ethernet ports are numbered with the 40GbE ports as 0 through 1. The 10GbE ports are numbered 2 through 17. See Brocade 7840 Extension Switch ports and Brocade SX6 Extension Blade ports for illustrations of the port numbering. Port numbers contained in the Ethernet port groups are shown in the following table.

TABLE 17 Brocade 7840 and Brocade SX6 Ethernet Port Groups

Port Number	Port Group
0, 1, 13, 17	1
2, 6	2
3, 7	3
4, 8	4
5, 9	5
10, 14	6

TABLE 17 Brocade 7840 and Brocade SX6 Ethernet Port Groups (continued)

Port Number	Port Group
11, 15	7
12, 16	8

Note that port group 1 contains the two 40GbE ports (0 and 1) and 10GbE ports 13 and 17. The remaining port groups contain the 10GbE ports from 2 to 16. Consider the following when using ports from these port groups:

- A port can block any port in its port group, but it cannot block a port outside of its port group.
- A port could affect another port in the same group due to differences in port speed or if the port is back-pressured due to Ethernet pause from an external switch. A blocked port may result from a slow-draining device or other congestion.

To avoid these effects on ports within the same port group, it is best that you do not mix speeds for ports within the group. Recommendations for the port groups are as follows:

- In port group 1, because the 40GbE ports are fixed at 40Gb/s, use either the 40GbE ports or the 10GbE ports at 10Gb/s or 1Gb/s.
- In port groups 2 through 8, which contain all 10GbE ports, configure the ports at either 10Gb/s or 1Gb/s.

  NOTE
  The table applies to ports configured in WAN mode. If the ports are configured as LAN ports, the grouping and blocking does not apply. As a recommended best practice, allocate a LAN port out of the same group as a WAN port.

Fibre Channel Port Groups

All Fibre Channel (FC) ports in a Brocade trunk must be in a single port group. The Brocade 7840 Switch, Brocade 7810 Switch, and the Brocade SX6 Blade provide FC port groups for configuring ISL trunk groups or Brocade trunks. Each port group contains eight FC ports.

Brocade 7840 Switch FC Port Groups

The trunk port groups on the Brocade 7840 Switch are as follows:

- Port group 0: ports 0–7
- Port group 1: ports 8–15
- Port group 2: ports 16–23

Brocade 7810 Switch FC Port Groups

The trunk port groups on the Brocade 7810 Switch are as follows:

- Port group 0: ports 0-7
- Port group 1: ports 8-11

Brocade SX6 Blade FC Port Groups

The trunk port groups on the Brocade SX6 Blade are as follows:

- Port group 0: ports 0–7
- Port group 1: ports 8–15

FC Port Group Considerations

The following requirements apply to forming trunk groups for the FC ports:

- All ports in a trunk group must belong to the same port group. For example, to form an 8-port trunk, select all eight ports from FC port group 0 or port group 1. You cannot use ports from each port group for the trunk.
- You can use from one to eight ports in a port group to form a trunk.
- Ports must be running at the same speed.
- Ports must be configured for the same distance.
- Ports must have the same encryption, compression, QoS, and FEC settings.
- Trunk groups must be created between Brocade switches (or Brocade adapters in the case of F_Port trunking). Brocade trunking is proprietary and is not supported on M-EOS or third-party switches.
- A direct connection must exist between participating switches.

For full details on trunking requirements and configuration, refer to the *Brocade Fabric OS Administration Guide*.

Network DP Components

The Brocade 7840 Switch, and the Brocade SX6 Blade share a similar design architecture. The following information discusses platform operation in FCIP mode, not Hybrid mode, which is when both FCIP and IP extension are enabled.

NOTE
The Brocade 7810 Switch operates only in Hybrid mode.

The following figures illustrate components and connections for each data processing (DP) complex when the platform is enabled in 10VE or 20VE modes. All 10-, 20-, and 40-Gb/s connections shown in the illustrations are full-duplex and internal to the platform. For more information about 10VE and 20VE port modes, see 10VE and 20VE Port Distribution on page 54.

NOTE
The following figures apply to the extension platform when it is in FCIP mode and not Hybrid mode.

FIGURE 8 DP Components and VE_Port Distribution in 10VE Mode

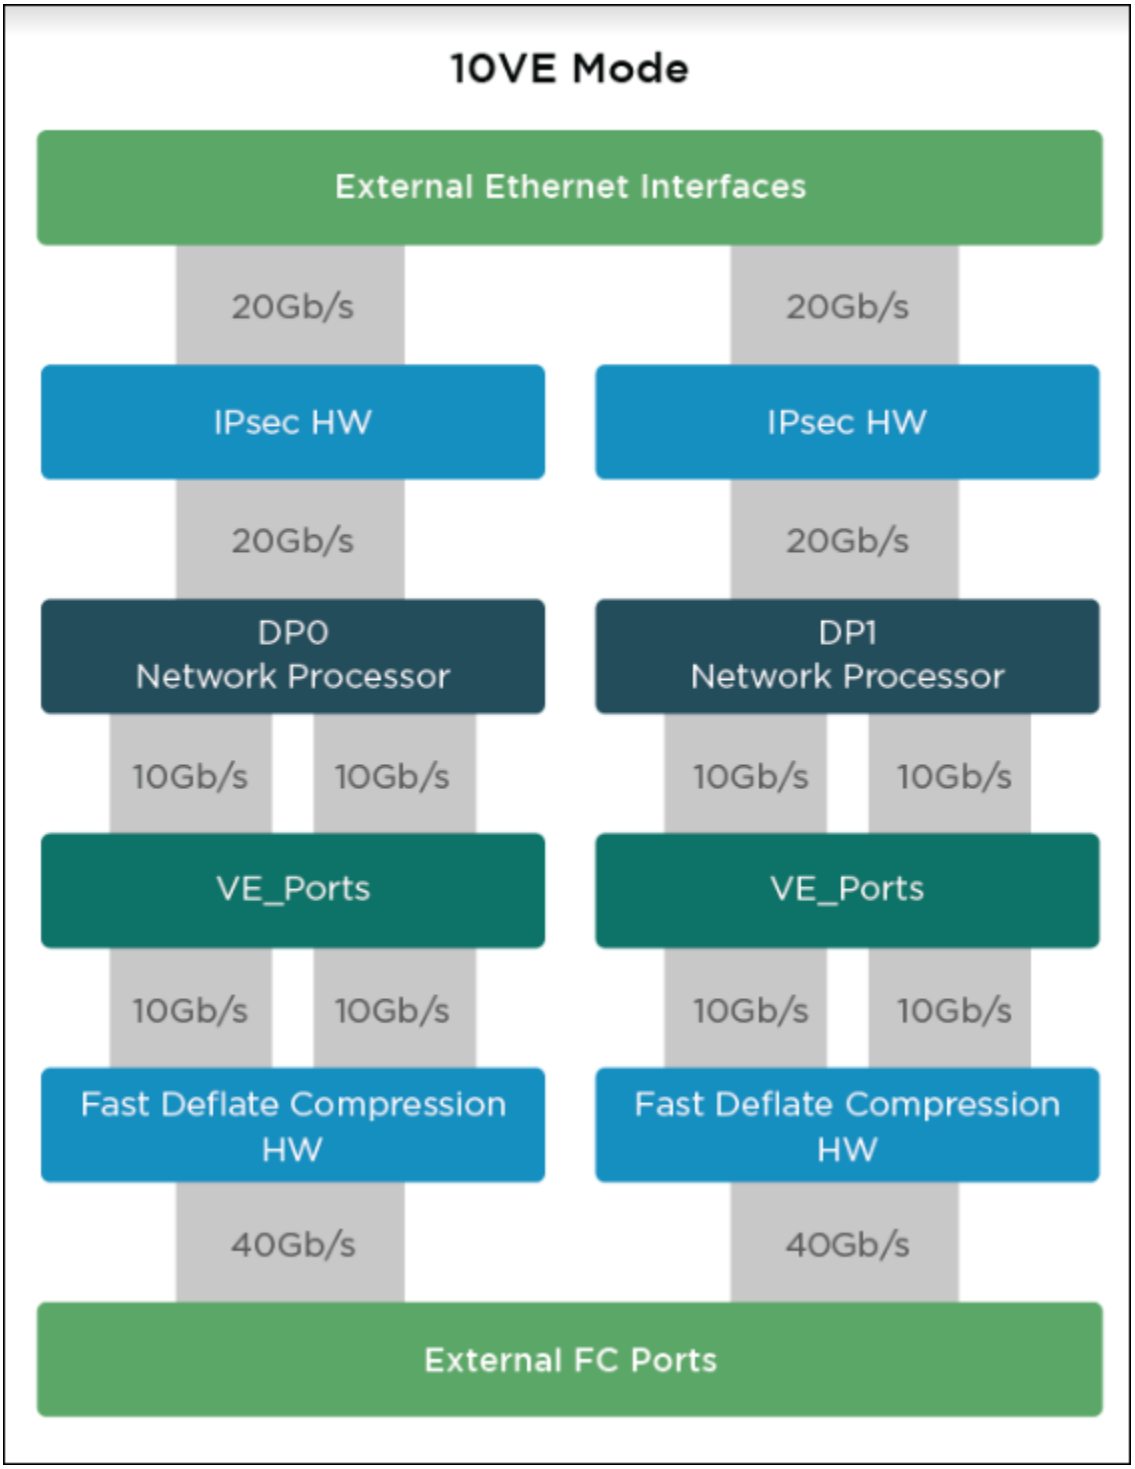

FIGURE 9 Brocade 7840 Switch DP Components and VE_Port Distribution in 20VE Mode

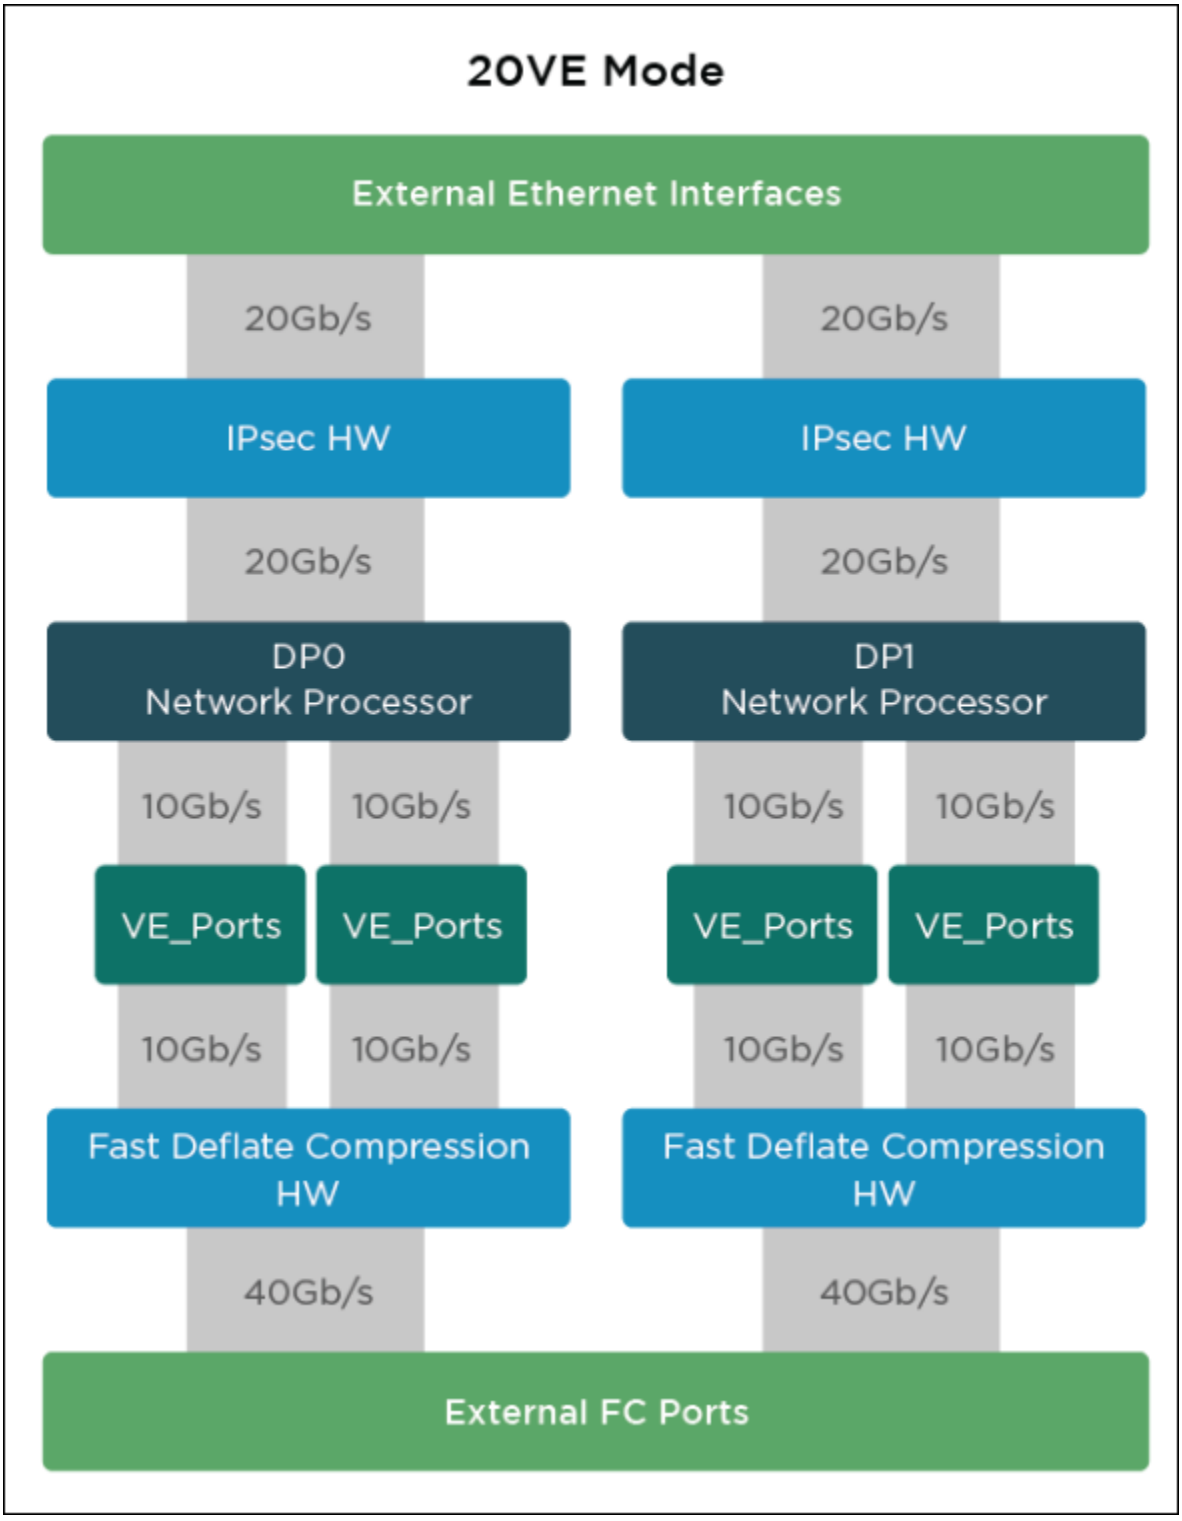

As shown in the illustrations:

- There is a 40Gb/s full-duplex connection between the FC switching ASIC and external FC ports and each DP. On the Brocade 7840 Switch, the external FC ports are Gen 5 (16Gb/s), and on the Brocade SX6 Blade, they are Gen 6 (32Gb/s).
- Fibre Channel (FC) frames can be compressed with the fast deflate compression hardware. Other software-based compression options can be configured.
- 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 the IPsec hardware using high-speed, low-latency, hardware-based encryptors. Each DP network processor can produce 20Gb/s of data flow going toward or coming from the external Ethernet interfaces and the WAN.
- In 10VE mode, each DP complex supports a total of five VE_Ports. 10VE mode is required for Hybrid mode operation.
- In 20VE mode, each DP complex supports a total of ten VE_Ports. 20VE mode is allowed in FCIP mode only.

If a 2:1 compression ratio is achieved using fast deflate compression, 40Gb/s per DP and 80Gb/s per platform is available to external FC ports (applicable to the Brocade 7840 Switch and the Brocade SX6 Blade but not the Brocade 7810 Switch).

The maximum compression ratio depends on a number of factors and is not guaranteed.

> NOTE
> Typical deflate compression may achieve different compression ratios. Brocade makes no promises as to the achievable compression ratios for customer-specific data.

The Adaptive Rate Limiting (ARL) aggregate of all circuit maximum values on a single DP complex cannot exceed 40 Gb/s. The ARL aggregate of all circuit minimum values for a single DP complex cannot exceed 20Gb/s. All circuits includes all circuits from all tunnels, not just all circuits from a single tunnel. ARL is used only when minimum and maximum bandwidth values are configured for a circuit.

Specific VE_Ports are associated with each DP complex. The VE_Port that you configure for a tunnel also selects the DP complex that is used for processing. See 10VE and 20VE Port Distribution on page 54 for more information.

For additional specifications and requirements for switch ports, tunnels, and circuits, see Tunnel and Circuit Requirements for Brocade Extension Platforms on page 62.

10VE and 20VE Port Distribution

In the Brocade 7840 Switch and Brocade SX6 Blade, each data processor (DP) complex, DP0 and DP1, can support either five or ten VE_Ports depending on how the VE mode is configured. In 10VE mode, each DP supports five VE_Ports, for a total of ten VE_Ports per switch or blade. In 20VE mode, each DP supports up to ten VE_Ports, for a total of twenty VE_Ports per switch or blade. The following table shows the VE_Ports supported in 10VE and 20VE mode. The VE_Port number is used when you configure a tunnel. The number of tunnels that you can configure per switch or blade is determined, in part, by the VE mode and the VE_Ports supported by each DP.

TABLE 18 10VE and 20VE Port Distribution

Extension Platform	VE Mode	DP0/DP1	Supported VE_Ports
Brocade SX6	VE10	DP0	16-20
		DP1	26-30
	VE20	DP0	16-25
		DP1	26-35
Brocade 7840	VE10	DP0	24-28
		DP1	34-38
	VE20	DP0	24-33

TABLE 18 10VE and 20VE Port Distribution (continued)

Extension Platform	VE Mode	DP0/DP1	Supported VE_Ports
		DP1	34–43

In 10VE mode, a VE_Port (and the total bandwidth of its circuits) can use all Fibre Channel bandwidth available to the DP complex where it resides, a maximum of 20 Gb/s.

In 20VE mode, a single VE_Port (and the total bandwidth of its circuits) on a DP complex can use half the Fibre Channel bandwidth available to the DP complex where it resides, to a maximum of 10 Gb/s. This option allows you to use more VE_Ports, but at a lower maximum bandwidth.

10GbE and 40GbE Port and Circuit Considerations

Enhanced 10GbE and 40GbE port operations require special considerations when configuring circuits, tunnels, failover operations, and bandwidth.

On a single VE port on a Brocade 7840 Switch or Brocade SX6 Blade, each tunnel that you create is limited to a maximum of ten circuits. (The limit on a Brocade 7810 Switch is six circuits.) The maximum committed rate of a single circuit is 10 Gb/s, whether configured on a 10GbE or 40GbE port.

For a complete list of tunnel, circuit, and IP address requirements and capacities, see Tunnel and Circuit Requirements for Brocade Extension Platforms on page 62.

Brocade 7840 License Options

Important Brocade Extension features and FICON extension capabilities of the Brocade 7840 Switch require the feature licenses shown in the following table. Use the `licenseshow` command to display license keys and licenses currently installed.

TABLE 19 Brocade 7840 Feature Licenses

Feature	Purpose	License (**licenseShow** output)
WAN Rate Upgrade 1	Increases bandwidth available to all extension tunnels configured on the switch from 5Gb/s for the base hardware to 10Gb/s.	WAN Rate Upgrade 1 license
WAN Rate Upgrade 2	Allows unlimited bandwidth for all tunnels configured on the switch. This also enables the 40GbE ports so that they can be used for configuring IP addresses. **NOTE** You must have a WAN Rate Upgrade 1 license to activate the WAN Rate Upgrade 2 license.	WAN Rate Upgrade 2 license
Advanced FICON acceleration	Enables accelerated tape read/write, IBM z/OS Global Mirror, and Teradata emulation features in FICON environments. Slot-based license.	Advanced FICON Acceleration (FTR_AFA) license
Advanced Extension License	This is enabled on the Brocade 7840 Switch at the factory. Required for multiple-circuit tunnels, Trunking, and ARL.	Advanced Extension (FTR_AE) license

For complete information about the licenses described in the preceding table and additional licenses available for the Brocade 7840 Switch, refer to the *Brocade Fabric OS Software Licensing Guide*.

Brocade 7810 License Options

The Brocade 7810 Switch supports a base model that you can upgrade with an upgrade license. Prior to release FOS 8.2.0, we offered distinct licenses ('Upgrade 1' and 'Upgrade 2') as well as 'Advanced Extension License', and an 'Advanced FICON Acceleration License'. With release 8.2.1, we simplify the process with a single 'Upgrade License' for the Brocade 7810. The base model and the base model w/ upgrade license will support the following.

NOTE

Be aware that with the simplified process, the base model w/upgrade license is identical to a fully-configured model.

TABLE 20 Brocade 7810 License Options

Functionality	Base Brocade 7810 Switch	Base Brocade 7810 Switch w/Upgrade License
FC Port Limits	4	12
VE Port Limits	2	4
GE WAN Port Limits	2	6
GE LAN Port LImits	4	4
GE Port Speed Limit	1Gb/s	10Gb/s
WAN Bandwidth Limit	1Gb/s	2.5Gb/s
Max IPEX flows (LAN TCP connections)	128	128
Max UDP connections	32	32
IPsec Support	Yes	Yes
Compression Support	Yes	Yes
FC Trunking	No	Yes Yes
Extension Trunking	No	Yes
Fabric Vision	No	Yes
Adaptive Rate Limiting	Yes	Yes
FICON Support	No	No
FC Routing	No	Yes
Extended Fabrics	No	Yes

Brocade SX6 License Options

The Brocade SX6 Blade is supported in the Brocade X6 Director platform. The Brocade X6-4 and X6-8 Directors support the following licenses:

- Extended Fabrics
- FC Trunking
- FICON Management Server
- Inter Chassis Link
- Integrated Routing Ports on Demand
- Fabric Vision and IO Insight

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

Brocade FX8-24 Extension Blade Overview

This section provides information on ports, circuits, and tunnels specific to the Brocade FX8-24 Extension Blade.

The blade can be installed in a Brocade DCX 8510-8 or DCX 8510-4 chassis.

The following figure shows the FC ports, GbE ports, and 10GbE ports on the Brocade FX8-24. There are 12 FC ports, numbered 0 through 11. The FC ports can operate at 1, 2, 4, or 8Gb/s. There are ten GbE ports, numbered 0 through 9. Ports xge0 and xge1 are 10GbE ports.

FIGURE 10 Brocade FX8-24 Extension Blade

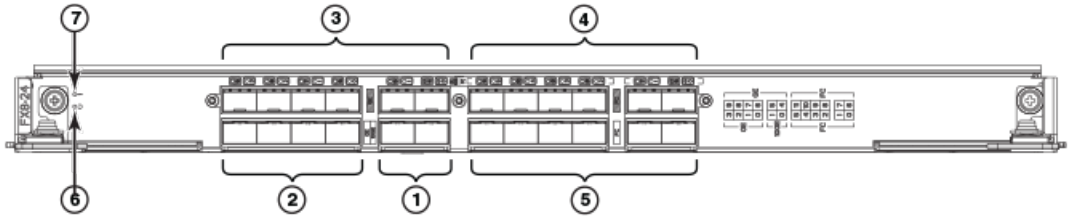

1. 10GbE XGE ports 0–1 (right to left)
2. 1GbE ports 0–3 (right to left)
3. 1GbE ports 4–9 (right to left)
4. FC ports 6–11 (right to left)
5. FC ports 0–5 (right to left)
6. Power LED
7. Status LED

Brocade FX8-24 Operating Modes

The Brocade FX8-24 operates in one of three modes. Depending on the mode, different combinations of ports and speeds are available. The three modes are:

- 1Gb/s mode: You can use all ten GbE ports (0 through 9). Both XGE ports are disabled.
- 10Gb/s mode: You can use the xge0 and xge1 ports. GbE ports (0 through 9) are disabled.
- Dual mode: You can use GbE ports 0 through 9 and port xge0. The xge1 port is disabled.

Brocade FX8-24 Data Processor Complexes

The Brocade FX8-24 contains two FCIP data processor (DP) complexes, DP0 and DP1. With two DP complexes, you can configure one tunnel on each DP and achieve a maximum of 20Gb/s full-duplex bandwidth for tunnel connections on each blade. Each DP complex has an associated local 10GbE XGE port. Each DP complex controls a specific range of GbE and VE_Ports.

The DP0 interface and ports are as follows:

- 10GbE XGE port 0
- VE_Ports 22 through 31
- 10Gb/s maximum bandwidth
- Operates in 10Gb/s mode and dual mode

The DP1 interface and ports are as follows:

- 10GbE XGE port 1 (10Gb/s mode only)
- VE_Ports 12 through 21
- 10Gb/s maximum bandwidth

- GbE ports 0 through 9 while operating in 1-Gb/s mode and dual mode

Removing the Brocade FX8-24 Extension Blade

You must perform the following steps before removing or changing the location of a Brocade FX8-24 Blade in a chassis. Otherwise, unexpected results might occur.

> ATTENTION
> If you are permanently removing a blade from a Brocade DCX, DCX-4S, DCX 8510-8, or DCX 8510-4 chassis to relocate to another slot in the chassis or if you are removing the blade from the chassis entirely, you must follow these procedures *before removing the blade*.

1. Delete all fciptunnel configurations. Use the `portcfg fciptunnel` *`slot/ve_port`* command.
2. Delete all IP routes defined on the blade. Use the `portcfg iproute` command.
3. Delete all IPIFs defined on the blade. Use the `portcfg ipif [`*`slot/geX`*`]|`*`xgxX`* command.
4. If logical switches are used on the switch, move all FX8-24 ports to the default logical switch. Use the `lscfg --config` *`FID slot/port`* command.
5. Remove the blade from the chassis.

Brocade FX8-24 Blade License Options

Important capabilities of the Brocade FX8-24 Blade require the feature licenses shown in the following table. Use the `licenseShow` command to display license keys and licenses currently installed.

TABLE 21 Brocade FX8-24 Feature Licenses

Feature	Purpose	License (**licenseShow** output)
10GbE support	Allows 10Gb/s operation on 10GbE ports. Slot-based license.	10Gigabit FCIP/Fibre Channel (FTR_10G) license
Advanced FICON acceleration	Enables accelerated tape read/write, IBM z/OS Global Mirror, and Teradata features in FICON environments. Slot-based license.	Advanced FICON Acceleration (FTR_AFA) license
Integrated routing (IR)	Required to configure EX_Ports and VEX_Ports to support Fibre Channel Routing (FCR). Chassis-based license.	Integrated Routing license
Advanced extension	Required for multiple-circuit tunnels, trunking, Adaptive Rate Limiting (ARL), and other features. Slot-based license.	Advanced Extension (FTR_AE) license

For complete information about the licenses described in the preceding table and additional licenses available for the switch, refer to the *Brocade Fabric OS Software Licensing Guide*.

Brocade FX8-24 Blade Multi-gigabit Circuits

For each 10GbE XGE port on a Brocade FX8-24 Blade, you can configure multi-gigabit circuits. You can allocate bandwidth for each circuit in 1Gb/s increments up to the maximum number of circuits and bandwidth allowed. For example, you can configure one 10GbE XGE port with two 5Gb/s circuits or a single 10Gb/s circuit. A maximum of ten circuits can be configured on a single XGE port. The maximum committed rate for a circuit between 10GbE XGE ports is 10 Gb/s. The Brocade FX8-24 Blade at each end of the tunnel must be running the same version of Fabric OS 7.0 or later if the committed rate for circuits exceeds 1 Gb/s.

To use both 10GbE XGE ports, the Brocade FX8-24 Blade must be in 10Gb/s mode. When the blade is in dual mode, you can use port xge0 for multi-gigabit circuits, but not port xge1. When the blade is in 1Gb/s mode, the 10GbE ports are not available.

NOTE
There is no difference in latency or throughput performance for single or multi-gigabit circuits.

Crossports and Failover

A cross-port is the non-local DP XGE port. The Brocade FX8-24 Blade provides two data processing (DP) complexes identified as DP0 and DP1. Each DP has a local 10Gb/s XGE port, xge0 and xge1 that corresponds to DP0 and DP1. You can configure a DP to use its non-local XGE port, which is done to provide an alternate traffic path if the local XGE port fails for some reason. Cross-ports are supported only on the Brocade FX8-24 Blade and are available when the blade is configured for 10Gb/s mode.

For DP0 and its local xge0 port, the cross-port is xge1. Likewise, for DP1 and its local xge1 port, the cross-port is xge0.

Typically, IP interface addresses (IPIFs) used by ge0 through ge9 and xge1 are used for any circuits that use VE_Ports 12 through 21. The xge1 port is the local XGE interface for VE_Ports 12 through 21. Likewise, IP addresses configured for xge0 are used by circuits for VE_Ports 22 through 31.

Configure a cross-port by assigning an IP address to the remote XGE port that can be used by the local XGE port. For example, assigning an IP address to xge0 as a cross-port makes the address available on the remote xge0 for VE_Ports 12 through 21 on the local xge1.

You can also assign IP routes (iproutes) used by the local port, VLAN tagging, and circuits with metrics to the remote XGE port to allow failover to the cross-ports.

Cross-ports contain the IPIFs and IP routes that belong to the remote interface. To use crossports, both XGE ports must be configured in 10Gb/s mode.

Bandwidth Allocation and Restrictions

There are specific bandwidth allocations and restrictions for the Brocade FX8-24 Blade that are important to review when configuring tunnels and circuits.

Front-end and Back-end Bandwidth

The Brocade FX8-24 Blade contains an internal port complex with 1Gb/s ports to support the blade's VE_Port groups, data processor (DP) complexes, GbE ports, and XGE ports.

Each DP complex has 10Gb/s (full-duplex) of available bandwidth. Therefore, each VE_Port group (VE_Port 22-31 and VE_Port 12-21) has 10Gb/s of bandwidth available to the internal port complex back-end ports. When the tunnels using VE_Ports in a specific VE_Port group consume the group's back-end bandwidth, additional circuits cannot be created for those tunnels. The port complex has 10Gb/s of front-end bandwidth available for each of the XGE ports. Tunnels (VE_Ports) cannot consume more than 10Gb/s of bandwidth over an XGE port. The internal port complex has another 10Gb/s of bandwidth available for the cross-port.

The following figure illustrates the internal DP complex with VE_Port groups, internal port complex, front-end and back-end port areas, and the cross-port (xport) on a Brocade FCX8-24 Blade.

FIGURE 11 Internal Port and DP Complexes on a Brocade FX8-24 Blade

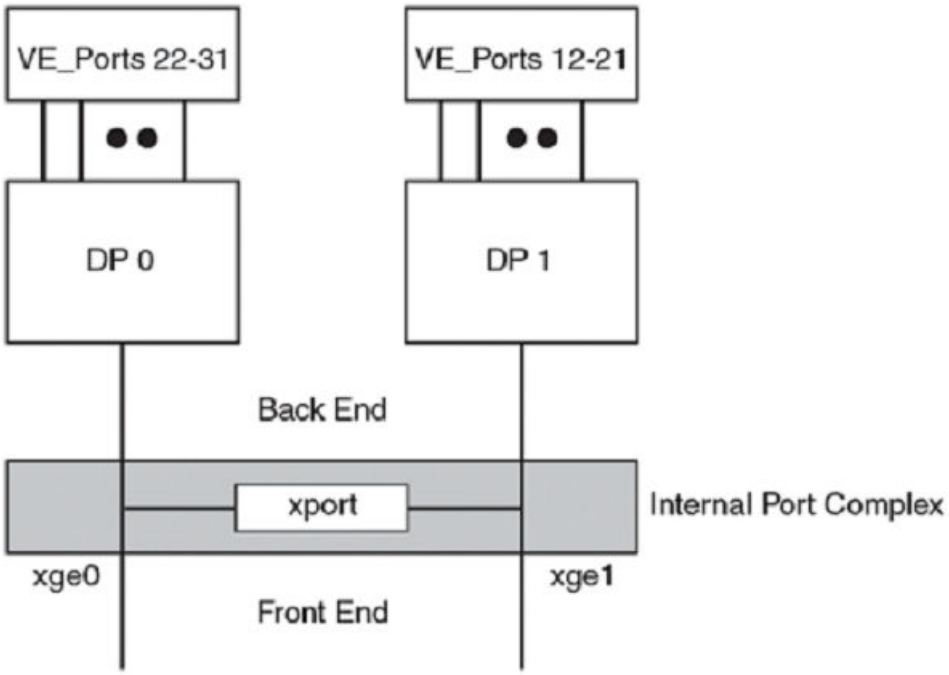

Calculating Back-end Bandwidth

The following are the ways to configure the back-end bandwidth for a tunnel and DP complex:

- To calculate the consumed bandwidth for a tunnel, round the maximum committed rates for all metric 0 circuits up to the next whole rate (for example 1.5Gb/s becomes 2Gb/s) and add them up. Then add the rounded-up maximum committed rates for all metric 1 circuits. The greater of the two values is the consumed bandwidth for a tunnel.
- To calculate the total consumed back-end port bandwidth for a DP complex, add the consumed bandwidth for each tunnel in the DP complex VE_Port group. The total cannot exceed 10Gb/s.
- Back-end bandwidths are always rounded up for each VE_Port group. For example, a circuit defined as 1.5Gb/s will consume 2Gb/s of back-end bandwidth.

Calculating Front-end Bandwidth

The following are the ways to calculate the front-end bandwidth for a tunnel and a XGE port:

- To calculate the front-end bandwidth usage on a per-tunnel and per-XGE port basis, add the consumed bandwidth for all metric 0 circuits for a tunnel using xge0 or xge1. Add the total consumed bandwidth for all metric 1 circuits for the tunnel. The greater of the two values is the total front-end port bandwidth usage for an xge0 or xge1 tunnel. See Circuit Failover on page 78 for more information on assigning metrics to circuits.
- Each XGE port is allocated 10Gb/s of front-end bandwidth. The total consumed front-end port bandwidth cannot exceed 10Gb/s per XGE port.

Calculating Cross-port Bandwidth

The DP complexes share only one cross-port, so total available bandwidth is 10Gb/s for all VE_Ports on the blade, regardless of the DP complex to which the VE_Ports belong. For more information on cross-ports, see Crossports and Failover on page 59.

- To calculate the bandwidth consumed by a cross-port on a per-tunnel basis, add the consumed bandwidth for all metric 0 circuits in the tunnel that use the cross port. Add the total consumed bandwidth for all metric 1 circuits in the tunnel that use the cross-port. The greater of the two values is the total cross-port-consumed bandwidth for the tunnel .
- The total crossport-consumed bandwidth is the total of the bandwidth for the tunnels using VE_Ports 12 through 31. The total cross-port-consumed bandwidth cannot exceed 10Gb/s.

ARL Limits

Bandwidth allocations are subject to the minimum committed rate (-b) and maximum committed rate (-B) set for circuits and tunnels using the Adaptive Rate Limiting (ARL) feature. For more information on ARL and ARL restrictions, refer to Adaptive Rate Limiting on page 27.

Failover Circuits and Groups

When considering the 10Gb/s bandwidth limit for each DP complex on a Brocade FX8-24 Blade, you must also consider failover circuits configured for VE_Ports in each complex. For example, you cannot create a circuit to use an address for a cross-port if a 10Gb/s failover circuit is assigned to the VE_Port on that cross-port. If the failover circuit were to come online, there would be no available bandwidth for the new circuit.

Failover groups allow you to define a set of metric 0 and metric 1 circuits that are part of a failover group. When all metric 0 circuits in the group fail, metric 1 circuits take over operation, even if there are metric 0 circuits still active in other failover groups. Typically, you would configure only one metric 0 circuit in a failover group. For detailed information, see Circuit Failover on page 78.

When calculating total bandwidth usage for a tunnel, you must also add the total bandwidth usage per failover group.

To calculate the total bandwidth usage for the failover group in the tunnel, for each failover group (0 through 9), perform the following steps:

1. Add the consumed bandwidth for all metric 0 circuits in the failover group.
2. Add the total consumed bandwidth for all metric 1 circuits in the failover group.

The greater of the two values is the total bandwidth usage for the failover group in the tunnel.

Bandwidth Allocation Example

The basis for all bandwidth calculations is determining how much bandwidth a given tunnel is consuming. The next step is determining where the tunnel is consuming that bandwidth. You must consider the 10Gb/s limits for back-end ports, front-end ports, and cross-ports. A tunnel, at the least, uses back-end and front-end port bandwidth. If cross-port circuits are configured, then the tunnel uses cross-port bandwidth as well.

For example, suppose that two 10Gb/s circuits are configured for a tunnel on VE_Port 12. Circuit 0 has a metric of 0 on xge1, and circuit 1 is a failover circuit with a metric of 1 on xge0 (refer to the figure under Front-end and Back-end Bandwidth on page 59). Note that configuring circuit 1 on xge0 is a cross-port configuration. Although this configuration is allowed, you cannot create additional circuits for VE_Port group 12 through 21 or group 22 through 31 for the following reasons:

- For VE_Port group 12 through 21, VE_Port 12 is consuming the maximum 10Gb/s of allocated back-end port bandwidth. See Calculating Cross-port Bandwidth on page 61.
- You cannot create a cross-port so that VE_Ports 22 through 31 use xge1 because VE_Port 12 is consuming the maximum 10Gb/s of cross-port bandwidth for its failover circuit. See Calculating Cross-port Bandwidth on page 61.

- If VE_Port 12 fails, all 10Gb/s traffic will flow over the cross-port and the xge0 front-end port. If additional circuits were already configured for the VE_Port 22 through 31 group, the front-end port bandwidth would exceed the 10Gb/s limit for xge0. See Calculating Front-end Bandwidth on page 60.

Tunnel and Circuit Requirements for Brocade Extension Platforms

Each Brocade Extension platform has specific considerations for tunnel and circuit requirements. The general tunnel and circuit requirements that apply to Brocade Extension platforms are as follow:

- You can define multiple addresses on Ethernet ports to configure multiple circuits. Multiple circuits can be configured as a trunk, which provides multiple source and destination addresses to route traffic across an IP network, provide load leveling, and provide failover capabilities.
- The committed rate for a circuit associated with a physical port cannot exceed the rate of the port, or 10Gb/s, whichever is lower.
- In a scenario where a tunnel has multiple circuits of different metrics (0 or 1), circuits with higher metrics (1) are treated as standby circuits and are only used when all lower metric (0) circuits fail. Using Circuit Failover Grouping, you can better control which metric 1 circuits will be activated if a metric 0 circuit fails.
- When the spillover is configured, circuit metrics determine which metric 1 circuits to use when capacity exceeds the metric 0 circuits.
- A circuit defines source and destination IP addresses on each end of a tunnel.
- If the circuit source and destination IP addresses are not on the same subnet, an IP static route (iproute) must be defined on both sides of the tunnels that designates the gateway IP addresses.
- As a best practice, all tunnel and circuit settings should be identical on both sides of the tunnel. This includes committed bandwidth, IPsec, compression, ARL minimum and maximum, Fastwrite, OSTP, FICON tunnel, and keepalive timeout values (KATOV). You must configure the tunnel and circuit parameters correctly on both sides of the network, otherwise the tunnel or circuit will fail to come online.
- VE_Ports or VEX_Ports cannot connect in parallel to the same domain at the same time as Fibre Channel E_Ports or EX_Ports.
- When load-leveling across multiple circuits, the difference between the ARL minimum data rate set on the slowest circuit in the trunk and the fastest circuit should be no greater than a factor of four. For example, a 100Mb/s circuit and a 400Mb/s circuit will work, but a 10Mb/s and 400Mb/s circuit will not work. This ensures that the entire bandwidth of the trunk can be utilized. If you configure circuits with the committed rates that differ by more than a factor of four, the entire bandwidth of the trunk might not be fully utilized.
- If the settings are not the same on both sides of the tunnel, Op Status displays `UpWarn` as shown in the following example:

```
switch:admin> portshow fciptunnel all -c
-------------------------------------------------------------------------------
 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
-------------------------------------------------------------------------------
 8/12   -         Up       -----i-    2d54m    0.00    0.00    4      -      -/-
 8/12   0 8/xge1  UpWarn   ---4--s    2d54m    0.00    0.00    4  5000/5000  0/-
 8/24   -         Empty    -------       0s    0.00    0.00    0      -      -/-
-------------------------------------------------------------------------------
 Flags:  tunnel: c=compression m=moderate compression a=aggressive compression
                 A=Auto compression f=fastwrite t=Tapepipelining F=FICON
                 T=TPerf i=IPSec l=IPSec Legacy
 Flags: circuit: s=sack v=VLAN Tagged x=crossport 4=IPv4 6=IPv6
                 L=Listener I=Initiator
```

Brocade 7840 Switch, Brocade 7810 Switch, and Brocade SX6 Blade

This section lists requirements and specifications for tunnels, circuits, and ports on the Brocade 7840 Extension Switch, the Brocade 7810 Extension Switch, and the Brocade SX6 Blade.

IP addresses and routes:

- You can configure maximum 60 IP addresses per DP complex.
- You can define up to 128 routes per GbE port; however, you can define only 120 IP routes per DP. For example, you can configure 64 IP routes defined on ge2.dp0 and another 64 IP routes defined on ge2.dp1.

VE_Ports and VE_Port groups:

NOTE
When the extension platform operates in Hybrid mode, 20VE mode is not allowed.

NOTE
VE mode is not supported on the Brocade 7810 Extension Switch.

TABLE 22 10VE and 20VE Port Distribution

Extension platform	VE mode	DP0/DP1	Supported VE_Ports
Brocade 7840	VE10	DP0	24–28
		DP1	34–38
	VE20	DP0	24–33
		DP1	34–43
Brocade SX6	VE10	DP0	16–20

NOTE

- You can have a maximum 20 VE_Ports on the switch. In the default 10VE mode, only 10 VE_Ports are enabled. In 20VE mode, all 20 VE_Ports are enabled.
- On the Brocade 7840 Switch, there are two VE_Port groups in 10VE mode. Each port group can share 20 Gb/s.
  - DP0 controls VE_Ports 24-28.
  - DP1 controls VE_Ports 34-38.
  - The remaining VE_Ports 29-33 and 39-43 are disabled.
- In 20VE mode on the Brocade 7840 Switch , there are four VE_Port groups. Each port group can share 10Gb/s.
  - DP0 controls VE_Ports 24-28 and VE_Ports 29-33.
  - DP1 controls VE_Ports 34-38 and VE_Ports 39-43.
- On the Brocade SX6 Blade, there are two VE_Port groups in 10VE mode. Each port group can share 20Gb/s.
  - DP0 controls VE_Ports 16-20.
  - DP1 controls VE_Ports 26-30.
  - The remaining VE_Ports 21-25 and 31-35 are disabled.
- In 20VE mode on a Brocade SX6 Blade, there are four VE_Port groups. Each port group can share 10Gb/s.
  - DP0 controls VE_Ports 16-20 and VE_Ports 21-25.
  - DP1 controls VE_Ports 26-30 and VE_Ports 31-35.
- On a Brocade 7810 Switch, DP0 controls VE_Ports, 12-15.
- VE_Ports are not associated with a particular Ethernet port.
- As a best practice, do not utilize multiple VE_Ports between the same domains. Instead, create a single tunnel (one VE_Port) with multiple circuits between the domains.

- As a best practice, do not mix VE_Ports and E_Ports (or VEX_Ports and EX_Ports) between the same domains.
- VEX_Ports are not supported.

Bandwidths, maximum and minimum rates:

- For a VE_Port group, the sum of the minimum committed rates of that group's circuits cannot exceed 10Gb/s when the switch is in 20VE mode and 20Gb/s when the switch is in 10VE mode.
- The minimum committed rate for all VE_Ports in one DP complex cannot exceed 20 Gb/s. The maximum rate for all VE_Ports in one DP complex cannot exceed 40Gb/s.
- The minimum committed rate for a circuit is 20Mb/s.
- The maximum committed rate for a circuit is 10Gb/s.
- With compression, total bandwidth cannot exceed 80Gb/s (40 Gb/s per DP) on the Fibre Channel side.
- The difference between the guaranteed (minimum) and maximum bandwidth for a tunnel cannot exceed the 5:1 ratio.
- The Brocade 7810 Switch allows a maximum of 2.5Gb/s across all VEs.

Circuits:

- The number of circuits that you can configure on an Ethernet port is limited only by the number of IP address pairs available and how the addresses are allocated. Each circuit requires a unique IP address pair. A unique pair means that one of the two addresses (local IPIF and remote IPIF) be unique in the pair. For example, the following address pairs use the same source address in each pair but the destination addresses are different, therefore each 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
- You can configure a maximum of 10 circuits for a trunk (VE_Port).
- You can configure a maximum of 40 circuits per DP.
- The Brocade 7810 Switch allows a maximum of 6 circuits per tunnel.

Brocade FX8-24 Requirements

The Brocade FX8-24 has the following requirements and specifications for tunnels, circuits, and ports.

IP addresses and routes:

- You can define up to 10 IP addresses for a 10GbE port and an additional 10 addresses on crossports when operating in 10Gb/s mode.
- You can define up to eight IP addresses for a 1GbE port.

VE_Ports, VE_Port groups, VEX_Ports:

- A Brocade FX8-24 Blade can support 20 VE_Ports, and therefore 20 extension tunnels.
- There are two VE_Port groups. DP1 controls ports numbered 12 through 21 and DP0 controls ports numbered and 22 through 31.
- Each tunnel is identified with a VE_Port number.
- VE_Ports are not associated with a particular Ethernet port.
- The blade also supports VEX_Ports to avoid the need to merge fabrics.
- VE_Port versus Ethernet port usage depends on the blade operating mode as follows:
  - 1Gb/s mode: VE_Ports 12 through 21 are available to use GbE ports 0 through 9. VE_Ports 22 through 31, xge0, and xge1 are not available.

  - In 10Gb/s mode, VE_Ports 12 through 21 are available to use xge1; VE_Ports 22 through 31 are available to use xge0. GbE ports 0 through 9 are not available.
  - In 10Gb/s mode, you can also configure VE_Ports 12 through 21 to use port xge0 as a crossport and VE_Ports 22 through 31 to use port xge1 as a crossport.
  - In dual mode, VE_Ports 12 through 21 are available to use GbE ports 0 through 9; VE_Ports 22 through 31 are available to use xge0. Port xge1 is not available.

Circuits:

- A limit of 20 circuits can be configured per VE_Port group (12 through 21 or 22 through 31) when using a 10GbE port. For the 20 circuits, 10 are configured on local ports and 10 on crossports.
- You can configure up to 10 circuits for a trunk (VE_Port).
- The Brocade FX8-24 Blade contains two 10GbE ports. You can define up to 10 circuits per trunk spread across the 10GbE ports.
- A limit of 10 circuits can be configured on a single 10GbE port. Each circuit requires a unique IP address.
- The blade contains ten 1GbE ports. You can define up to 10 circuits per trunk spread across the GbE ports.
- A limit of four circuits can be configured on a single 1GbE port. Each circuit requires a unique IP address.
- On the Brocade FX8-24 Blade, a unique pair means that **both** of the addresses (local IPIF and remote IPIF) must be unique and cannot be reused. For example, the following address pairs use the same source address in each pair, only the destination addresses are different. Because the source addresses are the same, the pair is not 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

Bandwidths, maximum and minimum rates:

- For a Brocade FX8-24 Blade with a VE_Port group on a 10GbE port, the sum of the maximum committed rates of that group's circuits cannot exceed 10Gb/s.
- For ARL, configure minimum rates of all the tunnels so that the combined rate does not exceed 20Gb/s for all VE_Ports on the blade.
- For ARL, you can configure maximum rate of 10Gb/s for all tunnels over a single 10GbE port and 10Gb/s for any single circuit.
- The minimum committed rate for a circuit is 10Mb/s.
- A circuit between 1GbE ports cannot exceed the 1Gb/s capacity of the interfaces rate.

Brocade IP Extension

Brocade IP Extension is supported on the following platforms:

- Brocade 7810 Extension Switch
- Brocade 7840 Extension Switch
- Brocade SX6 Extension Blade installed in a supported platform, such as the Brocade X6-4 Director or Brocade X6-8 Director

IP Extension provides Layer 3 extension for IP storage replication. The IP extension platform acts as the gateway for the LAN, but there is no Layer 2 extension, which means each side of the network must be on a different subnet.

The extended IP traffic receives the same benefits as does traditional FCIP traffic:

- Compression
- High speed encryption

- Frame based load leveling and lossless failover
- Network bandwidth management through rate shaping and QoS

IP Extension requires that you configure the switch or blade to operate in Hybrid mode, which only supports 10VE mode. Configuring Hybrid mode is disruptive because a reboot is required to load the hybrid mode image.

NOTE
This is not applicable to the Brocade 7810 Extension Switch, which is always in Hybrid mode.

Internal connections are re-mapped to provide 20Gb/s of LAN traffic,10Gb/s of FC traffic, and a maximum of 20Gb/s of WAN traffic on the Brocade 7840 Switch and Brocade SX6 Blade, and 10Gb/s of LAN traffic, 20Gb/s of FC traffic, and a maximum of 2.5Gb/s of WAN traffic on the Brocade 7810 Switch

When in Hybrid mode, the switch or blades allows up to eight of the 10GbE ports to be configured as LAN ports. LAN ports are not grouped, as opposed to WAN ports which are grouped. In general, LAN ports do not block each other, and LAN ports do not block WAN ports. The recommended best practice is to pick one port from each port group to be a LAN port.

Tunnels and Hybrid Mode

A tunnel on a VE_Port must be configured to enable IP Extension. A tunnel that supports IP traffic provides additional QoS priorities for IP Extension. The tunnel can carry both FC and IP traffic. When a tunnel is running in FC-only mode, it is compatible with a Brocade 7840 Switch or Brocade SX6 Blade running in FCIP mode (that is, not in Hybrid mode). FC traffic will be limited to 10Gb/s. The Brocade 7810 Switch provides only Hybrid mode and FC traffic is limited to 10Gb/s.

Compression is supported on a tunnel in Hybrid mode. With FC traffic, all compression modes are supported: fast deflate, aggressive deflate, and deflate compression. IP traffic is limited to two compression modes: aggressive deflate and deflate compression.

NOTE
The Brocade 7810 Switch supports only aggressive deflate and deflate for FC and IP traffic.

Out-of-Order Delivery on a Tunnel

Head of line blocking (HoLB) is mitigated for the extended IP flows through the tunnel. Each flow in a TCP connection receives an independent stream in the tunnel. The data must be delivered in-order for the stream; however, data can be delivered out-of-order for the tunnel. This allows the WAN to pass up any data that is received out of order because packet loss recovery is isolated on the WAN to the impacted stream, or connection, that the lost data belongs to.

IP Extension and Traffic Control Lists

A traffic control list (TCL) defines how LAN traffic is mapped to specific tunnels. When you create a TCL, you create a rule that identifies specific characteristics of the LAN traffic. Examples of these characteristics, include but are not limited to IP addresses, layer 4 protocols and ports, and VLAN tags. Each rule functions as an input filter for the DP and can either allow or deny a specific traffic flow from being passed through the IP extension tunnel. Multiple TCL rules, arranged by priority, provide a high level of control over the LAN traffic flow through a particular DP.

NOTE
Limits on the maximum-allowed TCLs differ among extension platforms. Brocade 7840 Switch and Brocade SX6 Blade provide 1024 Defined and 128 per DP (total of 256) Active TCLs where as the Brocade 7810 Switch provide 256 Defined and 32 Active TCLs.

NOTE

TCL rules must be configured. One default rule exists, which is to deny all traffic. This default rule cannot be removed or modified. It is the lowest priority rule, 65535, so it will be the last rule enforced. To have traffic over an IP extension tunnel, you must configure one or more rules that allow traffic to pass through.

Each TCL rule is identified by a name assigned to the rule when it is created. Thereafter, rules are modified or deleted by name. A TCL name contains 31 or fewer alphanumeric characters. Each name must be unique within the IP extension platform (such as Brocade 7810 Switch, a Brocade 7840 Switch, or a Brocade SX6 Blade). Rules are local to each platform and are not shared across platforms.

When traffic is allowed, the TCL rule specifies which tunnel and which QoS to use for that traffic type. When the rule denies traffic, it applies to both DP complexes unless a specific DP complex is selected.

The TCL priority number provides an order of precedence to the TCL rule within the overall TCL list. The priority value must be a unique integer. Smaller numbers are higher priority and larger numbers are lower priority. You can modify the priority to reposition the TCL rule to a different location within the list. When removing a rule, or even when creating a new rule, you can leave gaps in the priority numbers to allow subsequent in-between entries.

NOTE

The priority value must be unique across all active TCL rules within an IP extension platform. For example, if a chassis has multiple Brocade SX6 blades installed, the priority must be unique across all blades. If a TCL is defined as priority 10, that same priority cannot be used for another TCL rule, even if that rule would be assigned to another DP. A check is performed when the rule is enabled to ensure the priority value is unique.

The TCL input filter inspects a set of parameters to help identify the input traffic. It is based on several of the fields found in the IP, TCP and other protocol headers. The TCL input filter identifies a particular host, device, or application by means of the Layer 4 protocol encapsulated within IP, Layer 4 destination ports, Layer 3 source or destination IP addresses and subnets, DSCP/802.1P QoS values, or VLAN tagging.

When defining a TCL rule, the TCL action and TCL target will determine the behavior with regard to the DPs. For the TCL action, the following actions are possible:

- When the action is set to “allow,” the target must be an IP Extension-enabled VE_Port tunnel.
- TCL rules consist of either two or three main parts depending on if it is an “allow” or “deny” action.

Filter parameter	Action	Action
Priority number	Allow	Deny
Filter definition	Allow	Deny
Tunnel target	Allow	N/A

- The “allow” priority rules are activated on the DP that is associated with the selected VE_Port target.

DP complex	Brocade 7840 Switch Native VE_Ports 10VE mode	Brocade SX6 Blade Native VE_Ports 10VE mode
DP0	24–28	16–20
DP1	34–38	26–30

NOTE

For Brocade 7810 Switch, DP0 comprises ports 12-15. DP1 is not applicable.

- When the action is set to “deny” the rule is pushed to both DPs to deny matching traffic. 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 both DPs, thus accounting for traffic that may arrive at either DP. Traffic not specifically destined for one of the unicast LAN IPIF addresses cannot be internally forwarded to a specific DP LAN IPIF. For example, BUM (Broadcast, Unknown unicast and Multicast) traffic has no specific known destination, which means the specific DP LAN IPIF to forward the traffic to is unknown. Because the LAN ports cannot determine which IPIF to forward such traffic to, both DPs receive the TCL “deny” rule. If no particular LAN IPIF can be determined, both DPs must be capable of handling this type of traffic.
- If you must select a specific DP for a deny rule, you can configure the deny rule for that specific DP. The rule is pushed to the specified DP, denying any matching traffic on that DP only.

Once traffic is matched to a TCL rule, it is sent to the tunnel, and further TCL processing stops for that traffic. The TCL target parameter specifies which VE_Port tunnel the matched traffic will be sent to. Optionally, you can specify a traffic priority. 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 for allow and deny actions one time only when a TCP stream performs its handshake to form the traffic stream. If you make any changes to the TCL, including changes to priority rules or disabled rules, the particular stream those changes apply to will have no effect on the existing streams. There is no effect on an existing stream because that stream is already formed and the TCL is no longer being evaluated. If you do not see any traffic changes after changing a TCL, it is because the traffic stream is 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 TCP stream, forming a new stream. 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.

A recommended approach is to configure the fewest number of “allow” rules that will pass specific traffic. Unmatched traffic will encounter the default “deny-all” rule. You can prevent a subset of the allowed traffic by an “allow” rule with a “deny” 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 for the TCL to process. However, you can configure specific “deny” rules to ensure that certain devices or addresses cannot communicate across the IP extension tunnel. Always limit how you use “deny” rules. Configuring (and trouble-shooting) complex sets of “deny” rules results in more work and more chance of error.

In some situations, all traffic that appears at the IP Extension LAN interface is intended to pass through the tunnel:

- 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

In these cases, the simplest solution is to configure an “allow” rule that allows all traffic to pass through.

A default TCL is always created when the IP extension platform is configured for Hybrid mode. It has priority 65535, which is the lowest possible, and is set to deny all traffic. This rule cannot be deleted or modified. If no other rule is created, all traffic will be dropped. Therefore, it is critical that you create at least one TCL rule and set it to target an IP Extension enabled VE tunnel.

NOTE

When using Virtual Fabric Logical Switches (VF LS), it is important to know that the IP extension LAN Ethernet interfaces must be in the default switch. IP extension LAN Ethernet interfaces cannot be separated and assigned to different logical switches. The lan.dp0 and lan.dp1 IPIFs behind the IP extension LAN Ethernet interfaces are not part of VF LS contexts and cannot be assigned to an LS. If a datagram arrives at a IP extension LAN IPIF, it will be processed and is subject to the TCL.
Even though a VE_Port can reside within a VF LS, there is no requirement for the IP extension LAN Ethernet interfaces to also reside within that VF LS. In fact, the IP extension LAN Ethernet interfaces must remain in the default switch. The TCL directs traffic to the specified VE_Port regardless of the VF LS in which it resides. The VE_Port must be native to the DP hosting the lan.dp0 or lan.dp1. If the VE specified in the TCL is not native to DP hosting the LAN IPIF, the TCL will not find a match and the traffic will be dropped.

Matching IP Extension Traffic to Multiple VE_Ports

Multiple tunnels are typically used to go to multiple data centers. Normally, only one tunnel is used per data center between two IP extension platforms (such as a pair of Brocade 7840 Switches or a pair of Brocade SX6 Blades). A single tunnel can leverage Extension Trunking for aggregated bandwidth using multiple circuits, failover and failback, lossless link loss (LLL), and other benefits.

You can configure TCL rules that allow traffic to go to a specific tunnel when more than one tunnel exists.

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 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 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 figure). Priority rule 20 is evaluated next. That rule allows IP address 10.10.10.0 with 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 default priority rule 65535. The final rule, which cannot be altered or removed, is to deny all traffic. Any traffic processed by this final priority is dropped.

FIGURE 12 TCL Rules and Multiple Tunnels

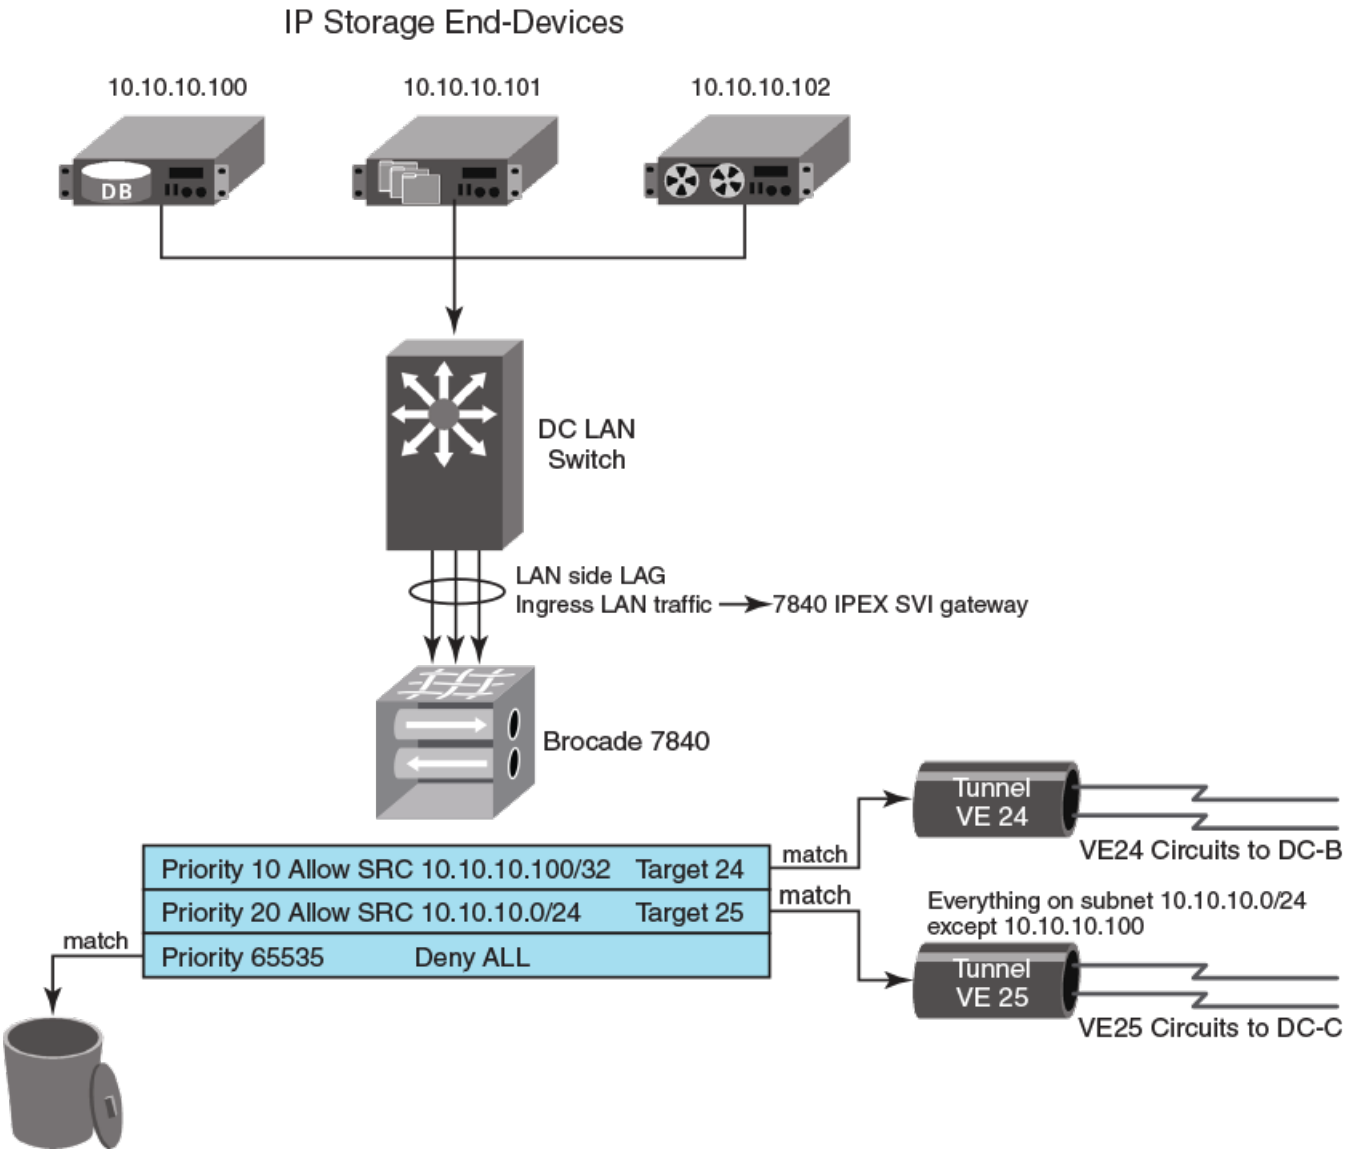

You can use the `portshow tcl` command to display the configured TCL rules.

```
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 TCL

You can create a TCL rule for traffic that does not terminate at the Brocade 7840 Switch, Brocade 7810 Switch, or the Brocade SX6 Blade. The non-terminated TCL option allows TCP traffic to be sent as-is to the other endpoint over a tunnel. This traffic is unlike terminated TCP traffic in that it does not count as one of the allowed 512 TCP connections at the DP for the Brocade 7840 Switch and the Brocade SX6 Blade or the allowed 128 TCP connections for the Brocade 7810 Switch.

The non-terminated TCL option is meant for use by low bandwidth control-type connections. It is recommended that you use the highest TCL priority value possible for the non-terminated traffic based the traffic handling requirements. This option can be enabled or disabled. However, you must stop and restart the traffic from the application or host to ensure proper handling after the TCL is updated.

IP Extension and QoS

Three fabric QoS priorities are supported for FC, which are high, medium, and low. You can configure a custom QoS distribution or use the default distribution of 50 percent for high, 30 percent for medium, and 20 percent for low.

For LAN traffic through IP Extension, three additional QoS priorities are created and default distributions are provided. The three priorities are IP-High, IP-Medium, and IP-Low. The QoS is added when a tunnel is enabled for IP Extension. Again, you can change the QoS configuration or use the default.

With IP Extension introduced in Fabric OS 7.4.0, a second level of QoS is introduced, called QoS groups. The IP group is for IP Extension traffic. The FC group is for Fibre Channel (or FCIP) traffic. This allows you to prioritize your traffic between FC and IP as needed. For example, you can specify a QoS distribution group ratio of 60 percent FC and 40 percent IP. The default group distribution is 50 percent FC and 50 percent IP. In other words, QoS distributions are specified separately for FC traffic and IP traffic when the Brocade 7840 Switch or Brocade SX6 Blade are operating in Hybrid mode.

> NOTE
> Minimum allocation for a single QoS type (high, medium, or low) should be 10 percent. QoS allocations within a group must total 100 percent. In addition, allocation for either FC or IP cannot exceed 90 percent.

IP Extension and Compression

The Brocade 7840 Switch and the Brocade SX6 Blade support four compression levels:

- None
- Fast deflate (not supported for IP traffic in Hybrid mode)
- Deflate
- Aggressive deflate

The Brocade 7810 Switch supports three compression levels.

- None
- Deflate
- Aggressive deflate

Compression can be configured at the QoS group / protocol level when the switch or blade are operating in Hybrid mode. Different compressions can be applied to FC extension traffic and IP extension traffic on the same tunnel. With the protocol selection, you can use fast deflate for FC traffic while the IP traffic is using deflate or aggressive deflate compression.

Compression is configured at the tunnel level, for all traffic on FC and IP protocols, but you can override the tunnel settings and select different compression at the QoS group / protocol level.

IP Extension and IP LAN Deployment

IP Extension is supported in policy-based routing (PBR) topologies and topologies with directly-connected devices on the Brocade 7840 Switch, the Brocade 7810 Switch, and the Brocade SX6 Blade. This allows you to connect IP storage to the LAN ports directly or thru Layer 2 switches.

Link aggregation groups (LAGs) are supported between the LAN ports and a Layer 2 Switch.

In Fabric OS 7.4.1 and subsequent releases, NIC teaming or other similar methods are not supported.

In Fabric OS 8.0.1 and subsequent releases, policy-based routing redirection is supported for LAN interfaces.

When using IP storage arrays, the IP storage can be connected either to a Layer 2 switch or directly to a Brocade 7840 Switch, Brocade 7810 Switch, or Brocade SX6 Blade.

In the following figures we use the Brocade 7840 Switch to illustrate LAN deployment as direct connect and as a Layer 2 switch connect.

FIGURE 13 IP Extension Direct Connect to IP Storage

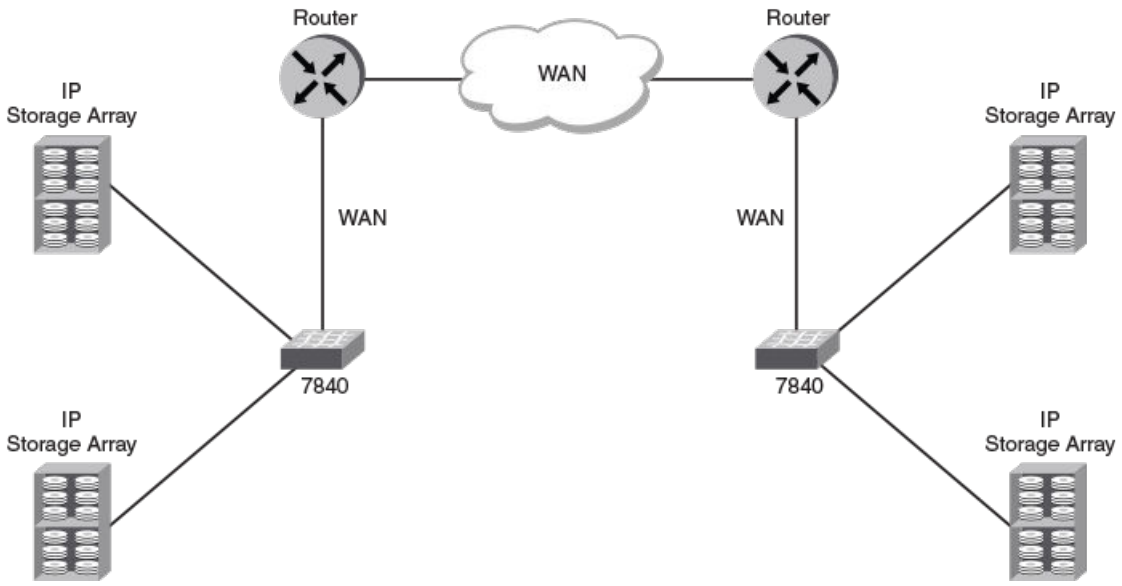

As a guideline, configure the IP storage array with the SVI on one of the DP complexes as the next hop gateway. Based on the next hop configuration, the storage device will learn the MAC address of the Brocade 7840 Switch SVI IP address through an ARP or Neighbor Discovery (ND) protocol request.

When IP storage devices are connected to a Layer 2 switch, as shown in the following figure, you can use LAGs to connect to the Brocade 7840 Switch LAN ports. The maximum number of ports in a LAG group is four. The maximum number of LAG groups is eight. Make sure that there is only one path between a single Layer 2 domain and the Brocade 7840 Switch.

NOTE

These maximum values apply to the Brocade SX6 Blade as well. For the Brocade 7810 Switch, however, the maximum values are 2 and 2, respectively.

NOTE
Beginning with the Fabric OS 7.4.0, static LAG is supported.

FIGURE 14 IP Extension Direct Connect to Layer 2 Switch

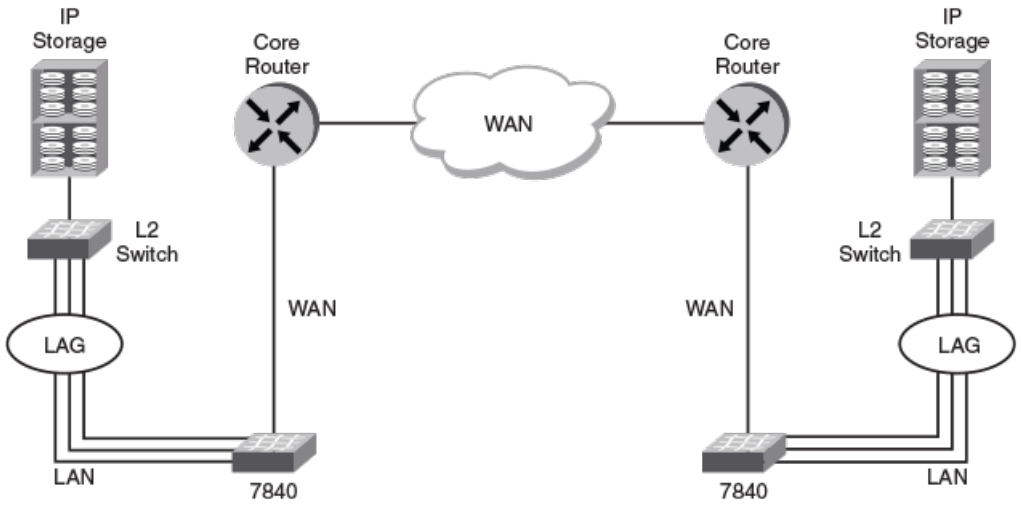

As shown in the figure, the Brocade 7840 Switch ports are connected to the IP storage array by means of the Layer 2 switch, and the WAN ports are connected to the WAN gateway. You must configure at least one SVI LAN IP address for each DP complex that is used. The router can be used as the WAN gateway, but it cannot be used on the LAN unless it is acting as a Layer 2 device.

The following figure shows IP extension used with a router that is configured with PBR to direct specific traffic to the Brocade 7840 Switch. This allows a data center router to support WAN traffic and to take advantage of the IP extension features for IP storage arrays. The PBR configuration simplifies the IP storage array and Brocade 7840 Blade connectivity. With the PBR configuration, the IP storage array and a Brocade 7840 Switch SVI IP address can be in different subnets. Configuration in a data center can keep existing IP address and router configurations. The data center router configuration is modified to add policy-based routing to send specific traffic to the Brocade 7840 Switch instead of the WAN core router.

FIGURE 15 IP Extension PBR Connect

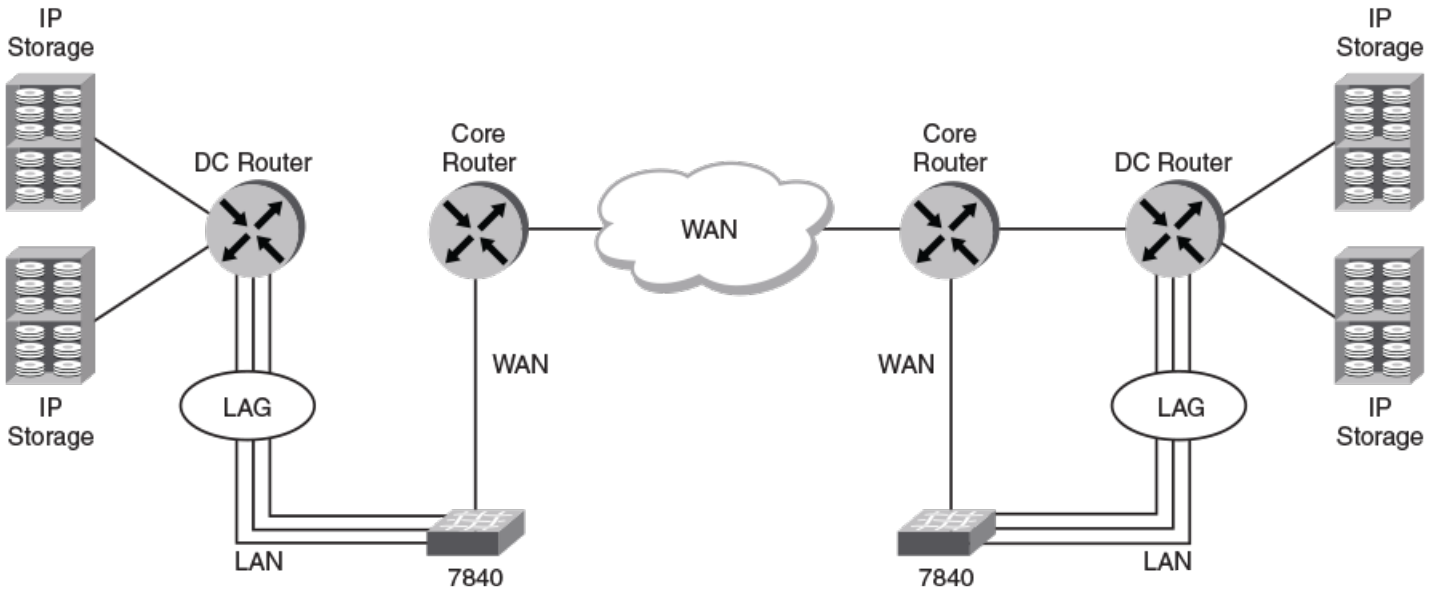

IP Extension Limitations and Considerations

For TCP traffic, the following considerations apply:

- Flow control for the LAN TCP sessions is handled by means of internal management of the receive window. The receive window allows up to 256 KB of data onto the LAN. However, the window will degrade to 0 if the WAN becomes congested or the target device experiences slow draining.
- TCP connections are limited to 512 terminated or accelerated connections per DP complex. Additional connection requests are dropped. A RASlog message is displayed when 500 connections are up and active.

  NOTE
  The Brocade 7810 Switch only supports 128 connections. A RASlog message is displayed when 116 connections are up and active.

- For the Brocade 7840 Switch and the Brocade SX6 Blade, up to 512 TCP open connection requests per second are allowed. For the Brocade 7810 Switch, up to 128 TCP open connection requests per second are allowed. Additional open connection requests are dropped. By limiting the number of connection requests per second, this helps prevent Denial of Service (DoS) attacks.
- Statistics are available for the number of dropped connections.

Extension Platform and L2 Protocols

Trunking on LAN Ports Using LACP

IP Extension supports up to eight IP addresses per DP complex for LAN traffic. IPv4 and IPv6 traffic are supported. Each DP complex keeps its own LAN interface as well as its own MAC for LAN traffic. In this manner, the DP complex acts as a switch virtual interface (SVI) for the LAN.

All GbE ports that are configured as LAN ports can access the SVI addresses of each DP complex. This allows for multiple GbE ports to access a single IP gateway. In addition, link aggregation groups (LAGs) are supported. A single LAG can contain up to four ports. A total of eight LAGs are supported. The eight supported LAGs can be any combination of static and dynamic.

TABLE 23 Maximum Supported LAGs for Default Switch Configuration

Platform	Max Static LAGs	Max Dynamic Lags	Max LAGs Supported (Hard Limit)	Max Number of GbE Ports per LAG	Max LAGs per Brocade SX6 *slot*	Max LAN Ports per Brocade Blade/Brocade Switch	Max Brocade SX6 Blades per Chassis
Brocade 7840	8	8	8 (combination of static and dynamic)	4	N/A	8	N/A
Brocade 7810	2	2	2	2	N/A	4	N/A
Brocade X6 Director	32	96 (0 to 32) on Brocade SX6 Extension Blade) (64 to 96) on Brocade FC32-64 Port Blade)	96 (combination of static and dynamic)	4	8	8	4

NOTE
In the Fabric OS 7.4.0 release through Fabric OS 8.1.0, only static LAGs are supported. In Fabric OS 8.2.0 and later, both dynamic and static LAGs are supported.

Other considerations for LAN and WAN ports:

- Link Level Discovery Protocol (LLDP) is supported with Fabric OS 8.2.0 and later. With LLDP, network devices advertise their identity, capabilities, and configuration on an IEEE 802 LAN.
- Jumbo frames are supported for LAN traffic. A jumbo frame can be up to 9216 bytes.
- VLAN tagging is supported in IP Extension. Stacked tagging (IEE802.1ad) is not supported.

Neighbour Discovery on GbE Ports using LLDP

LLDP (Link Level Discovery Protocol) is a Layer 2 protocol used by devices to advertise their identities, capabilities and configurations to directly connected peers and to learn and store information about the peers. It is a vendor-neutral protocol and formally referred to as IEEE standard 802.1AB-2009. LLDP is supported in both FCIP and Hybrid modes on the GbE ports on the following extension platforms:

- Brocade 7840 Extension Switch
- Brocade 7810 Switch
- Brocade X6 with the SX6 Extension Blade

The KAP Support for LACP and LLDP

In Fabric OS 8.2.0 and later, the control processor (CP) uses a keepalive procedure to control transmission of keepalive packets (KAP) on a per-port basis for LACP and LLDP.

The following table shows the timeout values and transmission intervals for LACP and LLDP protocols. The timeout values are the target recovery time from Extension Hot Code Load (eHCL) events to ensure non-disruptive upgrades. Keepalive support is designed to minimize the time during which no KAP are sent.

The following table shows the timeout values and transmission intervals for LACP and LLDP protocols.

TABLE 24 LACP and LLDP Timeout Values

Protocol	Timeout Counter	Hello Interval (minimum)	Minimum Timeout	Hello Interval (maximum)	Maximum Timeout	Default Hello Interval
LACP	3	1 sec (short)	3 sec	30 sec (long)	90 sec	30 sec (long)
LLDP	4 (2 to 10)	4 sec	8 sec	180 sec	1800 sec	30 sec

NOTE

Although the KAP timeout interval and count for the Brocade 7810 matches that for Brocade 7840 or the Brocade SX6 blade, the KAP on the Brocade 7810 cannot send keep-alive frames through an eHCL. So, the Brocade 7810 cannot support HA for IP traffic on LAG.

Upgrade and Downgrade Considerations for LAG and LLDP

For Fabric OS 8.2.0 and later, there are considerations for firmware upgrade and downgrade.

NOTE

The Brocade 7810 Extension Switch requires FOS 8.2.1 or later.

Link Aggregation Group (LAG) Considerations

When upgrading to Fabric OS 8.2.0 and later from an earlier release, consider the following:

- Firmware migration from an earlier release to Fabric OS 8.2.0 is non-disruptive. Static LAG configurations are migrated and configuration information is transformed to the Fabric OS 8.2.0 CLI. The `portcfg lag` and `portshow lag` commands are deprecated from Fabric OS 8.2.0.

When downgrading from Fabric OS 8.2.0 or later to an earlier release, consider the following:

- Firmware migration to an earlier release is non-disruptive.
- Firmware migration to an earlier release is not allowed if the Fabric OS 8.2.0 contains dynamic LAG configurations. All dynamic LAG configurations must be removed before downgrade can occur.
- It is not recommended to download a configuration file for Fabric OS 8.2.0 firmware to a switch running a lower version firmware, for example Fabric OS 8.1.0.
- When a valid configuration file is downloaded, reboot is mandatory for the new configuration file to be applied.
- Before downloading a configuration file, make sure you enter the `switchdisable` command. A configuration file should be downloaded only when the patform is in `switchdisable` mode.

Link Level Discovery Protocol (LLDP) Considerations

When upgrading to Fabric OS 8.2.0 and later from an earlier release, consider the following:

- Firmware migration from an earlier release to Fabric OS 8.2.0 is non-disruptive.
- After migrating to Fabric OS 8.2.0, LLDP functionality is available.

When downgrading from Fabric OS 8.2.0 and later to an earlier release, consider the following:

- Firmware migration to an earlier release is non-disruptive.

- Firmware migration to an earlier release is not allowed if the Fabric OS 8.2.0 contains non-default or user defined LLDP configurations. All LLDP configurations must be removed before downgrade can occur.
- It is not recommended to download a configuration file for Fabric OS 8.2.0 firmware to a switch running a lower version firmware, for example Fabric OS 8.1.1.
- When a valid configuration file is downloaded, reboot is mandatory for the new configuration file to be applied.
- Before downloading a configuration file, make sure you enter the `switchdisable` command. A configuration file should be downloaded only when the patform is in `switchdisable` mode.

Extension Hot Code Load for the Brocade 7840 and the Brocade SX6

The Brocade 7840 Switch and the Brocade SX6 Blade support Extension Hot Code Load (eHCL). (The Brocade 7810 Switch does not.) Extension HCL allows non disruptive firmware updates over the extension tunnels. eHCL benefits mainframe environments by supporting nonstop applications such as data replication and tape backups. Extension HCL maintains device to mainframe connectivity while the firmware downloads, without disrupting active I/O. eHCL operates in both FCIP mode and Hybrid mode. For information about eHCL operation, see Extension HCL Operation on page 23. Refer to the *Configuring Extension Features* chapter for configuration information.

Path MTU Discovery

Path maximum transmission unit (PMTU) discovery is supported on the Brocade 7810 Extension Switch, the Brocade 7840 Extension Switch, and Brocade SX6 Extension Blade. PMTU is the process of sending Internet Control Message Protocol (ICMP) datagrams of various known sizes across an IP network to determine the supported maximum datagram size.

Based on the largest ICMP Echo Reply datagram received, the PMTU discovery process sets the IP MTU for that circuit's IP interface (ipif). Each circuit initiates the PMTU discovery process prior to coming online. This is required because the circuit may have gone offline due to a link failure, rerouted to a new path, and now has a different MTU. If a circuit bounces, the PMTU discovery process will be initiated when attempting to re-establish the circuit. The PMTU discovery process can result in more time for the circuit establishment. The smallest supported MTU size is 1280 bytes. On the Brocade 7840 Switch, Brocade 7810 Switch, and the Brocade SX6 Blade, the largest supported IP MTU size is 9216 bytes; however, PMTU discovery will not discover an MTU greater than 9100 bytes. If the IP network's MTU is known, the best practice is to set it manually in the `portcfg ipif` command. This will avoid values determined by PMTU discovery that are less than the exact MTU of the IP network.

PMTU requires that ICMP is permitted across all IP network devices and the WAN. A rudimentary check would be if you could ping devices across this network. Brocade PMTU discovery uses ICMP Echo Requests. If there are no firewalls most likely ICMP is free to traverse the network. If PMTU discovery cannot communicate with the peer switch, the circuit will not be established.

Enable 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 show the configuration of the MTU parameter and `portshow fcipcircuit --detail` command to display the actual discovered PMTU value being used. You can also initiate PMTU discovery using the `portcmd --pmtu` command.

Circuit Failover

Circuit failover provides an orderly means to recover from circuit failures. To manage failover, each circuit is assigned a metric value, either 0 or 1, which is used in managing failover from one circuit to another.

Trunking with metrics uses lossless link loss (LLL), and no in-flight data is lost during the failover. If a circuit fails, Trunking first tries to retransmit any pending send traffic over another lowest metric circuit. In the following figure, circuit 1 and circuit 2 are both lowest metric circuits. Circuit 1 has failed, and transmission fails over to circuit 2, which has the same metric. Traffic that was pending at the time of failure is retransmitted over circuit 2. In-order delivery is ensured by the receiving extension switch or blade.

FIGURE 16 Link Loss and Retransmission over Peer Lowest Metric Circuit

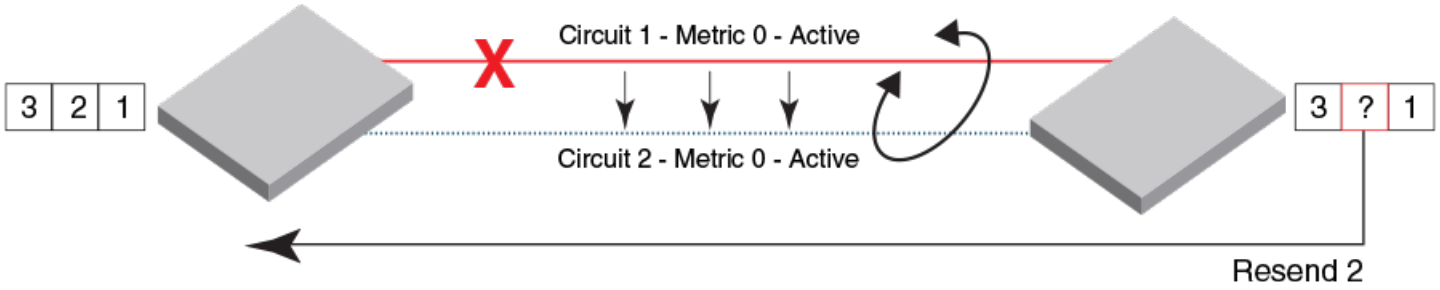

NOTE
Modifying a circuit metric disrupts traffic.

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. In this case, circuit 2 is not used until no lowest metric circuits are available. If all lowest metric circuits fail, then the pending send traffic is retransmitted over any available circuits with the higher metric. Failover between like metric circuits or between different metric circuits is lossless.

FIGURE 17 Failover to a Higher Metric Standby Circuit

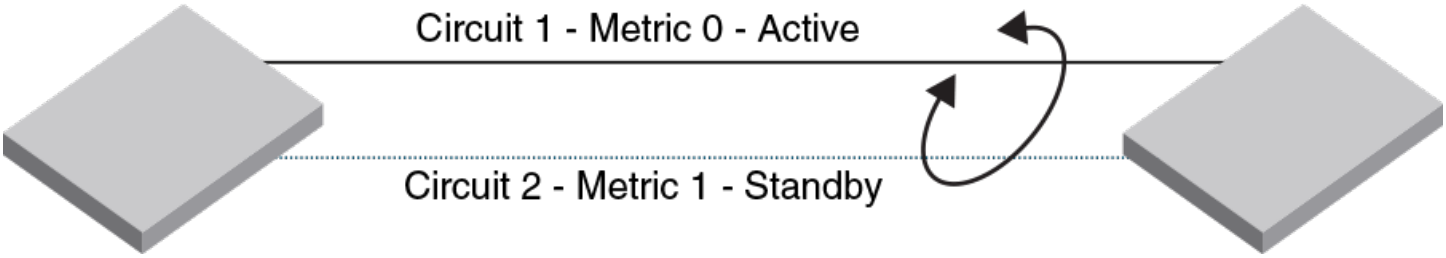

Only when all metric 0 circuits fail do available metric 1 circuits cover data transfer. If the metric 1 circuits are not identical in configuration to the metric 0 circuits, then the metric 1 circuits will exhibit a different behavior. Additionally, if the metric 1 WAN path has different characteristics, these characteristics define the behavior across the metric 1 circuits. Consider configuring circuit failover groups to avoid this problem.

Circuit Failover Grouping

With circuit failover groups, you can better control which metric 1 circuits will be activated if a metric 0 circuit fails. To create circuit failover groups, you define a set of metric 0 and metric 1 circuits that are part of the same failover group. When all metric 0 circuits in the group fail, metric 1 circuits will take over data transfer, even if there are metric 0 circuits still active in other failover groups.

Typically, you would only define one metric 0 circuit in the group so that a specific metric 1 circuit will take over data transfer when the metric 0 circuit fails. This configuration prevents the problem of the tunnel operating in a degraded mode, with fewer than the defined circuits, before multiple metric 0 circuits fail.

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 group must fail before the metric 1 circuits in that failover group are used.
- All metric 1 circuits in a group are used if all metric 0 circuits in the group fail or there is no metric 0 circuit in the group.
- Circuits can be part of only one failover group
- Circuit failover groups are only supported by Fabric OS 7.2.0 or later.
- Both ends of the tunnel must have the same circuit failover groups defined.
- Tunnel and circuit states will indicate a misconfiguration error if circuit failover group configurations are not valid.
- Modifying of the failover group ID is a disruptive operation, similar to modifying the metric.
- Circuit failover groups are not used to define load balancing over metric 0 circuits (*only* failover rules).
- When no circuit failover groups are defined, failover reverts to the default operation: all metric 0 circuits must fail before failing over to metric 1 circuits.
- 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 valid failover groups defined per tunnel is limited by the number of circuits that you can create for the switch model as follows:
  - For a Brocade 7840 Switch, you can configure up to 5 valid groups on a 10-circuit tunnel.
  - For a Brocade 7810 Switch, you can configure up to 3 valid groups on a 6-circuit tunnel.
  - For a Brocade SX6 Blade, you can configure up to 5 valid groups on a 10-circuit tunnel.
  - For an Brocade FX8-24 Blade, you can configure up to 5 valid groups on a 10-circuit tunnel.
- Consider available WAN bandwidth requirements when configuring failover circuit groups. See Bandwidth Calculation during Failover on page 80.

Examples of Circuit Failover in Groups

The following table illustrates circuit failover in a tunnel with two failover groups, each with two circuits. All data through the tunnel is initially load-balanced over circuits 1 and 2. The following occurs during circuit failover:

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

TABLE 25 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	1,500	If active, yes.
Circuit 2 Metric 0	2	1000Mb	1,000	If active, yes.
Circuit 3 Metric 1	1	500Mb	1,500	Only when circuit 1 fails.
Circuit 4 Metric 1	2	1000Mb	1,000	Only when circuit 2 fails.

The following table illustrates circuit failover in a tunnel with circuits in failover groups and circuits that are not part of failover groups. In this configuration, all data is initially load-balanced over circuit 1, circuit 2, and circuit 3 (when they are all active). The following occurs during circuit failover:

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

  Reason: Circuit 1 fails over to circuit 4 (both are in failover group 1) and circuits 2 and 3 are active with 500Mb bandwidth.
- If circuit 2 fails, data is load-balanced over circuit 1 and circuit 3, and no other circuit becomes active.

  Reason: Circuits 1 and 3 are the only active circuits, because circuits 4 and 5 only become active when circuits 1 or 3 fail.
- If circuit 2 and circuit 3 fail, circuit 5 becomes active and data is load-balanced over circuit 1 and circuit 5.

  Reason: Ungrouped circuits 2 and 3 fail over to ungrouped circuit 5, which has a metric of 0.
- If circuit 1, circuit 2, and circuit 3 fail, circuit 4 and circuit 5 become active and data is load-balanced over both.

  Reason: Circuit 1 fails over to circuit 4, which is the failover circuit for group 1 with a metric of 0. Ungrouped circuit 5 is the failover circuit for ungrouped, failed circuits 2 and 3.

TABLE 26 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	1,500	If active, yes.
Circuit 2 Metric 0	Not defined.	500Mb	1,500	If active, yes.
Circuit 3 Metric 0	Not defined.	500Mb	1,500	If active, yes.
Circuit 4 Metric 1	1	500Mb	1,500	Only when circuit 1 fails.
Circuit 5 Metric 1	Not defined.	1000Mb	1,000	Only when circuits 2 and 3 fail.

Bandwidth Calculation during Failover

The bandwidth of higher metric circuits is not calculated as available bandwidth on a tunnel until all lowest metric circuits have failed.

Assume the following 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, and circuit 1 is created with a maximum transmission rate of 500Mb/s. Together, circuits 0 and 1 provide an available bandwidth of 1.5 Gb/s.
- Circuits 2 and 3 are created with a metric of 1. Both are created with a maximum transmission rate of 1Gb/s, for a total of 2Gb/s. This bandwidth is held in reserve.

The following actions occur during circuit failures:

- If either circuit 0 or circuit 1 fails, traffic flows over the remaining circuit while the failed circuit is being recovered. The available bandwidth is still considered to be 1.5Gb/s.
- If both circuit 0 and circuit 1 fail, there is a failover to circuits 2 and 3, and the available bandwidth is updated as 2Gb/s.

- If a low metric circuit becomes available again, the high metric circuits return to standby status, and the available bandwidth is updated again as each circuit comes online. For example, if circuit 0 is recovered, the available bandwidth is updated as 1 Gb/s. If circuit 1 is also recovered, the available bandwidth is updated as 1.5 Gb/s.

10-GbE Lossless Link Loss (FX8-24 Blade)

Circuit failover, or lossless link loss (LLL), is supported between 10GbE circuits on FX8-24 blades when both 10GbE ports are on the same logical switch and are operating in 10Gb/s mode. You can configure higher metric circuits for failover from lower metric circuits (see Circuit Failover on page 78). You can also configure IP addresses for a failover crossport. Crossports are IP addresses (and routes) that belong to the other 10GbE port's VE group. The crossport for xge0 is xge1 and the crossport for xge1 is xge0. For more information on crossports and configuring crossports, see Crossports and Failover on page 59.

LLL is supported per VE_Port on the VE_Port's DP complex. Because a VE_Port cannot span GbE and 10GbE interfaces, neither can LLL. LLL is supported on both GbE and 10GbE interfaces, just not together.

Benefits and limitations of 10GbE lossless link loss (LLL) failover include the following:

- LLL provides failover to protect against link or network failure and 10GbE port disable.
- Data will not be lost due to failover.
- Failover supports active-passive and active-active configurations.
- Dual mode is not supported for 10GbE port failover.
- Failover does not protect against failure of a DP complex.
- Disabling a VE_Port will not use LLL. In this case, route failover will occur at the FC level based on APT policy, if there is another route available, and may cause loss of FC frames.

  NOTE
  All circuits and data must belong to a single VE_Port to benefit from LLL.

Circuit Spillover

Circuit spillover is a load-balancing method introduced in Fabric OS 8.0.1 that allows you to define circuits that are used only when there is congestion or a period of high bandwidth utilization. Circuit spillover is configured on QoS tunnels on a VE_Port.

When using spillover load balancing, you specify a set of primary circuits to use all the time, and a set of secondary circuits (spillover circuits) that are used during high utilization or congestion periods. Primary and secondary circuits are controlled using the metric field of the circuit configuration. For example, when a tunnel is configured for spillover, traffic uses the metric 0 circuits until bandwidth utilization is reached on those circuits. When bandwidth in metric 0 circuits is exceeded, the metric 1 circuits are used. Conversely, when the utilization drops, the metric 1 spillover circuits are no longer used.

The failover behavior of a tunnel remains intact with regard to the lower metric circuits. If a tunnel is configured for spillover and the lower metric circuits become unavailable, the higher metric circuits function as the backup circuits. Metric 1 circuits always behave as backup circuits whether configured for spillover or failover.

Failover load balancing, in comparison to spillover, requires that all lower metric circuits become unavailable before the higher metric circuits are used. When you have multiple lower metric circuits, this means that all of them must fail before any of the higher metric circuits are used.

Circuit failover groups can still be configured and will behave as before, however, only the metric value determines whether a circuit is considered a spillover circuit or a primary circuit. For example, if a tunnel is configured with four circuits, with two metric 0 and two metric 1 circuits in two separate failover groups, both metric 0 circuits are used, and only when they become saturated are the metric 1 circuits used. The failover grouping is used for failover scenarios only.

Because of how spillover is implemented, the observed behavior may not appear as expected. For example, consider QoS traffic flows of high, medium, and low with a bandwidth percentage assigned to each priority. QoS traffic is carried by the active circuits according to the percentage allocated to each priority. Say the QoS low priority traffic reaches saturation before hitting its bandwidth allocation limit, it will spill 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 circuit (metric 1) is used even though the metric 0 circuits are not saturated. When the metric 0 circuit is saturated, additional traffic will spillover to the metric 1 circuit.

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

> NOTE
> 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.

If one end of the spillover circuit is using a version of Fabric OS software prior to FOS 8.0.1, the mismatch is identified and the circuit behavior defaults to circuit failover.

Understanding Circuit Spillover Utilization

To verify 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 spillover usage of metric 1 circuits. Using and `--qos/-q` option provides more detail about which QoS tunnel the spillover is occurring on. Some commands provide PDU/byte counts, and others provide throughput information:

```
portshow fciptunnel -c
portshow fciptunnel -c -q
portshow fciptunnel 24 -c
portshow fcipcircuit 24 -c [-q]
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 is 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)
```

Some understanding of how resources are allocated will help understand the utilization numbers. A VE has multiple QoS tunnels, each of which can be 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 period for each individual 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. If utilization is present, it most likely indicates that some pushup back to the host is occurring.

The values displayed for the tunnel and circuit represent the highest level measured within that object.

The values displayed for the QoS tunnel level represent the range of these measurements, lowest and highest measured.

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

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

To accurately see why or why not throughput spillover is active, use and `--qos` option. By looking at the displayed utilization range, you can determine if the resources are being evenly used, or which QoS is being used or not. If the low-high utilization range is tight, it means that the resource groups are being evenly used for that QoS.

Examples

The following series of various `--perf` options shows the utilization reaching the spillover thresholds for one QoS and 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  1/-
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/-
-------------------------------------------------------------------------------
```

The next series shows the utilization not quite reaching the spillover thresholds, thus no spillover.

```
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

Many considerations for circuit spillover and circuit failover bandwidth and network requirements are similar, as follows:

- In the `portcfg fciptunnel` command, the value for *`max-comm-rate`* cannot exceed 5 times the value for *`min-comm-rate`*.
- The minimum committed rate value is 20000.
- There is no restriction on the bandwidth difference between circuits.
- The round-trip time must be less than 250 ms.
- Packet loss must be less than 1 percent.

The following are specific to circuit spillover:

- When the tunnel is configured as spillover, all metric 1 circuits are treated as metric 0 (primary) circuits when calculating the aggregate bandwidth restrictions.

Service-Level Agreement

A service-level agreement (SLA) works in conjunction with the existing WAN Tool features to provide automatic testing of a circuit before it is placed into service. SLAs are supported on the Brocade 7840 Switch, Brocade 7810 Switch, and the Brocade SX6 Blade.

The primary purpose of an SLA is to provide automated testing of a circuit before it is placed into service. The SLA checks the circuit for the packet loss percentage. If you need to verify the circuit for additional network performance, such as throughput, congestion, or out-of-order delivery, use WAN Tool to run tests manually. See Using WAN Tool on page 219 for information.

You must configure an SLA session at each end of the circuit being tested. The SLA session uses information from the circuit configuration to configure and establish the SLA connections. 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 minor mismatch. When the session is established, traffic starts automatically. For the duration of the test, the traffic must remain under the specified loss percentage before the circuit is placed into service. On the Brocade 7840 Switch and the Brocade SX6 Blade, up to 20 SLA sessions can be defined per DP. On the Brocade 7810 Switch, the limit is 12. See Configuring a Service-Level Agreement on page 108 for information about configuring and using an SLA.

> NOTE
> On the Brocade 7810 Switch, when you have SLA's 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 may exceed the 2.5Gb/s max of the platform. In this type of configuration, SLA use should be avoided or used with caution.

In addition to packet loss, the SLA can also test for timeout duration. If the timeout value is reached during the SLA session, the session is terminated and the circuit is put into service. A timeout value of "none" means that the test runs until the runtime and the packet-loss values are met.

Interaction with an SLA while it is running is limited. You can view statistics, and you can abort an active session. Any attempt to modify a session while it is active is blocked, which means the WAN Tool commands cannot be used while an SLA session is running.

Whenever a tunnel or circuit goes offline and comes back online, or when a circuit is administratively disabled then enabled, the SLA session is started and tests the link before allowing the circuit to go back into service. Configured SLA sessions are persistent across reboots, because circuit configurations are persistent across reboots and the SLA is part of the circuit configuration. However, user-configured WAN Tool sessions are not persistent.

After configuring an SLA, you assign the SLA to a specific circuit with the `portcfg fcipcircuit` command.

NOTE

During an eHCL reboot, the SLA is disabled and no new SLA sessions can be created until all eHCL operations are complete. After all eHCL operations are complete, the SLA is re-enabled.

Configuring Extension Features

Configuration Overview

Configuration consists of two main phases. The first phase is the planning and preparation phase. The second phase involves logging into the Brocade extension device and issuing commands that configure the features available on that device.

For the planning and preparation phase, see Configuration Prerequisites on page 88. You can use that information as a checklist for preparation.

When you configure the Brocade extension platform, the recommended sequence is as follows:

1. Configure the platform mode. Select the application modes and VE modes. On the Brocade FX8-24 Extension blade, configure the GbE port mode and VEX mode.
2. On the Brocade 7840 Extension switch, the Brocade 7810 Extension switch and the Brocade SX6 Extension blade, configure the GbE port speed on the 10GbE ports.
3. Configure the IP parameters. This includes the IP interface (IPIF), the IP route information where needed, and the VLAN.
4. Configure the service-level agreement (SLA) function.
5. Configure IPsec policies for the Brocade 7840 Switch, Brocade 7810 Switch, and the Brocade SX6 Blade. IPsec policies are not supported on the Brocade FX8-24 Blade. Instead, IPsec is configured when the circuit is created or modified.
6. Configure the extension tunnels for FCIP.
7. Configure the Extension Hot Code Load (eHCL) feature. eHCL is supported on the Brocade SX6 Blade and the Brocade 7840 Switch but not on the Brocade 7810 Switch.
8. Configure IP Extension if you are using this feature. IP Extension is supported on the Brocade 7840 Switch, Brocade 7810 Switch, and the Brocade SX6 Blade.
9. Finalize the configuration.

Configuration Prerequisites

Before you begin, review Tunnel and Circuit Requirements for Brocade Extension Platforms on page 62 then consider the following points.

- Determine the amount of bandwidth that is required for the extension features deployed in your network.
- Confirm that the WAN links are provisioned and tested for integrity.
- Confirm that the LAN links are provisioned.
- Make sure that cabling within the data center is completed.
- Make sure that switches and other devices are physically installed and powered on.
- Make sure that you have admin access to all switches and blades that you need to configure.
- For the Brocade FX8-24 Blade, determine which of the three possible GbE or XGE port operating modes will be used.
- For the FX8-24 Blade, determine which 10-GbE crossports to use for active-active or active-passive configurations.
- For the Brocade 7840 Switch and Brocade SX6 Blade, determine the VE_Port operating modes that will be used (10VE mode or 20VE mode).

  NOTE
  The Brocade 7810 Switch does not support the different VE operating modes.

- Determine which Ethernet ports will be used. The Ethernet ports on the Brocade 7840 Switch, the Brocade 7810 Switch, and the Brocade SX6 Blade are in groups, and connections should be spread across the groups and not within the same group if possible.
- Obtain subnets and assign IP addresses for each circuit endpoint that you intend to use, plus the netmask and IP MTU size. The IP MTU size may be smaller than 1500 if there is an IPsec device or similar device in the path. If the IP MTU is larger than 1500, use the following guidelines for your extension product:
  - For the Brocade 7840 Switch, the Brocade 7810 switch, and the Brocade SX6 blade, the IP MTU size must be at least 1280. If the supported maximum IP MTU size in the network is larger than 9216, the IP MTU must be 9216. You can use Path MTU Discovery to automatically set the IP MTU size for the circuit's IP interface.
- Determine the gateway IP address as needed for each route across the WAN. The gateway IP address will be on the same IP subnet as the subnet used for the IPIF interface that will use that gateway. The route will be the subnet and netmask on the remote side.
- Determine the VE_Port numbers that you want to use. The VE_Port numbers serve as tunnel IDs. Typically, the first one is used.
- Determine source and destination IP addresses for circuit 0 and the minimum and maximum rates for ARL. These values are set by the `portCfg fciptunnel create` command. If ARL is not being used, set the minimum and maximum committed rates to the same value.
- Determine how many additional circuits you want to create. You will need the source and destination IP addresses for each circuit and the minimum and maximum rates for ARL. You will need to know if you intend to assign metrics to circuits so that lower metric circuits fail over to circuits with higher metrics. For all circuits except circuit 0, these values are set by the `portCfg fcipcircuit` create command.
- If you are using IP Extension, make sure that you have the LAN information available for your site.
- When configuring tunnels to support large numbers of devices, consider the memory limitations of the extension switch or blade if you are enabling any type of emulation feature, such as FCP or FICON. If too many devices are present or activated at one time, acceleration operations can be negatively impacted. See Memory Use Limitations for Large-Device Tunnel Configurations on page 37.

Configuring Platform Modes

Platform mode configuration consists of setting modes that affect how the platform operates.

- FCIP mode (Brocade 7840 Switch and Brocade SX6 Blade) or Hybrid mode (Brocade 7840 Switch and Brocade SX6 Blade).

  NOTE
  The Brocade 7810 Switch does not support mode configuration. It is in Hybrid mode by default and this setting cannot be changed.

- VE mode. This mode determines how many VE ports are available as well as their total throughput. (Brocade 7840 Switch and Brocade SX6 Blade)
- 7810 GE port mode. This mode can be set to either copper or optical.
- SX6 blade. When you remove a blade or change the blade slot, platform information for the blade must be cleared.
- FX8-24 blade GbE mode. This mode can be configured to support various throughput.
- FX8-24 virtual EX (VEX) mode. This mode is supported only on the Brocade FX8-24 Blade.

Configuring FCIP or Hybrid Mode

The Brocade 7840 Switch and the Brocade SX6 Blade can assume both modes though the default is FCIP. The Brocade 7810 Switch operates solely in Hybrid mode. If you want to use IP Extension features, Hybrid mode must be configured on the switch or blade.

You can configure the Brocade 7840 Switch and Brocade SX6 Blade in either FCIP or Hybrid mode using the `extncfg --app-mode` command, which is disruptive and requires a reboot. The Brocade 7810 Switch only operates in Hybrid mode.

The following switch modes are available.

- With **FCIP mode**, only FCIP traffic is sent over the extension tunnels. Each DP allows up to 20G of FC traffic (up to 40G with 2:1 compression). You can select either 10VE or 20VE port modes.
- With **Hybrid mode**, both FCIP traffic and IP traffic can be sent over the extension tunnels.
  - Brocade 7840 Switch or Brocade SX6 Blade. For each DP, Hybrid mode enables up to 10G of FC traffic (up to 20G FC traffic with 2:1 compression) and up to 20G of IP traffic. In Hybrid mode, only 10VE port mode is available.
  - Brocade 7810 Switch. This switch is always in Hybrid mode. It does not support switching modes (FCIP, 10VE, and 20VE).

  NOTE
  If conflicting configurations exist when you are switching modes, the `extncfg` command will fail. For example, if you are in 20VE mode on a Brocade 7840 switch and have a tunnel configured on VE_Port 30, you cannot change to 10VE mode because VE_Port 30 is unavailable in 10VE mode.

1. Connect to the switch and log in using an account assigned to the admin role.
2. To set the switch to FCIP mode, use the `extncfg --app-mode fcip` command.

   ```
   switch:admin> extncfg --app-mode fcip
   ```

   FCIP is the default operating mode.
3. To set the switch or blade operating mode to hybrid, use the `extncfg --app-mode hybrid` command.

   ```
   switch:admin> extncfg --app-mode hybrid
   ```

   Changing operating mode between FCIP and hybrid is disruptive.

4. To display information about the current mode, use the `extncfg --show` command.

```
switch:admin> extncfg --show
```

Configuring VE Mode

The VE mode on a Brocade 7480 Switch or Brocade SX6 Blade controls how many VE_Ports are available.

In FCIP mode, you can configure the Brocade 7480 Switch or the Brocade SX6 Blade in either 10VE mode (default) or 20VE mode using the `extncfg --ve-mode 10VE||20VE` command. This command is disruptive; it requires a reboot.

NOTE
While the Brocade 7810 Switch does not support VE modes, it does provide between 12-15 VE ports.

You can configure the following VE modes on a Brocade 7480 Sswitch or Brocade SX6 Blade.

- **10VE mode**: In this mode, 10 of the 20 total VE_Ports on the switch are enabled. A single VE_Port on a DP complex can use all Fibre Channel 20Gb/s of bandwidth available to the complex.
- **20VE mode**: This mode is available when the switch is in FCIP mode. In this mode, all 20 VE_Ports are enabled. A single VE_Port on a DP complex can use half of the available Fibre Channel bandwidth available to the DP complex, a maximum of 10 Gb/s. This option allows use of more VE_ports, but at a lower maximum bandwidth.

NOTE
On the Brocade 7840 Switch or the Brocade SX6 Blade, you should configure only the maximum number of VE_Ports. 10VE mode will accommodate nearly all environments and is the default.

The following table lists the available ports in 10VE and 20VE mode.

TABLE 27 10VE and 20VE Mode Ports

Product	10VE mode VE_Ports (default)	20VE mode VE_Ports
Brocade 7840 Switch	DP0 - 24 through 28 DP1 - 34 through 38	DP0 - 24 through 33 DP1 - 34 through 43
Brocade SX6 Blade	DP0 - 16 through 20 DP1 - 26 through 30	DP0 - 16 through 25 DP1 - 26 through 35

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `extncfg --ve-mode` command to select the VE mode.
   - To set the operating mode to 20VE, enter the following command.

```
switch:admin> extncfg --ve-mode 20VE
```

   - To set the operating mode to 10VE, enter the following command.

```
switch:admin> extncfg --ve-mode 10VE
```

3. To display the current operating mode, enter the following command.

```
switch:admin> extncfg --show
VE-Mode: configured for 20VE mode
```

   The display shows the switch is in 20VE mode.

Clearing the SX6 Blade Configuration

When you remove a Brocade SX6 blade from a chassis, you can clear all the configuration information for that blade. You can also reset the blade to a default configuration when it is not removed from the chassis.

You can clear the configuration for an SX6 blade when the blade is removed from the chassis by using the `extncfg --config` command with the `-clear` option. When the configuration is cleared, any additional configuration information for that blade is removed from the configuration database for the blade that is no longer in the chassis. If you clear the configuration for a blade that is still inserted in a slot, only the configuration database is affected. To fully clear the configuration for an inserted blade, use the `slotpoweroff` command and the `slotpoweron` command.

Another way to clear the blade configuration is to reset the blade to its initial default configuration. Use the `extncfg --config` command with the`-default` option to reset the blade as though it was newly inserted into the chassis. This clears the active running configurations on an extension blade that is active in the chassis. This command is valid only for enabled/active extension blades in the chassis. The command clears the active configurations and removes all extension-related configurations.

1. Connect to the switch and log in using an account assigned to the admin role.
2. To clear the configuration in a chassis-based system when the SX6 blade is removed from slot 2, enter the following command.

   ```
   switch:admin> extncfg --config -clear -slot 2
   ```

   A slot power-off/power-on must be used to clear the configuration entirely.
3. To clear the configuration in a chassis-based system when the SX6 blade remains in an enabled slot, enter the following command sequence. In this example, the blade is in chassis slot 3.

   ```
   switch:admin> slotpoweroff 3
   Slot 3 is being powered off
   switch:admin> extncfg --config -clear -slot 3
   switch:admin> slotpoweron 3
   Powering on slot 3.
   ```

   The `slotpoweroff` and `slotpoweron` commands must be used to clear the configuration entirely.
4. To reset the default values for an SX6 blade that is inserted in slot 3 of a chassis-based system, enter the following command.

   ```
   switch:admin> extncfg --config -default -slot 3
   ```

   The result is as if the blade were newly inserted into the chassis and there were no extension configurations.

Configuring GE Mode on the Brocade 7810 Switch

The Brocade 7810 Extension Switch is configured with eight GE ports: two extension RJ-45 copper GbE ports and six optical ports that support SFP+ transceivers. GE0 and GE1 are always in copper mode and GE2 thru GE7 are always in optical mode. All eight ports are visible to you in the output of the`switchshow` command but only six are usable at a time. That is because GE0 thru GE1 and GE2 thru GE3 are mutually exclusive. That is, at any moment, only GE0, GE1 or GE2, GE3 can be used (see below).

NOTE
The Brocade 7810 Switch default port speed is 1Gb/s.

The Ethernet interfaces can be enabled in copper or optical mode using the `extncfg` command as follows:

- For copper mode, enter `extncfg --ge-mode copper`. In this mode, GE0 and GE1 will be enabled and usable but ports GE2 and GE3 will be disabled and unusable.
- For optical mode, enter `extncfg --ge-mode optical`. In this mode, GE0 and GE1 will be disabled and unusable where as GE2 and GE3 will be enabled and usable.

Ports that are unused cannot have any configurations on them. So when transitioning between modes, you are must delete any configurations on those ports prior to allowing the mode switch to occur.

NOTE
Regarding switching ports on mode change, per-port selection is not supported.

Consider the following for copper and optical operating modes:

- Switching between copper and optical operation is a non disruptive operation. The switch does not reboot.
- During a mode switch, no configurations may be present on the GE ports. If a configuration exists on a port that will be disabled, changing modes will be blocked, and the configuration must first be deleted.

```
switch:admin> extncfg –show
APP Mode is HYBRID (FCIP with IPEXT)
VE-Mode: None
GE-Mode: Optical
switch:admin> extncfg --help

Usage: extncfg <action> [options]

action:
  --ve-mode 10VE|20VE               - Set VE-Mode to 10VE or 20VE mode.
  --ge-mode copper|optical          - Set GE-Mode to copper or optical.
                                      (Brocade 7810 Only)
  --app-mode fcip|hybrid            - Set APP-Mode to FCIP or HYBRID (FCIP
                                      with IPEXT).
  --show                            - Display APP & VE mode details.
  --config -default                 - Default the Extension configuration.
  --config -clear                   - Clear the Config for Extension
                                      configuration.
  --fwld-prep [-version <version>]  - Prepare the switch for firmware
                                      download to the target version.

  --auth-error-monitor enable\disable - Enable or disable the IPsec Auth
                                      Error Monitor (Brocade FX8-24 Only)

  -h,--help                         - Print this usage statement.

switch:admin> extncfg –ge-mode copper
Operation Succeeded
switch:admin> extncfg –show
APP Mode is HYBRID (FCIP with IPEXT)
VE-Mode: Not Applicable.
GE-Mode: Copper
switch:admin>
```

This is a sample output of `switchshow` in optical mode.

```
switch:admin> switchshow
switchName:     sw0
switchType:     178.0
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:c4:f5:7c:01:34:48
zoning:         OFF
switchBeacon:   OFF
FC Router:      OFF
FC Router BB Fabric ID: 1
Address Mode:   0
HIF Mode:       OFF

Index Port Address  Media Speed   State       Proto
==================================================
   0   0   010000   id    N8      Online      FC  F-Port  10:00:00:05:33:26:72:f1
   1   1   010100   id    N8      Online      FC  F-Port  10:00:00:05:33:26:72:f0
```

```
  2   2   010200   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  3   3   010300   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  4   4   010400   id    N16     Online      FC  F-Port  20:02:00:11:0d:5a:00:00
  5   5   010500   id    N16     Online      FC  F-Port  20:03:00:11:0d:5a:01:00
  6   6   010600   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  7   7   010700   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  8   8   010800   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  9   9   010900   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
 10  10   010a00   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
 11  11   010b00   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
 12  12   010c00   --    --      Offline     VE
 13  13   010d00   --    --      Offline     VE
 14  14   010e00   --    --      Offline     VE
 15  15   010f00   --    --      Offline     VE
     ge0           id    1G      Offline     FCIP  Copper  Disabled (Unsupported blade mode)
     ge1           id    1G      Offline     FCIP  Copper  Disabled (Unsupported blade mode)
     ge2           id    1G      Online      FCIP
     ge3           id    1G      Online      FCIP
     ge4           id    1G      No_Sync     FCIP
     ge5           id    1G      No_Sync     FCIP
     ge6           id    10G     Online      FCIP
     ge7           id    10G     Online      FCIP
```

This is sample output for `switchshow` in optical mode and ge0 as persistent disabled.

```
switch:admin> switchshow
switchName:     sw0
switchType:     178.0
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:c4:f5:7c:01:34:48
zoning:         OFF
switchBeacon:   OFF
FC Router:      OFF
FC Router BB Fabric ID: 1
Address Mode:   0
HIF Mode:       OFF

Index Port Address  Media Speed   State       Proto
==================================================
   0   0   010000   id    N8      Online      FC  F-Port  10:00:00:05:33:26:72:f1
   1   1   010100   id    N8      Online      FC  F-Port  10:00:00:05:33:26:72:f0
   2   2   010200   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
   3   3   010300   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
   4   4   010400   id    N16     Online      FC  F-Port  20:02:00:11:0d:5a:00:00
   5   5   010500   id    N16     Online      FC  F-Port  20:03:00:11:0d:5a:01:00
   6   6   010600   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
   7   7   010700   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
   8   8   010800   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
   9   9   010900   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  10  10   010a00   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  11  11   010b00   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  12  12   010c00   --    --      Offline     VE
  13  13   010d00   --    --      Offline     VE
  14  14   010e00   --    --      Offline     VE
  15  15   010f00   --    --      Offline     VE
      ge0           id    1G      Offline     FCIP  Copper  Disabled (Persistent and Unsupported blade mode)

      ge1           id    1G      Offline     FCIP  Copper  Disabled (Unsupported blade mode)
      ge2           id    1G      Online      FCIP
      ge3           id    1G      Online      FCIP
      ge4           id    1G      No_Sync     FCIP
      ge5           id    1G      No_Sync     FCIP
      ge6           id    10G     Online      FCIP
      ge7           id    10G     Online      FCIP
```

This is a sample output for `switchshow` in copper mode.

```
switch:admin> switchshow
switchName:     sw0
switchType:     178.0
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:c4:f5:7c:01:34:48
zoning:         OFF
switchBeacon:   OFF
FC Router:      OFF
FC Router BB Fabric ID: 1
Address Mode:   0
HIF Mode:       OFF

Index Port Address  Media Speed   State       Proto
==================================================
   0   0   010000   id    N8      Online      FC  F-Port  10:00:00:05:33:26:72:f1
   1   1   010100   id    N8      Online      FC  F-Port  10:00:00:05:33:26:72:f0
   2   2   010200   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
   3   3   010300   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
   4   4   010400   id    N16     Online      FC  F-Port  20:02:00:11:0d:5a:00:00
   5   5   010500   id    N16     Online      FC  F-Port  20:03:00:11:0d:5a:01:00
   6   6   010600   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
   7   7   010700   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
   8   8   010800   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
   9   9   010900   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  10  10   010a00   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  11  11   010b00   --    N32     No_Module   FC  (POD license not assigned or reserved yet)
  12  12   010c00   --    --      Offline     VE
  13  13   010d00   --    --      Offline     VE
  14  14   010e00   --    --      Offline     VE
  15  15   010f00   --    --      Offline     VE
      ge0           id    1G      Online      FCIP  Copper
      ge1           id    1G      Online      FCIP  Copper
      ge2           id    1G      No_Sync     FCIP  Disabled (Unsupported blade mode)
      ge3           id    1G      No_Sync     FCIP  Disabled (Unsupported blade mode)
      ge4           id    1G      No_Sync     FCIP
      ge5           id    1G      No_Sync     FCIP
      ge6           id    10G     Online      FCIP
      ge7           id    10G     Online      FCIP
```

Configuring GbE Mode on the Brocade FX8-24

The Brocade FX8-24 Blade can be configured to support 1GbE, 10GbE, or both. The Brocade FX8-24 Blade is supported in a Brocade DCX 8510 chassis.

The GbE ports on a Brocade FX8-24 Blade can operate in one of three ways:

- **1 Gb/s mode**: In this mode, GbE ports 0 through 9 are enabled as 1GbE ports, with the XGE ports disabled. The 10GbE (FTR_10G) license is not required.
- **10 Gb/s mode**: In this mode, 10GbE ports xge0 and xge1 are enabled, with GbE ports 0 through 9 disabled. The 10GbE (FTR_10G) license is required and must be assigned to the slot in which the Brocade FX8-24 Blade resides.

  NOTE
  Switching between 10Gb/s mode and 1Gb/s mode disrupts traffic.

- **Dual mode**: In this mode, GbE ports 0 through 9 and the 10GbE port xge0 are enabled. The 10GbE port xge1 is disabled. The 10GbE (FTR_10G) license is required and must be assigned to the slot in which the Brocade FX8-24 Blade resides.

Auto-negotiation is enabled by default in 1G mode. In 10G mode it is disabled and not supported. When the port is set for 1G mode, you can disable auto-negotiation using the `portcfgge` *ge* `--disable -autoneg` command.

The following table shows which VE_Ports are available in each mode.

TABLE 28 Brocade FX8-24 GbE Modes and VE_Ports

Mode	GE ports	VE_Ports
1Gb/s	GbE 0 – GbE 9	VE_12 – VE_21
Dual	GbE 0 – GbE 9	VE_12 – VE_21
	XGE0	VE_22 – VE_31
10Gb/s	XGE1	VE_12 – VE_21
	XGE0	VE_22 – VE_31

1. Connect to the switch and log in using an account assigned to the admin role.
2. After deleting the port's configuration, use the `bladeCfgGeMode --set` *mode* `slot` *slot-number* command to set the GbE port mode of operation for the Brocade FX8-24 Blade.

   The following example enables 1g mode for GbE ports 0 through 9 on an FX8-24 blade in slot 8. Ports xge0 and xge1 are disabled.

   ```
   switch:admin> bladecfggemode --set 1g -slot 8
   ```

   The next example enables 10g mode for ports xge0 and xge1 on an FX8-24 blade in slot 8. Ports GbE ports 0 through 9 are disabled.

   ```
   switch:admin> bladecfggemode --set 10g -slot 8
   ```

   The following example enables dual mode for GbE ports 0 through 9 and xge0 on an FX8-24 blade in slot 8. Port xge1 is disabled.

   ```
   switch:admin> bladecfggemode --set dual -slot 8
   ```

3. Use the `bladecfggemode --show` command to display the GbE port mode for the Brocade FX8-24 Blade.

   The following example shows 10Gb/s mode is configured for the blade in slot 8.

   ```
   switch:admin> bladecfggemode --show -slot 8
   bladeCfgGeMode: Blade in slot 8 is configured in 10GigE Mode
   10GigE mode: only xge0 and xge1 are enabled (ge0-9 ports are disabled)
   The blade in slot81 supports IP Sec tunnels on only VEs 12 through 21 (xge1 or ge0-ge9)
   ```

Configuring VEX_Ports on the FX8-24

A virtual EX_Port that connects a Fibre Channel router to an edge fabric. From the point of view of a switch in an edge fabric, a VEX_Port appears as a normal VE_Port. It follows the same Fibre Channel protocol as other VE_Ports. However, the router terminates VEX_Ports rather than allowing different fabrics to merge as would happen on a switch with regular VE_Ports.

If you are going to use a VEX_Port in your tunnel configuration, use the **portCfgVEXPort** command to configure the port as a VEX_Port.

If the fabric is already connected, disable the Ethernet ports and do not enable them until after you have configured the VEX_Port. This prevents unintentional merging of the two fabrics.

VEX_Ports are described in detail in the*Brocade Fabric OS Administration Guide*. Refer to that publication if you intend to implement a VEX_Port.

The following example configures a VEX_Port, enables admin, and specifies fabric ID 2 and preferred domain ID 220.

```
switch:admin> portcfgvexport 18 -a 1 -f 2 -d 220
```

Configuring Ports

On the Brocade 7810 Extension Switch, the Brocade 7840 Extension Switch, and the Brocade SX6 Extension Blade, you can configure port speed of 1GbE or 10GbE on the ports that support 10GbE. You cannot change port speed on the ports that support 40GbE.

On the Brocade FX8-24 Extension Blade, available ports and speed are configured for the blade platform. See Configuring GbE Mode on the Brocade FX8-24 on page 94 for more information.

Configuring Port Speeds

You can configure the speed of 10GbE ports on the Brocade 7840 Switch and the Brocade SX6 Blade to 1Gb/s or 10Gb/s (default) using the `portcfgge` command.

The *base* Brocade 7810 Switch has two GE interfaces, each is limited to 1 Gb/s mode. An upgraded Brocade 7810 Switch can have 6GE interfaces, which can be set to 1Gb/s or 10Gb/s. On the Brocade 7810 Switch's optical SFP ports, 1Gb/s is the default speed.

Auto-negotiation is enabled by default in 1G mode only. In 10G mode it is disabled and not supported. You can disable auto-negotiation on a port set for 1G mode with the `portcfgge` *ge* `--disable -autoneg` command.

> NOTE
> Auto-negotiation is for 1G GE PHY negotiation. It is not a speed negotiation. The GE port can be set to either 1G mode or 10G mode. A port set in auto-negotiate mode is negotiating full duplex and use of pause frames (802.3X) with the attached switch. The port will come up if there is a mismatch in the negotiated parameters. However, the port will not come up if one end is has auto-negotiate enabled and the other end has auto-negotiate disabled.

Use the following steps to configure port speed on the Brocade 7840 Switch or the Brocade SX6 Blade 10GbE ports. The example commands are for a Brocade 7840 Switch because no slot number is used.

> NOTE
> If a port has active circuits and you change the port speed, you will disrupt traffic.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Perform one of the following steps:
   - To set the port speed at 1Gb/s for port ge4, enter the following:

     ```
     switch:admin> portCfgGe ge4 --set -speed 1G
     ```

   - To set the port speed at 10Gb/s for port ge4, enter the following:

     ```
     switch:admin> portCfgGe ge4 --set -speed 10G
     ```

   - To disable auto-negotiation on port ge4, enter the following:

     ```
     switch:admin> portCfgGe ge4 --disable -autoneg
     ```

   - To enable auto-negotiation on port ge4, enter the following:

     ```
     switch:admin> portCfgGe ge4 --enable -autoneg
     ```

3. To display the current port speed configuration for ge4, enter the following:

```
switch:admin> portCfgGe ge4 --show
 Port         Speed     Flags     Lag ID    Channel
-----------------------------------------------------
 ge4          10G       ----      ----      N/A
-----------------------------------------------------
Flags: A:Auto-Negotiation Enabled  C:Copper Media Type
L:LAN Port G: LAG Member

switch:admin> portshow ge4
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: 10G
```

Configuring Layer 2 Protocols

Configuring Global LLDP Parameters

The Link Level Discovery Protocol (LLDP) discovers the Ethernet neighbor interface at the link level, negotiates with the peer for configured parameters, negotiates the capabilities for Converged Enhanced Ethernet (CEE), and manages the CEE status of the port. To enable GbE connectivity, LLDP primarily maintains the GbE priority and logical link status, and negotiates them with the peer device.

NOTE
LLDP is enabled by default on GbE ports and the global parameters are applied to all ports.

To configure LLDP global parameters, follow these steps:

1. LLDP is enabled by default. If LLDP is not enabled, enable the LLDP protocol on the switch using the `lldp --enable` command.

```
switch#admin> lldp --enable
```

2. Configure the LLDP global parameters using the `lldp --config` command.

```
switch#admin> lldp --config -sysname Ext_Switch_15
switch#admin> lldp --config -sysdesc 7840_15_on_default_vf
switch#admin> lldp --config  -mx 5
switch#admin> lldp --config  -txintvl 40
```

3. Enable the required Type Length Values (TLVs) globally using the `lldp --enable -tlv` command.

```
switch#admin> lldp --enable  -tlv dcbx
```

NOTE
The `dcbx` and `sys-name` TLVs are enabled by default globally.

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

With Fabric OS 8.2.1, for the TLV port-desc, the string is now of the format: "Switch Model Name (in the case of a fixed-port switch)/Slot Model Name (in the case of a chassis): Mode + Speed + Slot/Port" (for example, Brocade FC32-64: ETH 10G 8/0 or Brocade SX6: WAN 10G 4/ge13).

- sys-cap—System Capabilities
- sys-desc—System Description

  With Fabric OS 8.2.1, for the TLV sys-desc, the default system description string in the TLV advertised for the switch is now of the format: "Switch Model Name, Firmware Version" (for example, Brocade 7840, Fabric OS Version 8.2.1).
- sys-name—System Name

4. You can use the following commands 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
ge1        120           110            7/ge1          0027.f8f0.a8d0 622 449 ven110
ge2        120           94             0024.389c.003c 0024.389c.0000 631 439 MLXe-33
ge4        120           106            0024.389c.003e 0024.389c.0000 622 439 MLXe-33
ge5        120           104            0024.389c.003f 0024.389c.0000 622 439 MLXe-33
ge7        120           106            0024.389c.0009 0024.389c.0000 622 439 MLXe-33
ge10       120           120            8/ge10         0027.f8f0.a8d0 600 425 ven110
ge11       120           112            ge11           0005.3365.7c42 622 450 SB116
ge13       120           95             port0          8c7c.ff21.9a2e 622 436 (null)
ge16       120           114            0024.389c.0041 0024.389c.0000 544 359 MLXe-33
ge17       120           112            0024.389c.0042 0024.389c.0000 622 439 MLXe-33
<output truncated>
```

Configuring and Activating an LLDP Profile for a Group of Ports

LLDP is enabled by default. You can also create multiple customized LLDP profiles with different parameters and apply them to specific groups of ports. If no specific LLDP profile is applied on a port, the global parameters are applied by default.

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
```

As of Fabric OS 8.2.1, a group of ports can now be expressed as a range rather than as a string of ports separated by commas. For 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

Switch#admin> lldp --show -stats ge17
LLDP Interface statistics for ge17
Frames transmitted: 8603
Frames Aged out:    0
Frames Discarded:   0
Frames with Error:  0
Frames Recieved:    0
TLVs discarded:     0
TLVs unrecognized:  0
```

Configuring Static and Dynamic LAGs Using LACP

The Brocade 7840 Switch, Brocade 7840 Switch, and the Brocade SX6 Blade platforms support both static and dynamic link aggregation groups (LAGs).

Whereas the Brocade 7840 Switch and the Brocade SX6 Blade platforms allow a maximum combination of static and dynamic LAGs of 8, the Brocade 7810 Switch supports 2. Furthermore, the supported speeds for LAG configurations on the Brocade 7810 Switch are 1G and 10G. Ethernet ports should be configured as LAN ports before adding them as part of any LAG. The `autoneg` configuration option is applicable only to the 1Gb/s LAG speed. 1Gbe RJ-45 copper ports on the Brocade 7810 Switch can also participate in LAG (both static and dynamic) formation, and the LAG could have a mix of copper and optic GbE ports.

When you create a LAG, you assign a name to it. Ports can be added to and removed from a named LAG. The port speed and auto-negotiation parameters must match for all ports being added to a LAG. Port speed and link auto-negotiation can be set for the LAG to control the setting of the individual LAG member ports. The port speed and link auto-negotiation setting of LAG member ports cannot be set individually. You can enable or disable individual ports in a LAG.

Auto-negotiation is part of IEEE 802.3 and applies only to 1Gb/s Ethernet (GE) interface mode. 10Gb/s and 40Gb/s interfaces have no auto-negotiation settings. Auto-negotiation does not apply to link speed, because the speed is not being negotiated. Auto-negotiation is specific to the use of Ethernet pause frame flow-control and full-duplex link settings. These settings can be determined by auto-negotiation between the two endpoints. If the data center LAN switch is configured for auto-negotiation, the IP extension platform must also be set for auto-negotiation. Otherwise, the links will not come online, because both ends of the link must be set identically.

NOTE

On 1-Gb/s interfaces, auto-negotiation is enabled by default. Beginning with Fabric OS 8.1.0, half-duplex is not supported.

Configuring a LAG is optional. However, without a LAG, you can connect only a single Ethernet link from the IP extension platform (such as the Brocade 7840 Switch, Brocade 7810 Switch, or the Brocade SX6 Blade) to the LAN switch in your data center. However, if you use multiple VLANs, you can have one connection per VLAN, which provides one logical connection per Layer 2 domain or Layer 2

VLAN, and only one connection is required. A single link does not provide redundancy. A LAG provides redundant links with redundant cables and optics. The recommended practice is to configure a LAG when connecting to a data center LAN switch.

On a Brocade 7840 Switch or a Brocade SX6 Blade, a LAG cannot have more than four interfaces (links) assigned. For the Brocade 7810 Switch, the limit is 2. Each LAG is configured with its own LAG name and ID.

A LAG treats multiple connections between two components logically as a single connection. The GbE port must be configured to operate in LAN mode before you can configure a LAG. On chassis that have multiple extension blades installed, GbE ports from different blades cannot be combined into a single LAG.

NOTE

Fabric OS 7.4.1 through Fabric OS 8.1 support static LAGs. Fabric OS 8.2.0 and subsequent releases support both dynamic and static LAGS. In Fabric OS 8.2.0, the `portcfg lag` and `portshow lag` commands are replaced with the `portchannel` command for LAG settings.

The following steps configure a link aggregation group (LAG) for a GbE LAN port.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `lacp --config --sysprio` command to set the global priority between 0 and 65535 for Link Aggregation Control Protocol (LACP) on the platform followed by `lacp --show` to view the settings. This step is optional, and it applies only to dynamic LAGs.

```
switch:admin> lacp --config --sysprio 100
switch:admin> lacp --show
LACP system priority: 100
LACP System ID: 0x8000,00-05-33-74-85-42
```

   The LACP system priority is set to 100. The default system priority is 32768.

3. Use the `portchannel --create` command to create a group for LAG. The commands show how to create a static group and a dynamic group.

```
switch:admin> portchannel --create dlag101 -type dynamic -key 555
switch:admin> portchannel --create slag101 -type static -key 100

switch:admin> portchannel --show
Name                          Type        Oper-State   Port-Count   Member Ports
--------------------------------------------------------------------------------
dlag101                       Dynamic     Offline      0
slag101                       Static      Offline      0
```

4. Use the `portchannel --add` command to add GE ports to the groups.

```
switch:admin> portchannel --add dlag101 -port ge6

switch:admin> portchannel --add slag101 -port ge15
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 --add slag101 -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
--------------------------------------------------------------------------------
dlag101                         Dynamic     Online       1            ge6
slag101                         Static      Offline      3            ge15 ,ge16 ,ge17
```

   With FOS 8.2.1, a group of ports can now be expressed as a range rather than as a string of ports separated by commas.

   In static LAGs, the modified LAN GbE ports are disabled by default and must be enabled after completing all the configuration on the LAG. When you add a port to a disabled dynamic LAG, the port gets disabled. When you remove a port from a disabled dynamic LAG, the port gets enabled. If a disabled dynamic LAG is deleted using the `portchannel --delete`, all the member ports get enabled.

5. Use the `portchannel --config` command to set the speed for all the ports in the LAG. For 1G ports, you can also enable auto negotiation. Auto negotiation cannot be enabled for 10G ports.

```
switch:admin> portchannel --config slag101 -speed 1G
switch:admin> portchannel --config slag101 -autoneg on
```

   NOTE

   The 1G ports are disabled after changing the autoneg configuration. The ports must be manually enabled.

6. For dynamic LAG member ports, you can configure the KAP timeout and priority.

```
switch:admin> portchannel --config  -port ge6-7 -priority 250
switch:admin> portchannel --config  -port ge6 -timeout s
switch:admin> portchannel --config  -port ge7 -timeout l
```

7. Use the `portchannel --enable` command to enable a LAG.

```
switch:admin> portchannel --enable slag101

switch:admin> portchannel --show -all
Name                            Type        Oper-State   Port-Count   Member Ports
--------------------------------------------------------------------------------
dlag101                         Dynamic     Online       1            ge6
slag101                         Static      Online       3            ge15 ,ge16 ,ge17
switch:admin>
```

   NOTE

   By default, the admin state of a LAG is enabled.

8. Use the `portchannel --show -detail` command to show detailed information about the LAG groups.

```
switch:admin> portchannel --show
Name                            Type        Oper-State   Port-Count   Member Ports
--------------------------------------------------------------------------------
test                            Dynamic     Online       2            ge5*,ge6
static                          Static      Offline      2            ge0 ,ge1

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         Enabledd
```

For additional information about additional options for the `portchannel` command, refer to *Brocade Fabric OS Command Reference*.

Configuring IPIF and IP

IP configuration consists of the following tasks.

- Configuring IP interface (IPIF) addresses.
- Configuring IP routes, if required.
- Configuring VLAN tag IDs on the Brocade 7840 Switch, Brocade 7810 Switch, and the Brocade SX6 Blade. On the Brocade FX8-24 Blade, VLAN IDs are configured when a tunnel is created or modified.
- Verifying IP connectivity.

Configuring IPIF

An IP interface (IPIF) is the end point of an extension tunnel or circuit. You must configure an IPIF for each circuit that you want to create. Depending on whether you intend to use eHCL, you might also need to configure IPIF for eHCL.

> NOTE
> The Brocade 7840 Switch and Brocade SX6 Blade support a maximum of 60 IPIFs per DP and 64 per Port. The Brocade 7810 Switch supports a maximum of 60 IPIFs per DP and 60 per Port.

On a Brocade 7840 Switch or Brocade SX6 Blade, each IPIF used in a circuit is associated with a GE port and data processor (DP) complex, either DP0 or DP1. The Brocade FX8-24 Blade does not use DP parameters when configuring circuits.

To configure an IPIF, use the `portCfg ipif create` command. The IPIF consists of an IP address, netmask, an IP MTU size, and other options depending on the extension switch or blade. On the Brocade X6 chassis or Brocade 7840 platform you can modify an existing IP address. On the Brocade DCX 8510, you must first delete an existing IP address and then create it with modified parameters.

Requirements and options for configuring IPIFs include the following:

- There are no addressing restrictions for IPv4 and IPv6 connections with both switches or blades in the tunnel running Fabric OS 7.0 and later. A tunnel can have both IPv4 and IPv6 circuits. However, each circuit in the tunnel must be either IPv4 at both ends or IPv6 at both ends.
- You can use CIDR notation for both the IPv4 addresses and IPv6 addresses.
- You can specify an optional IP MTU size. If not specified, the size will be set to 1500 bytes. The maximum supported MTU size is 9216.

  You can set the MTU manually or set it to AUTO. When set to AUTO, the circuit will use Path MTU Discovery (PMTUD) to determine an MTU.
- You can set a CIDR subnet mask of 31. When an IPv4 address with a 31-bit subnet mask is configured, using network addresses for broadcast purpose is not allowed. For example, there is a one-bit difference in the following two addresses:
  - Address A: 192.168.1.10/31
  - Address B: 192.168.1.11/31

  By using IP addresses with 31-bit prefixes (as defined in RFC 3021) you can reduce the number of IP subnets used by networking devices to establish IP connectivity to point-to-point WANs that they are connected to. The use of a 31-bit prefix allows the all-zeros and all-ones IP addresses to be assigned as host addresses on point-to-point networks. Prior to RFC 3021 the longest prefix in common use on point-to-point links was 30-bits, which meant that the all-zeros and all-ones IP addresses were wasted.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portCfg ipif create` command to create an IPIF. The following example shows an IPIF created on a Brocade 7840 Switch using GE port 2 and DP0.

   ```
   switch:admin> portcfg ipif ge2.dp0 create 192.168.1.24 netmask 255.255.255.0 mtu 1500
   ```

   The following example shows an IPIF created on a Brocade SX6 blade in slot 4 of a Brocade X6-8 Director.

   ```
   switch:admin> portcfg ipif 4/ge14.dp0 create 192.168.10.10/24 mtu 1500
   ```

   The following example shows an IPIF created on a Brocade FX8-24 blade in slot 8 of a Brocade DCX 8510.

   ```
   switch:admin> portcfg ipif 8/ge0 create 192.168.1.24/24 mtu 1500
   ```

   You can use either a CIDR prefix or a netmask value for IPv4 subnet masking. Any MTU larger than the default 1500 is considered a jumbo frame.
3. On the Brocade DCX 8510, you cannot modify an IPIF. To change an IPIF, you must first delete it. Use the `portcfg ipif delete` command to delete an IPIF.

   The following example shows an IPIF deleted on a Brocade FX8-24 blade in slot 8 of a Brocade DCX 8510. A DP parameter is not used.

   ```
   switch:admin> portcfg ipif 8/ge0 delete 192.168.1.24
   ```

   NOTE

   You cannot delete an IP interface if there is a tunnel or circuit configured to use it. Be sure to delete all tunnels, circuits, and IP routes using an interface before deleting it.

4. To modify an IPIF IP address on a Brocade 7840 Switch, a Brocade 7810 Switch, or a Brocade SX6 Blade, use the `portCfg ipif modify` command.

   The following example shows creating an IP address and then modifying the MTU value for that address on a Brocade 7840. The VLAN value can also be modified. The netmask value cannot be modified.

```
switch:admin> portcfg ipif ge3.dp0 create 192.168.10.10/24 mtu 1320
Operation Succeeded.
switch:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge3.dp0      192.168.10.10                 / 24   1320  0     U R M
 lan.dp0      192.168.10.1                  / 24   1500  0     U R M
 lan.dp1      1000:17:10::107               / 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

switch:admin> portcfg ipif ge3.dp0 modify 192.168.10.10 mtu 1500

!!!! 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.
switch:admin>
switch:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge3.dp0      192.168.10.10                 / 24   1500  0     U R M
 lan.dp0      192.168.10.1                  / 24   1500  0     U R M
 lan.dp1      1000:17:10::107               / 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
```

   If the IP address is a tunnel that is in use, the tunnel is brought down before the modifications are made. The tunnel is brought up after the modifications are applied.

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

   The following example shows IPIF information for a specific slot and port.

```
switch:admin> portshow ipif 11/ge17

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 11/ge17.dp0  192.168.10.10                 / 24   1500  0     U R M I
 11/ge17.dp0  FC00:20:32:11::17:3           / 64   1390  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
```

Configuring IP Route

Routing is based on the destination IP address presented by an extension circuit. If the destination address is not on the same subnet as the Ethernet port IP address, you must configure an IP route to that destination with an IP gateway on the same subnet as the local Ethernet port IP address.

When configuring IP routes, be aware of the following limitations:

- You can define up 128 routes per GbE port on a Brocade 7840 Switch or a Brocade SX6 Blade and a maximum of 120 routes per DP.
- You can define up 120 routes per GbE port on a Brocade 7810 Switch.
- You can define up to 32 routes for each GbE port on the Brocade FX8-24 Blade.
- When you use the `portCfg iproute` command, it can display more routes than those you configured after all routes are added.

To configure a route, use the `portCfg iproute create` command to specify the destination IP address, subnet mask, and address for the gateway router that can route packets to the destination address.

Optionally, on the Brocade FX8-24 blade, you can configure an IP route for a failover crossport using the `-x` or `--crossport` option. For information on configuring IP routes using crossport addresses, see Configuring IP Routes with Crossports on page 178.

The following figure shows an IP route sample configuration. The Brocade X6 Director connects to the WAN through a gateway, as does the Brocade 7840. Notice that on each side of the WAN, the IP addresses for gateway and switch are in the same subnet.

FIGURE 18 Configuring an IP Route

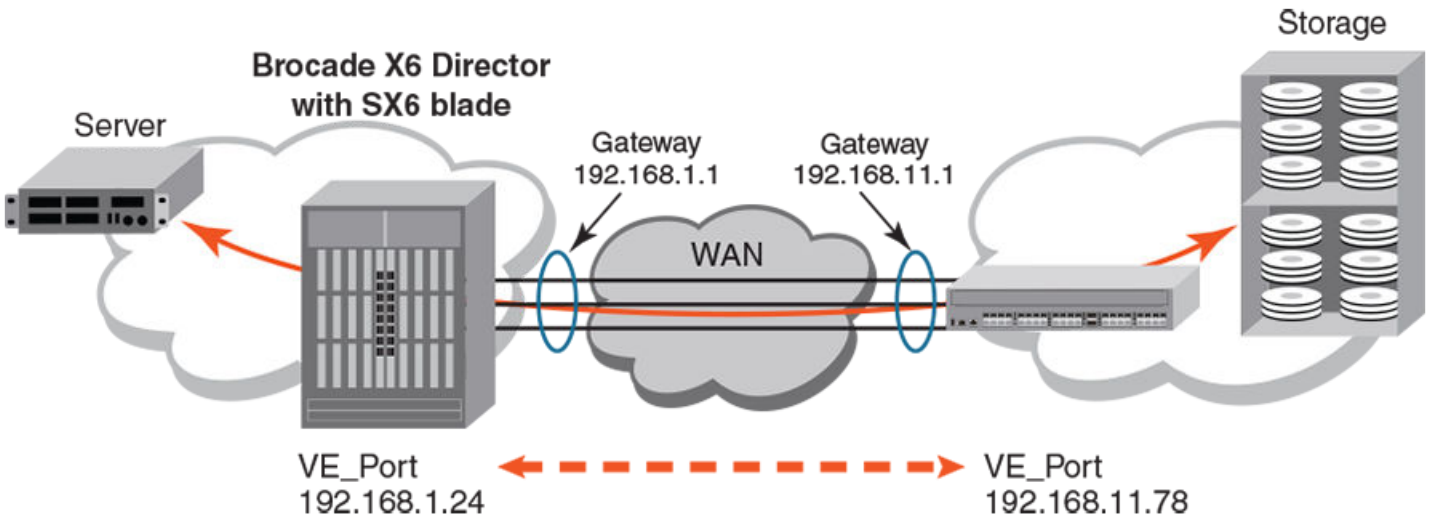

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portCfg iproute create` command to create a route on the Brocade X6 Director to the gateway.

   The following command creates an IP route to destination network 192.168.11.0 for port ge0 on the SX6 blade in slot 8 of the director. The route is through local gateway 192.168.1.1. After the destination address, either specify a pfx (prefix length) or network mask.

```
switch:admin> portcfg iproute 8/ge0.dp0 create 192.168.11.0/24 192.168.1.1
Operation Succeeded
```

3. Use the `portCfg iproute create` command to create a route on the Brocade 7840 to the gateway on its side of the network.

   The following command creates an IP route to destination network 192.168.1.0 for port ge2 on the Brocade 7840 switch. The route is through local gateway 192.168.11.1. Because Ethernet ports are shared between DP complexes, the ge1.dp0 option directs the command to a specific DP.

   ```
   switch:admin> portcfg iproute ge2.dp0 create 192.168.1.0/24 192.168.11.1
   Operation Succeeded
   ```

4. Use the `portshow iproute` command to display the configured IP routes.

   The following example shows IP routes on the Brocade 7840 Switch.

   ```
   switch:admin> portshow iproute

    Port         IP Address               / Pfx  Gateway               Flags
   ------------------------------------------------------------------------------
    ge2.dp0      192.168.1.0              / 24   192.168.11.1          U G S
    ge2.dp0      192.168.11.0             / 24   *                     U C
    ge2.dp0      192.168.11.1             / 32   *                     U H L
   ------------------------------------------------------------------------------
    Flags: U=Usable G=Gateway H=Host C=Created(Interface)
           S=Static L=LinkLayer X=Crossport
   ```

5. Use the `portcfg iproute modify` command on the Brocade 7840 Switch, Brocade 7810 switch, or the Brocade SX6 Blade to modify the local gateway address of an IP route. You cannot use this command to modify the destination network address. If this needs to be modified, you must delete the IP route, and then recreate it. Also, you cannot use this command to change the prefix length or network mask.

   The following example modifies the gateway IP address for an existing IP route on the Brocade 7840.

   ```
   switch:admin> portcfg iproute ge2.dp0 modify 192.168.1.0/24 192.168.11.5
   Operation Succeeded
   ```

Configuring VLANs

When a VLAN tag is created on a circuit, all traffic over that circuit will use the specified VLAN. When the layer 2 path for the IPIF is tagged in a VLAN, the extension circuit must be configured with the matching tag. This tagging ensures that all traffic using that circuit is sent with the appropriate VLAN tag and all incoming traffic for that circuit will be checked to ensure the correct tag is set.

The Brocade 7840 Extension Switch, the Brocade 7810 Extension Switch, and the Brocade SX6 Extension Blade use the `portcfg ipif` command to include the VLAN tag information when the IPIF is created. To change the VLAN tag, you must first delete the IPIF then create it with the new values.

The Brocade FX8-24 Extension Blade uses the `portcfg vlantag` command to add VLAN tag information to an interface IP address. You can also include the VLAN tag when you initially define the FCIP circuit on the Brocade FX8-24 Blade so that you need not perform the additional step of adding the tag after the circuit is defined. In addition, VLAN tags can be added to the XGE ports on a Brocade FX8-24 Blade when the blade is configured to use crossports. For more information, see Configuring VLAN Tags with Crossports on page 179.

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

2. On the Brocade 7840 Switch, Brocade 7840 Switch, or the Brocade SX6, use the `portcfg ipif` command to add VLAN tag information to an IPIF.

   The following example creates an IPIF with VLAN 200.

```
switch:admin> portcfg ipif ge2.dp0 create 192.168.5.20/24 mtu 1650 vlan 200
Operation Succeeded
```

3. Use the `portcfg ipif` command to display the port information.

   The following example shows VLAN 200. When the IPIF with VLAN 200 is used in a circuit, the destination IPIF must also use VLAN 200.

```
switch:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.5.2                   / 24   1500  0     U R M I
 ge2.dp0      192.168.5.20                  / 24   1650  200   U R M
 ge2.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
```

Verifying IP Connectivity

To verify that you have IP connectivity, use the `portcmd --ping` command to confirm the source and destination IPIF addresses can communicate. You must have IP connectivity between the two circuit endpoints for them to talk to each other and bring the circuit up.

1. Use the `portcmd --ping` command to verify connectivity.

   The following example shows a successful ping operation, which confirms network connectivity between the source and destination addresses.

```
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
```

2. Use the `portshow ipif` command to display port information.

   The following example shows port information on two different switches, Switch_A and Switch_B, which are connected back-to-back. You can see that Switch_A has VLAN 200 and Switch_B does not. The ping from Switch_A to Switch_B fails.

```
Switch_A:admin> portshow ipif
 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.5.20                  / 24   1650  200   U R M
-------------------------------------------------------------------------------
```

   Switch B

```
Switch_B:admin> portshow ipif
 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.1.20                  / 24   1650  0     U R M
-------------------------------------------------------------------------------
```

   Ping from Switch A to Switch B

```
Switch_A:admin> portcmd --ping ge2.dp0 -s 192.168.5.20 -d 192.168.1.20

PING 192.168.1.20 (192.168.5.20) with 64 bytes of data.
From 192.168.1.20: icmp_seq=1 Request timed out
From 192.168.1.20: icmp_seq=2 Request timed out
From 192.168.1.20: icmp_seq=3 Request timed out
From 192.168.1.20: icmp_seq=4 Request timed out

--- 192.168.1.20 ping statistics ---
4 packets transmitted, 0 received, 100% packet loss, time 12736 ms
```

Configuring a Service-Level Agreement

The primary purpose of a service-level agreement (SLA) is to provide automated testing of a circuit before it is placed into service. The SLA checks the circuit for packet loss percentage. If you need to verify the circuit for additional network performance, such as throughput, congestion, out-of-order delivery, use WAN Tool to run tests manually. See Using WAN Tool on page 219 for information.

The SLA feature is supported on the Brocade 7840 Switch, the Brocade 7810 Switch, and the Brocade SX6 Blade.

You must configure an SLA session at each end of the circuit being tested. The SLA session uses information from the circuit configuration to configure and establish the SLA connections. 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 minor mismatch. When the session is established, traffic starts automatically. For the duration of the test, the traffic must remain under the specified loss percentage before the circuit is placed into service. On a Brocade 7840 Switch or a Brocade SX6 Blade, up to 20 SLA sessions can be defined per DP. On a Brocade 7810 switch, a maximum of 12 sessions can be defined.

In addition to packet loss, the SLA can also test for timeout duration. If the timeout value is reached during the SLA session, the session is terminated and the circuit put into service. A timeout value of "none" means the test runs until the runtime and packet-loss values are met.

Interaction with an SLA while it is running is limited. You can view statistics and you can abort an active session. Any attempt to modify a session while it is active is blocked, which means the WAN Tool commands cannot be used while an SLA session is running.

Whenever a tunnel or circuit goes offline and comes back online, or when a circuit is administratively disabled then enabled, the SLA session is started and tests the link before allowing the circuit to go back into service. Configured SLA sessions are persistent across reboots, because circuit configurations are persistent across reboots and the SLA is part of the circuit configuration. However, user-configured WAN Tool sessions are not persistent.

After configuring an SLA, you assign the SLA to a specific circuit with the `portcfg fcipcircuit` command.

NOTE
On the Brocade 7810 Switch, when you have SLA's 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 may exceed the 2.5Gb/s max of the platform. In this type of configuration, SLA use should be avoided or used with caution.

NOTE
During an eHCL reboot, the SLA is disabled and no new SLA sessions can be created until all eHCL operations are complete. After all eHCL operations are complete, the SLA is reenabled.

The following steps show how to configure, display, and abort SLA sessions. You can configure multiple SLAs.

1. Use the `portcfg sla` command to create an SLA session. You must create an SLA session at each end of the circuit, but the session names need not match.

```
switch:admin> portcfg sla networkA create --loss 0.5
switch:admin> portcfg sla networkB create --loss 1 --runtime 10 --timeout 30
```

   The SLA named networkA is created with a packet-loss limit of 0.5 percent. It will run for the default 5 minutes and never timeout.

   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.

2. Use the `portcfg fcipcircuit` command to assign an SLA to a circuit. The following command modifies a circuit and assigns 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 fcipctunnel -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): c=Control h=HighPri m=MedPri l=LowPri I=IP-Extension
                 i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                 a=FastDeflate d=Deflate D=AggrDeflate
      (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-Configured
```

   The sample output assumes that both ends of the circuit are configured with matching SLAs, where packet loss, runtime, and timeout values are the same. Any difference in transmit speed between the endpoints will be negotiated to the lower value. The output shows tunnel 24, circuit 0 is under active test and has an SLA configured.

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 the SLA sessions complete their run, additional information is displayed in the `portcmd --wtool show --detail` command that shows when the session either completed or was stopped, and the completion reason. Partial command output provides an example showing the session 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 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

Alternatively, you can display the individual TCP Lite connections under each session with the
following command.
Here, 24.0 is a 1G AWT session that consists of 2 .5G connections.

switch:admin> portcmd --wtool show -c

 Session OperSt  LocalIp/lPort  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

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)
```

The WAN Tool session 24.0 is aborted.

7. Use the `portcmd --wtool stop-all` command to abort all running SLA sessions.

```
switch:admin> portcmd --wtool stop -all
```

Configuring IPsec

IPsec is enabled on the tunnel level, not on the circuit level. This means that all circuits in a tunnel use the same IPsec settings. Different tunnels can have different 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 through public key infrastructure (PKI).

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 come up and will prevent the extension tunnel from coming up.

The PSK requirements are as follow:

- For the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade, the pre-shared key must be a 16- to 64-character string.
- For the Brocade FX8-24 Blade, the pre-shared key must be a 32-character string.

The PKI requirements are as follow:

- ECDSA certificates are supported only on the Brocade 7840 Switch, the Brocade 7810 Switch, and the Brocade SX6 Blade.
- Non-ECDSA certificates are not supported.
- PKI support is restricted to key-size P384 and hash-type SHA384.
- X.509 certificates are supported.
- Before downgrading to a release before Fabric OS 8.2.0, you must remove all imported certificates.

NOTE
For information on configuring and managing certificates using the `secCertMgmt` command, refer to "SSL configuration overview" in the *Brocade Fabric OS Administration Guide*.

When running IPsec, it is recommended that both sides of the extension tunnel should be running the same Fabric OS version.

NOTE
Creating and using IPsec policies is recommended for the security of the data that is transmitted over the network.

On the Brocade 7810 switch, the Brocade 7840 switch, and the Brocade SX6 blade, you first define an IPsec policy though you can define multiple IPsec policies. The IPsec policy is enabled on that tunnel when configured. For information on how to configure a tunnel with IPsec policy, see Configuring IPsec on the Brocade 7810, the Brocade 7840, and the Brocade SX6 on page 112.

With the Brocade 7840 and the Brocade SX6, IPsec will infer the FPGA type and set the max number of available SAs at 1024. (As IPsec does not have direct access to the PCI bus to determine the FPGA type, IPsec will use the board type to infer the FPGA type.) However, on the Brocade 7810, we provide a maximum of 256 SAs.

NOTE
IKE requires four SAs per session and the Brocade 7810 switch provides a maximum of 48 IKE sessions. So, IKE will require no more than 192 concurrent SAs for data traffic, leaving 64 SAs available for re-keying.

On the Brocade FX8-24 blade, the IPsec policy is defined and enabled when you create or modify a tunnel on the blade. For information on how to configure IPsec on a Brocade FX8-24, see Configuring IPsec on the Brocade FX8-24 Blade on page 118.

Configuring IPsec on the Brocade 7810, the Brocade 7840, and the Brocade SX6

Before enabling IPsec on a tunnel on the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade, you must first define an IPsec policy. You can define multiple policies, but only one policy can be applied to each tunnel. All circuits in that tunnel use the same IPsec policy.

Beginning with Fabric OS 8.2.0, you can modify an IPsec policy while the policy is still assigned to a tunnel or WAN Tool session. You must provide the new profile and new authentication data. If you are modifying only the authentication data, you need only provide the authentication data. In some instances the local side and remote side can get out of sync and indicate an authentication error. See IPsec IKE Authentication Failures on page 116 for information on how to restart IKE authentication.

When you use pre-shared key (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.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg ipsec-policy` command to define a policy. The pre-shared key must be 16 to 64 characters long.

   The following example creates an IPsec policy that has the name “myPolicy1”. The pre-shared key is 16 characters long.

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

3. After you create the IPsec policy, you can use it when you configure the tunnel. Use the `portcfg fciptunnel` command to enable a policy on a tunnel.

   NOTE
   This command is disruptive when an existing tunnel with active circuits is modified.

   The following example uses the `portcfg fciptunnel modify` command to enable the policy “myPolicy1" for an existing tunnel on a Brocade 7840 Switch. Remember to enable IPsec on each end 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
   ```

4. 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       Authentication data
--------------------------------------------------------------------------------
 myPolicy1          123ashorttestkey
--------------------------------------------------------------------------------
Flags: *=Name Truncated. Use "portshow ipsec-policy -d for details".
```

   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       Authentication data
 IKE-ID   Oper  Flg Local-Addr      Remote-Addr     IKE Rekey      ESP Rekey
--------------------------------------------------------------------------------
 myPolicy1          ****************
  dp0.0   UP  I     192.168.1.2     192.168.5.2     5h59m51s  0    3h20m10s  0
  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".I=Initiator R=Responder
```

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

```
switch:admin> portshow ipsec-policy -d
 IPSec-policy: myPolicy1
-------------------------------------------------
  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:       3
```

5. To disable an IPsec policy on a tunnel, use the `portcfg 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
```

6. 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
```

7. To create an IPsec policy using public key infrastructure (PKI), use the `portcfg ipsec-policy policy1 create --profile pki` command. You must previously obtain a CA certificate.

   The following example creates a PKI policy.

```
switch:admin> portcfg ipsec-policy policy1 create --profile pki --key-pair sb127_kp
Operation Succeeded
```

8. Use the `portshow ipsec-policy --detail` command to display details.

```
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:         sb127_kp
 Certificate:      sb127_cert.pem
 Certificate Hash: aff6fea1b19d81ea43aa72f4275a9cf550edadc0
 Num IKE Session:  0
```

   The following example shows IPsec with active IKE sessions. The summary info for the IKE data will include the remote certificate requested and an indicator if the hash matches or not.

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

IPSec Policy       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
```

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

```
switch:admin> portcfg ipsec-policy psk1 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
switch:admin>
```

10. To modify a PSK-to-PKI policy, you must modify the profile and the authentication data. Both actions must occur at the same time.

```
switch:admin> portcfg ipsec-policy policy1 modify --profile pki --keypair sb65kp1

!!!! 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>
```

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

```
switch:admin> portcfg ipsec-policy policy1 modify --profile preshared --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
switch:admin>
```

12. To modify a PKI-to-PKI policy, only the authentication data must be modified.

```
switch:admin> portcfg ipsec-policy policy1 modify --keypair sb65kp2

!!!! 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>
```

IPsec IKE Authentication Failures

On the Brocade 7810 switch, the Brocade 7840 switch, or the Brocade SX6 blade, an IKE authentication error can require user intervention to correct. This error occurs if there is an IKE session parameter mismatch. When such an error occurs, the IKE session is put into a faulty state and remains there until manually corrected thereby preventing an intruder from using multiple attempts to authenticate. For example, if the preshared-key is incorrect during the initial IKE authentication exchange, it will trigger an authentication error and display an RASlog XTUN-2012 message. Refer to the *Brocade Fabric OS Message Reference* for details. You can enter the `portshow ipsec-policy --ike` command where the IKE Operational Status is reported as FAULT.

To recover from an IKE authentication error, you must restart the IKE authentication exchange. Prior to Fabric OS 8.1.0, to restart the IKE authentication, remove IPsec from the tunnel and add it back in. If a policy is in use, you must first disable it. Beginning with Fabric OS 8.1.0, you could restart the IKE authentication with the `portcfg ipsec-policy restart` command and name the IPsec policy that you want to restart.

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

Prior to Fabric OS 8.1.0, follow these steps.

2. To disable an IPsec policy on a tunnel, use the `portcfg modify` command.

   The following example disables the IPsec policy on tunnel 24.

```
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
```

3. 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
```

4. Use the `portcfg ipsec-policy` command to define a policy. The pre-shared key must be 16 to 64 characters long. After you define the policy, enable it on the tunnel.

Beginning with Fabric OS 8.1.0, follow these steps.

5. Connect to the switch and log in using an account assigned to the admin role.
6. You can restart IKE sessions under a policy that are in an inactive state. This is useful when an authentication error occurs on an IKE session in FIPS mode and you want to restart the session without creating and enabling duplicate policies on a tunnel.

   The following commands show an IKE session that is in FAULT state and restarting the IKE session for a specific IPsec policy.

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

 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 pol1 restart
Operation Succeeded
switch:admin> portshow ipsec-policy -i
 IPSec Policy       Authentication data
 IKE-ID   Oper  Flg Local-Addr      Remote-Addr     IKE Rekey      ESP Rekey
-------------------------------------------------------------------------------
 pol1               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
```

Configuring IPsec on the Brocade FX8-24 Blade

On the Brocade FX8-24 Blade, you use the `portcfg fciptunnel` command to define and enable the IPsec policy on a tunnel. The pre-shared key must be a 32 characters string.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `--ipsec` option of the `portcfg fciptunnel create` and `portcfg fciptunnel modify` commands to define a policy.

   The following examples are for the Brocade FX8-24 Blade. They show IPsec and IKE keys enabled for traffic from VE_Ports 2/12 and 2/13 across multiple circuits.

```
portcfg fciptunnel 2/12 create --remote-ip 192.168.0.90 --local-ip 192.168.0.80 -b 50000 -B 50000 \
-x 0 -d c0 -i -K12345678901234567890123456789012
portcfg fcipcircuit 2/12 create 1 --remote-ip 192.168.1.90 --local-ip 192.168.1.80 -b 50000 -B 50000 -x 0
portcfg fcipcircuit 2/12 create 2 --remote-ip 192.168.2.90 --local-ip 192.168.2.80 -b 50000 -B 50000 -x 0
portcfg fcipcircuit 2/12 create 3 --remote-ip 192.168.3.90 --local-ip 192.168.3.80 -b 50000 -B 50000 -x 0
portcfg fcipcircuit 2/12 create 4 --remote-ip 192.168.4.90 --local-ip 192.168.4.80 -b 50000 -B 50000 -x 0
portcfg fcipcircuit 2/12 create 5 --remote-ip 192.168.5.90 --local-ip 192.168.5.80 -b 50000 -B 50000 -x 0

portcfg fciptunnel 2/13 create --remote-ip 192.168.0.91 --local-ip 192.168.0.81 -b 50000 -B 50000 -x 0 -d \
c0 -I -K12345678901234567890123456789012 -l
portcfg fcipcircuit 2/13 create 1 --remote-ip 192.168.1.91 --local-ip 192.168.1.81 -b 50000 -B 50000 -x 0
portcfg fcipcircuit 2/13 create 2 --remote-ip 192.168.2.91 --local-ip 192.168.2.81 -b 50000 -B 50000 -x 0
portcfg fcipcircuit 2/13 create 3 --remote-ip 192.168.3.91 --local-ip 192.168.3.81 -b 50000 -B 50000 -x 0
portcfg fcipcircuit 2/13 create 4 --remote-ip 192.168.4.91 --local-ip 192.168.4.81 -b 50000 -B 50000 -x 0
portcfg fcipcircuit 2/13 create 5 --remote-ip 192.168.5.91 --local-ip 192.168.5.81 -b 50000 -B 50000 -x 0
```

   The `-l`(legacy) option specifies to use the IPsec connection process compatible with Fabric OS releases prior to FOS 7.0.0. Note that the `-l` option is a disruptive request that causes the tunnel to bounce.

FX8-24 Authentication Tag Error Monitor

With the option `--auth-error-monitor enable|disable` of the `extnCfg` command you can set the automatic recovery mode of the IPsec Authentication Error Monitor for all Brocade FX8-24 blades in the chassis. The mode is disabled by default.

```
switch:admin> extncfg --help

Usage: extncfg <action> [options]

action:
  --ve-mode 10VE|20VE               - Set VE-Mode to 10VE or 20VE mode.
                                      (7840 / SX6 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 - 7840 / SX6 only).
  --slot <#>                        - Specify slot number for operation.
  --show [-slot <#>|-all]           - Display APP/VE/GE mode details.
  --config -default [-slot <#>|-all] - Default the Extension configuration.
  --config -clear [-slot <#>|-all]   - clear the Extension configuration.
  --fwdl-prep [-version #.#.#] [-abort] - Prepare the switch for firmware
                                      download to the target version.
  --auth-error-monitor enable|disable  - Enable/Disable the IPSec Auth
                                      error monitor for FX8-24 blades.

  -h,--help                         - Print this usage statement.

switch:admin> extncfg --show
IPSec Auth-Error Reset Detection: Disabled
switch:admin> extncfg --auth-error-monitor enable
Operation succeeded.
switch:admin> extncfg --show
IPSec Auth-Error Reset Detection: Enabled
switch:admin>
```

If the mode monitor mode is enabled and the authentication error threshold is exceeded, the slot automatically resets to recover from the failure condition.

The IPsec Authentication Error Threshold Monitor is associated with two RASLOGs, which are generated whenever the authentication error monitor detects an error. The RASLOGs differ based on whether the recovery mode is enabled or disabled.

- When the authentication error monitor is disabled, you observe the following: `IPS-2000: IPSec authentication error threshold exceeded for slot %d`.
- When the authentication error monitor is enabled, you observe the following: `IPS-2001: IPSec authentication error threshhold exceeded for slot %d. Resetting blade`.

Configuring Extension Tunnels for FCIP

Before you begin configuring tunnels, make sure that you have configured the following:

- Platform modes
  - FCIP only (Brocade 7840 Switch and Brocade SX6 Blade) or Hybrid mode (Brocade 7840 Switch and Brocade SX6 Blade)

    NOTE
    The Brocade 7810 Switch is always in Hybrid mode; there is no configuration to change.

  - VE mode (Brocade 7480 Sswitch and Brocade SX6 Blade)
  - GE mode (Brocade 7810 Switch)
  - GbE mode (Brocade FX8-24 Blade)
  - VEX_Ports (Brocade FX8-24 Blade)
- Port speed
- IP information
  - IP interfaces (IPIFs)
  - IP routes
  - VLAN
- Service-level agreement (SLA)
- IPsec policies (Brocade 7840 Switch, Brocade 7810 Switch, and Brocade SX6 Blade)

In addition, you must verify IP connectivity before configuring a tunnel. Otherwise, the tunnel will not come up.

NOTE
A Brocade 7840 Switch, Brocade 7810 Switch, and Brocade SX6 Blade can connect with one-another through an extension tunnel but they cannot connect to a Brocade FX8-24 Blade.

When you configure a tunnel, you must configure the local side and the remote side before it comes up. A tunnel consists of one or more circuits. When you first configure a tunnel, it contains a circuit that is identified as circuit 0. You can then configure additional circuits for that tunnel.

Tunnels exist between end points. Each end point is identified by its IPIF address. Tunnels are established through the VE_Ports on the switch or blade. For example, when you configure tunnel 24 on the local side (the Brocade 7840 Switch), that tunnel is established through VE_Port 24. The tunnel configuration identifies the local and remote IPIFs. On the remote side, the local and remote IPIF addresses for that tunnel are flipped. The remote side can have a different VE_Port for its end of the tunnel. For example, on a Brocade SX6 blade, tunnel 11/16 is established on VE_Port 16 in chassis slot 11.

To help you stage a configuration without committing specific circuit parameters, we recommend that you first configure the tunnel with appropriate tunnel parameters only (no IP addresses or circuit options) using `portcfg fciptunnel` command. Then, configure circuit 0 and additional circuits using `portcfg fcipcircuit` commands.

Configuring VE_Ports to Persistently Disable

NOTE

Disabling the VE_Port is not required.

Persistently disable the VE_Port before you begin configuring the tunnel on the VE_Port. By default, the VE_Port is persistently enabled. You manually change the state of the VE_Ports from persistently enabled to persistently disabled. This action prevents unwanted fabric merges from occurring until the tunnel is fully configured. After the tunnels have been fully configured on both ends of the tunnel, you must persistently enable the ports.

Persistently disabled ports remain disabled across power cycles, switch reboots, and switch enables.

If you enter `portCfgPersistentDisable` and receive "command not allowed in fmsmode" or "command not found" messages, FICON Management Server (FMS) mode may be enabled. You cannot use the `portcfgpersistentdisable` or `portcfgpersistentenable` commands with FMS mode enabled. Use the `portdisable` and `portenable` commands instead.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfgpersistentdisable` command to disable any VE_Ports that you will use in the tunnel configuration.

   This example disables VE_Port 24.

   ```
   switch:admin> portcfgpersistentdisable 24
   switch:admin>
   ```

   There is no additional response or confirmation after entering the command.
3. Use the `portshow` command to show the status of a port.

   This example shows the status of VE_Port 24. You can see the status of persistent disable in the output. (The "<<<<<" is not part of the actual output.)

   ```
   switch:admin> portshow 24
   portIndex:  24
   portName: 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: No
   ```

4. When 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 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 FMS mode is enabled.

      This example disables port 24, a VE_Port in a Brocade 7840 switch.

      ```
      switch:admin> portdisable 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. (The "<<<<<" is not part of the actual output.)

   ```
   switch:admin> portshow 24
   portIndex:  24
   portName: 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:
   ```

Configuring Tunnels

Before configuring tunnels, make sure the following are in place:

- IPIFs are configured for the local switch and remote switch end points of the tunnel and each of its circuits.
- IP routes are configured, which is required when the end point are on different subnets.
- VLANs are configured (optional).
- IPsec policies are configured (optional but recommended).
- IP connectivity is verified.

The Brocade 7840 Switch and Brocade SX6 Blade support a maximum of 20 tunnels; the Brocade 7810 Switch supports a maximum of 4 tunnels.

The following table shows the platforms supported in Fabric OS 8.2.1 that you can connect with tunnels. For example, you can create a tunnel between two Brocade FX8-24 Blades, but not between an FX8-24 and a Brocade SX6 Blade.

TABLE 29 Extension Tunnel Connections

	Brocade 7840	Brocade 7810	Brocade SX6	Brocade FX8-24
Brocade 7840	Yes	Yes	Yes	No
Brocade 7810	Yes	Yes	Yes	No
Brocade SX6	Yes	Yes	Yes	No
Brocade FX8-24	No	No	No	Yes

Tunnels are created between the VE_Ports of a switch or blade. You use the IPIFs on each switch to identify the local and the remote endpoints of the tunnel. You configure each side of the tunnel so that the endpoints point to one another.

When a tunnel is created with the `portcfg fciptunnel create` command and you specify the local IP and remote IP address endpoints, it contains one circuit, which is circuit 0. To add additional circuits to the tunnel, use the `portcfg fcipcircuit` command. You can create a tunnel with no circuits, but no traffic will pass.

NOTE
Any options that you specify when you create a tunnel must match at each end. The VE_Port can be different.

Refer to the *Brocade Fabric OS Command Reference* for additional information on the options available with the `portcfg fciptunnel` command.

1. Connect to the local switch and log in using an account assigned to the admin role.
2. Use the `portcfg fciptunnel create` command to create a tunnel on the local and remote switch.
   a) Create the local tunnel.

      The following example creates a tunnel (and circuit) on the local Brocade 7840 Switch.

      ```
      Local_switch:admin> portcfg fciptunnel 24 create --local-ip 192.168.5.2 --remote-ip 192.168.1.2
      -b 5000000 -B 5000000 -k 1000
      ```

   b) Create the remote tunnel.

      The following example creates a tunnel (and circuit) on the remote Brocade 7840 Switch.

      ```
      Remote_switch:admin> portcfg fciptunnel 24 create --local-ip 192.168.1.2 --remote-ip
      192.168.5.2  -b 5000000 -B 5000000 -k 1000
      ```

   Notice that the local IP and remote IP on each switch point to one other.
3. Use the `portshow fciptunnel -c` command to display the tunnel and circuit information.

   The following example displays the tunnel and circuit information on the local switch. OpStatus shows **Up**, which indicates that both sides of the tunnel are configured correctly.

   ```
   Local_switch:admin> portshow fciptunnel -c

    Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
   ------------------------------------------------------------------------------
    24    -         Up       ---------    18h3m    0.00    0.00   5       -       -
    24    0 ge2     Up       ----ah--4    18h3m    0.00    0.00   5  5000/5000  0/-
   ------------------------------------------------------------------------------
    Flags (tunnel): 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 one circuit, circuit 0.

4. Use the `portshow fcipcircuit` command to display the circuit information for a specific tunnel.

   The following example shows the circuit information for tunnel 24 on the local switch.

```
Local_switch:admin> portshow fcipcircuit 24

 Circuit 24.0 (DP0)
 ==================================================
  Admin/Oper State     : Enabled / Online
  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: 5000000 / 5000000 kbps
  Peer Comm Rates      : 5000000 / 5000000 kbps
  Actual Comm Rates    : 5000000 / 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: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%
```

Keep-alive Timeout Values for Different FC Protocols

When FICON traffic will flow over an extension tunnel, consider the following items when configuring the keep-alive timeout value (KATOV):

- A tunnel that carries FICON traffic requires a KATOV of less than or equal to 1 second for each circuit added to a tunnel.
- If the tunnel is created first with the FICON flag, then the KATOV for all added circuits will be 1 second (recommended value for FICON configurations).
- If the tunnel is created with one or more circuits, and the tunnel is modified to be a FICON tunnel, then the circuits that were previously created must be modified to have the correct KATOV.
- Set the circuit KATOV to the same value on both ends. If local and remote circuit configurations do not match, the tunnel will use the shorter duratoin of the configured values.
- For normal extension tunnel operations over tunnels transporting FICON traffic, the KATOV for all circuit members of a VE_Port (tunnel) must be less than the overall I/O timeout for all FC exchanges. If the FC I/O timeout value is less than the KATOV, then inputs and outputs will time out over all available circuits without being retried.
- The default timeout values are not the same for every supported platform. The value for the Brocade SX6 blade and Brocade 7840 switch is 6 seconds. The Brocade FX8-24 blade value is 10 seconds. The default value will also differ depending on whether FICON emulation is enabled.

  NOTE

  The Brocade 7810 switch does not support FICON or FICON Emulation.

The KATOV should be based on application requirements. Check with your FC initiator or IP initiator providers to determine the appropriate KATOV for your application. The sum of KATOVs for all circuits in a tunnel should be close to the overall FC initiator I/O timeout value. As an example, a mirroring application has a 6-second I/O timeout. There are three circuits belonging to the VE_Port (3 circuit members in the tunnel). Set the KATOV to 2 seconds on each circuit. This will allow for maximum retries over all available circuits before an I/O is timed out by the initiator.

Changing the KATOV can be disruptive.

Refer to the *Brocade Fabric OS Command Reference* for additional information on the on option format and value range available with the **portcfg fcipcircuit** command.

1. Use the **portcfg fcipcircuit modify** command to change the KATOV value.

   The following example modifies circuit 0 on VE_Port 24 and changes the KATOV to 2 seconds (2000 milliseconds).

```
switch:admin> portcfg fcipcircuit 24 modify 1 -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 **portcfgshow fcipcircuit** command to display the tunnel configuration.

   The following example displays the tunnel values for tunnel 24 with the KATOV pointed out. (The output is truncated.)

```
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            << KATOV
  Metric               : 0

. . .
```

Configuring Emulation Features on Tunnels

FICON emulation supports FICON traffic over IP WANs using FCIP as the underlying protocol.

> NOTE
> The Brocade 7810 Switch does not support FICON Emulation.

FICON emulation can be extended to support performance enhancements for specific applications through use of the following licensed features:

- IBM z/OS Global Mirror emulation (formerly eXtended Remote Copy or XRC)
- FICON tape emulation (tape read and write pipelining)

Using the FICON emulation features requires the Advanced FICON Acceleration licenses installed on each slot where FICON emulation is configured.

Refer to the *Brocade FICON Administration Guide* for more information on FICON emulation and emulation features on a tunnel.

Configuring Compression Options

Compression can be set at the tunnel level and it can be set at the protocol level for FC and IP. The switch or blade must be in Hybrid mode (both FCIP and IP Extension enabled) to set the individual protocol-level compression. Changing the compression level is disruptive.

NOTE
Compression mode must be set the same on each end of the tunnel.

Follow the guidelines in the following table for assigning explicit compression levels for tunnels on the Brocade 7810 switch, the Brocade 7840, and the Brocade SX6.

NOTE
The Brocade 7810 Switch does not support fast deflate compression mode.

TABLE 30 Assigning Compression Levels

Total Tunnel Bandwidth on a DP	Compression Level
More than 4 Gb/s	Fast deflate
2 Gb/s to 4 Gb/s	Deflate
2 Gb/s or less	Aggressive deflate

The enhancements for IP Extension allow you to configure compression on the tunnel at a protocol level. The compression options override the main tunnel compression level and set the compression for the specified protocol to the desired mode. The available modes depend on the protocol, whether FC or IP.

TABLE 31 IP Extension Hybrid Mode Protocol Compression Choices

Compression Level	FC Protocol Support	IP Protocol Support
Fast deflate	Yes	No
Deflate	Yes	Yes
Aggressive deflate	Yes	Yes

The following table shows the compression choices for the Brocade FX8-24 Blade.

TABLE 32 Brocade FX8-24 Compression Choices

Total Effective Tunnels FC Side	Compression Level
More than 2Gb/s	Standard
More than 512Mb/s and less than or equal to 2Gb/s	Moderate
Equal to or less than 512Mb/s	Aggressive

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

2. Use the `portcfg fciptunnel --compression` command to set the tunnel compression mode.

   The following example uses the `extncfg` command to verify that the switch is in FCIP mode and the `portcfg fciptunnel modify` command to change the tunnel compression.

```
switch:admin> extncfg --show
APP Mode is FCIP
VE-Mode: configured for 10VE mode.

FCIP_Remote:admin> portcfg fciptunnel 24 modify --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
```

3. Use the `portcfg fciptunnel --ip-compression` command to set the IP protocol compression mode. The switch must be in Hybrid mode.

   The following example uses the `extncfg` command to verify that the switch is in Hybrid mode and the `portcfg fciptunnel modify` command to change the tunnel compression for IP protocol.

```
switch:admin> extncfg --show
APP Mode is HYBRID (FCIP with IPEXT)
VE-Mode: configured for 10VE mode.
e: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
```

Configuring WAN on Tunnels

Configuring WAN on tunnels includes setting values for the following:

- Bandwidth values for Adaptive Rate Limiting (ARL) (See Configuring ARL on page 127.)
- Quality of Service (QoS) (See Configuring QoS Priorities over a Tunnel on page 128.)
- Virtual LAN (VLAN) (See Configuring VLANs on page 106.)
- Differentiated Services Code Point (DSCP) (See Configuring DSCP on page 132.)

QoS refers to policies for handling differences in data traffic. These policies are based on data characteristics and delivery requirements. For example, ordinary data traffic is tolerant of delays and dropped packets, but real-time voice and video data are not. QoS policies provide a framework for accommodating these differences in data as it passes through a network.

> **NOTE**
> 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. Refer to the *Brocade Fabric OS Administration Guide* for additional information.

QoS for Fibre Channel traffic is provided through internal QoS priorities. Those priorities can be mapped to TCP/IP network priorities using zone name prefixes and VCs. The different priority TCP sessions can be marked upon egress. The TCP marking is done at the IP layer using Layer 3 DSCP or at the Ethernet layer within the 802.1Q tag header using 802.1P. There are two options for TCP/IP network-based QoS:

- DSCP

- VLAN tagging and Layer 2 Class of Service (L2CoS)

You can configure QoS, DSCP, and VLAN tagging at the tunnel and circuit level for data path traffic.

Configuring ARL

Adaptive Rate Limiting (ARL) is configured on a per-circuit basis because each circuit can have available different amounts of bandwidth. Any single circuit is limited to 10 Gb/s, unless the hardware imposes a lower bandwidth. ARL never attempts to exceed the maximum configured value and reserves at least the minimum configured value.

The ARL minimum and maximum bandwidth is configured when a circuit is created or modified. The minimum bandwidth is considered the floor and the maximum bandwidth is the ceiling. See ARL Considerations on page 28 for information about bandwidth values when configuring ARL.

On the Brocade 7840 Switch, the Brocade 7810 Switch, and the Brocade SX6 Blade, you can configure the type of ARL algorithm for backing off the traffic. The default is automatic and the ARL logic determines the best approach. The ARL choices for these platforms are as follows.

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

On the Brocade FX8-24 Blade, when the minimum and maximum bandwidth differ, the ARL algorithm performs a static reset to the transmission floor, and then ramps the traffic back up.

NOTE
Best practice is to configure the minimum and maximum bandwidth values to the same value if you are not using ARL.

1. Connect to the switch and log in using an account assigned to the admin role.
2. When you create a tunnel, circuit 0 is automatically created. Use the `portcfg fciptunnel` command to create a tunnel and assign upper and lower bandwidth values for ARL.

   The following example creates a tunnel and configures circuit 0 with lower and upper ARL bandwidth values of 3 Gb/s and 5 Gb/s. With an upper limit of 5 Gb/s on circuit 0, a 10-Gb/s tunnel can support an additional 5 Gb/ps.

   ```
   switch:admin> portcfg fciptunnel  24 create -b 3000000 -B 5000000
   Operation Succeeded
   ```

3. Use the `portcfg fcipcircuit` command to create an additional circuit on tunnel 24 and assign ARL bandwidth values.

   The following example modifies circuit 0 to lower the ARL bandwidth values and creates circuit 1 with ARL bandwidth values. Circuit 1 is modified with the ARL algorithm set to step-down.

   ```
   switch:admin> portcfg fciptunnel 24 modify -b 2500000 -B 3000000
   Operation Succeeded

   switch:admin> portcfg fcipcircuit 24 create 1 -b 1500000 -B 2000000
   Operation Succeeded

   switch:admin> portcfg fcipcircuit 24 modify 1 --arl-algorithm step-down
   Operation Succeeded
   ```

4. Use the `portshow fcipcircuit` command to display circuit information.

   The following example shows the information for circuits on tunnel 24 with the configured bandwidth values. The display is truncated. Notice the Online Warning, because the local and remote ARL bandwidth values are different. The ARL algorithm type is shown for circuit 1.

```
switch:admin> portshow fcipcircuit 24 -d

 Circuit 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: 2500000 / 3000000 kbps
  Peer Comm Rates      : 5000000 / 5000000 kbps
  Actual Comm Rates    : 2500000 / 3000000 kbps
  Keepalive (Cfg/Peer) : 6000 (6000 / 6000) ms
  Metric               : 0
  Connection Type      : Default
  ARL-Type             : Auto
[...]

 Circuit 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: 2000000 / 2000000 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
[...]
```

Configuring QoS Priorities over a Tunnel

Per-Priority TCP QoS (PP-TCP-QoS) prioritizes FC traffic flows between initiators and targets within a tunnel to optimize bandwidth and performance. Each circuit handles one of the following priority traffic types:

- F class - F class is the highest priority, and is assigned bandwidth as needed at the expense of lower priorities, if necessary. This is referred to as strict priority.
- QoS high - The default priority value is 50 percent of the available bandwidth.
- QoS medium - The default value is 30 percent of the available bandwidth.
- QoS low - The default value is 20 percent of the available bandwidth.

QoS high, medium, and low priority traffic are assigned a percentage of available bandwidth based on priority level. QoS priority is based on the Virtual Circuit (VC) that carries data into the DP complex. For example, if data enters on a high VC, it is placed on a high TCP connection; if it enters on a low VC, then it is placed on the low TCP circuit. Data is assigned to the proper VC based on zone name prefix.

The following figure illustrates the internal architecture of TCP connections that handle PP-TCP-QoS. Note that this illustrates a tunnel containing a single circuit only that is in FCIP mode. When the tunnel is in Hybrid mode (FCIP and IP extension), the three QoS priority tunnels exist for both the FCIP traffic and the IP extension traffic, a total of six QoS VCs.

FIGURE 19 TCP Connections for Handling QoS

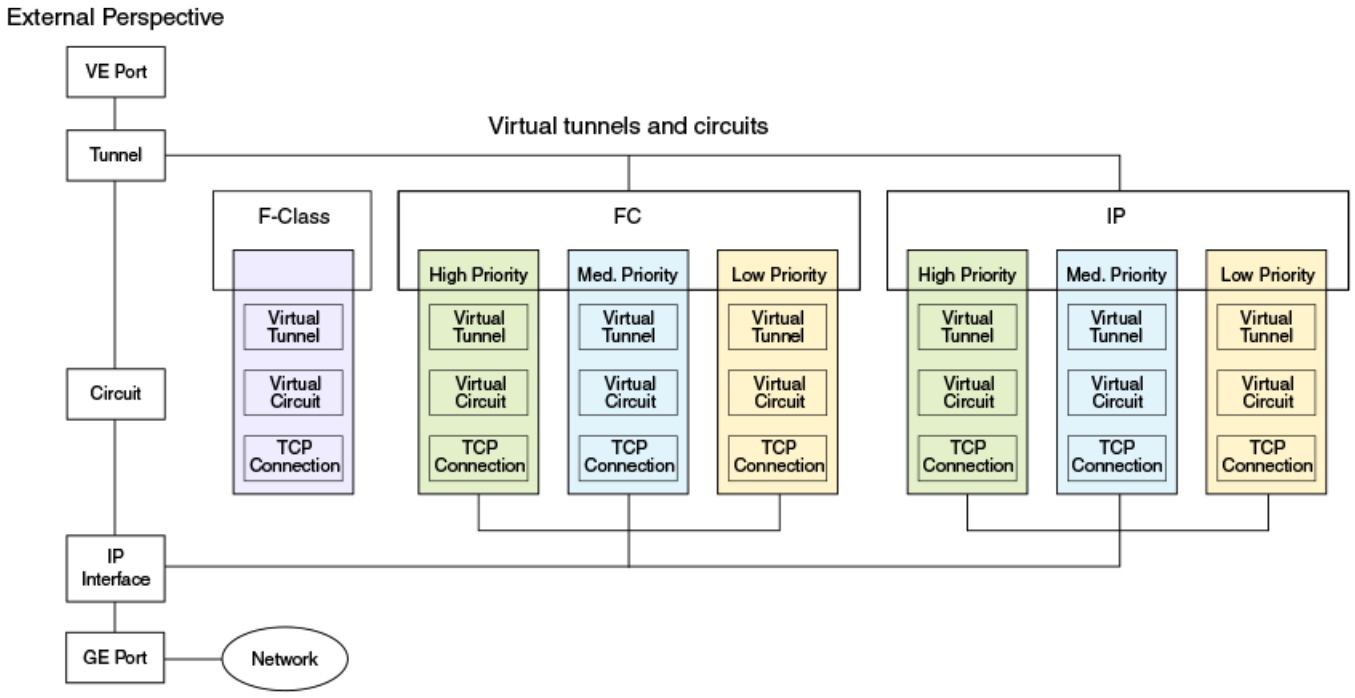

NOTE

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. Refer to the *Brocade Fabric OS Administration Guide* for additional information.

You can modify the default QoS priority values on Brocade extension switches and blades. This action only changes the QoS priority distribution in the tunnel; it does not reconfigure the fabric.

When the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade is configured in Hybrid mode, QoS traffic can be prioritized by protocol with a percentage assigned to FC and a percentage assigned to IP. (Recall that the Brocade 7810 Switch only operates in Hybrid mode.) The initial setting is 50/50. For each protocol, FC and IP, you can assign a percentage of the bandwidth to each of high, medium, and low levels. When configuring QoS percentages, remember the following:

- The three values (high, medium, low) must equal 100 percent.
- A minimum of 10 percent is required for each level.
- QoS priority settings must be the same on each end of the tunnel.

NOTE

Priorities are enforced only when there is congestion on the network. If there is no congestion, all traffic is handled at the same priority.

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

2. Use the `portcfg fciptunnel --distribution` command to change the default bandwidth values.

   The following example modifies the protocol bandwidth distribution to 60 percent for FC and 40 percent for IP. When you change bandwidth distribution, the QoS values are reset to their default.

   NOTE
   This command is disruptive.

```
switch:admin> portcfg fciptunnel 24 modify --distribution protocol,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
```

3. Use the `portcfg fciptunnel --qos-bw-ratio` command to change the ratio for high, medium, and low.

   The following example modifies QoS bandwidth ratio in an FCIP-only tunnel to 60 percent, 25 percent, and 15 percent for high, medium, and low. The three values total 100 percent.

```
switch:admin> portcfg fciptunnel 24 modify --qos-bw-ratio 60,25,15
Operation Succeeded
```

   The following example modifies QoS bandwidth ratio in a hybrid tunnel (FCIP and IP). The first triplet of values is for FCIP and the second triplet is for IP. Each triplet totals 100 percent.

```
switch:admin> portcfg fciptunnel 24 modify --qos-bw-ratio 60,30,10,50,30,20
Operation Succeeded
```

4. Use the `portshow fciptunnel` command to display the configured QoS values.

   The following example shows the QoS values for tunnel 24 configured to 60 percent, 25 percent, and 15 percent. The output is truncated.

```
switch:admin> portshow fciptunnel 24

 Tunnel: VE-Port:24 (idx:0, DP0)
 ======================================================
  Oper State           : Online
  TID                  : 24
  Flags                : 0x00000000
  IP-Extension         : Disabled
  Compression          : None
  QoS BW Ratio         : 60% / 25% / 15%
[...]
```

5. Use the `portshow fciptunnel` command to display the virtual circuits.

   The following example shows the QoS virtual circuit information. Over each physical circuit, there is a virtual circuit for each of the high, medium, and low QoS tunnels, and a VC for the F-class control data. In this example, there is only a single physical circuit configured, circuit 0. The switch is in FCIP mode, not Hybrid mode.

```
switch:admin>  portshow fciptunnel -q -c

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
-------------------------------------------------------------------------------
 24    -         Up       c--------   2d3h2m    0.00    0.00   3       -       -
 24    0 ge2     Up       ----ah--4   2d3h2m    0.00    0.00   3     0/5000  0/-
 24    -         Up       h--------   2d3h2m    0.00    0.00   3       -       -
 24    0 ge2     Up       ----ah--4   2d3h2m    0.00    0.00   3  2500/5000  0/-
 24    -         Up       m--------   2d3h2m    0.00    0.00   3       -       -
 24    0 ge2     Up       ----ah--4   2d3h2m    0.00    0.00   3  1500/5000  0/-
 24    -         Up       l--------   2d3h2m    0.00    0.00   3       -       -
 24    0 ge2     Up       ----ah--4   2d3h2m    0.00    0.00   3  1000/5000  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
```

Modifying QoS Default Priority Values on a FX8-24 Blade

You can modify the default QoS priority values for high, medium, and low on the Brocade FX8-24 blade. This action only changes the QoS priority distribution in the tunnel and does not reconfigure the fabric. Modifying the QoS default priority values affects the overall priorities for QoS traffic in the tunnel.

> NOTE
> Priorities are enforced only when there is congestion on the network. If there is no congestion, all traffic is handled at the same priority.

For information about changing the QoS priority values at the protocol level, which applies to the Brocade SX6 and Brocade 7840, see Configuring QoS Priorities over a Tunnel on page 128.

Change the priority percentages on 16-Gb/s and 32-Gb/s extension platforms using the optional `percentage` tunnel argument for the`portcfg fciptunnel create` and `portcfg fciptunnel modify` commands. When configuring QoS percentages for each level, remember the following:

- The three values must equal 100 percent.
- A minimum of 10 percent is required for each level.
- QoS priority settings must be the same on each end of the tunnel.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg fciptunnel` command to create or modify a tunnel and change the default QoS priority values.

   The following command sets the QoS priority ratios on VE_Port 12 to high (50 percent), medium (40 percent) and low (10 percent) priorities, respectively.

```
switch:admin> portcfg fciptunnel 1/12 create --qos 50,40,10
```

3. Use the `portshow fciptunnel` command to display tunnel values.

   The following command displays details of the tunnel configuration, including the QoS percentages for high, medium, and low priorities.

```
switch:admin> portshow fciptunnel 1/12
```

Configuring DSCP

Layer 3 Class of Service Differentiated Services Code Point (DSCP) refers to a specific implementation for establishing QoS policies as defined by RFC 2475. DSCP uses six bits of the Type of Service (TOS) field in the IP header to establish up to 64 different values to associate with data traffic priority.

DSCP settings are useful only if IP routers are configured to enforce QoS policies uniformly within the network. IP routers use the DSCP value as an index into a Per Hop Behavior (PHB) table. Control connections and data connections can be configured with different DSCP values. Before configuring DSCP settings, determine if the IP network you are using implements PHB, and consult with the WAN administrator to determine the appropriate DSCP values.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg fcipcircuit` command to modify the DSCP marking.

   The following command modifies circuit 0 on tunnel 24 and sets the DSCP high value to 1 for fibre channel.

   NOTE
   Only change DSCP values after consulting with your WAN administrator.

```
switch:admin> portcfg fciptunnel 24 modify --dscp-high 1
Operation Succeeded
```

3. Use the `portshow fcipcircuit` command to display the configured values for DSCP.

   The following example shows the DSCP high value on circuit 0 in tunnel 24 for FC. (The output is truncated.)

```
switch:admin> portshow fcipcircuit 24

 Circuit 24.0 (DP0)
 ==================================================
  Admin/Oper State     : Enabled
  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: 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)      : 1 / 8 / 4 (Ctrl:32)
  DSCP (IP:h/m/l)      : 0 / 0 / 0
[...]
```

Configuring Layer 2 Class of Service

VLAN traffic is routed using 802.1Q-compliant tags within an Ethernet frame. The tag includes a unique VLAN ID, and Class of Service (CoS) priority bits. The CoS priority scheme Layer 2 Class of Service (L2CoS) uses three Class of Service (CoS or 802.1P) priority bits, allowing eight priorities. Consult with your WAN administrator to determine usage.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg fcipcircuit` command to configure the L2CoS values.

   The following example modifies circuit 0 on tunnel 24 to configure L2CoS values.

   NOTE
   The circuit must have a VLAN tag before you can configure L2CoS values.

```
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` command to display the L2CoS values.

   The following example shows the L2CoS values for circuit 0 on tunnel 24 for FC traffic. The output is truncated.

```
 switch:admin> portshow fcipcircuit 24

 Circuit 24.0 (DP0)
 ==================================================
  Admin/Oper State     : Enabled / Online
  Flags                : 0x00000000
  IP Addr (L/R)        : 10.1.15.20 4/ge15 <-> 10.1.15.10
  HA IP Addr (L/R)     : 0.0.0.0  <-> 0.0.0.0
  Configured Comm Rates: 1000000 / 1000000 kbps
  Peer Comm Rates      : 1000000 / 1000000 kbps
  Actual Comm Rates    : 1000000 / 1000000 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              : 100
  L2Cos (FC:h/m/l)     : 16 / 8 / 4 (Ctrl:32)
  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
[...]
```

Configuring both DSCP and Layer 2 Class of Service

If a tunnel or circuit is VLAN tagged, both DSCP and Layer 2 Class of Service (L2CoS) may be tagged on ingress traffic unless the VLAN is end-to-end with no intermediate hops in the IP network. The following table shows DSCP priorities mapped to L2CoS priorities. This may be helpful when consulting with the network administrator. You can modify DSCP and L2CoS values for different priority traffic when configuring circuits for extension switches and blades.

TABLE 33 Default Mapping of DSCP Priorities to L2CoS Priorities

DSCP priority/bits	L2CoS priority/bits	Assigned to
7 / 000111	1 / 001	Medium QoS
11 / 001011	3 / 011	Medium QoS
15 / 001111	3 / 011	Medium QoS

TABLE 33 Default Mapping of DSCP Priorities to L2CoS Priorities (continued)

DSCP priority/bits	L2CoS priority/bits	Assigned to
19 / 010011	3 / 011	Medium QoS
23 / 010111	3 / 011	Medium QoS
27 / 011011	0 / 000	Class 3 Multicast
31 / 011111	0 / 000	Broadcast/Multicast
35 / 100011	0 / 000	Low Qos
39 / 100111	0 / 000	Low Qos
43 / 101011	4 / 100	High QoS
46 / 101110	7 / 111	Class F
47 / 101111	4 / 100	High QoS
51 / 110011	4 / 100	High QoS
55 / 110111	4 / 100	High QoS
59 / 111011	4 / 100	High QoS
63 / 111111	0 / 000	Reserved

Configuring Failover

There are two types of configuration supported:

- Active-active: Data will be sent on both 10-GbE ports to initiate weighted balancing of the batches across the trunk circuits.
- Active-passive: Data fails over using LLL to a passive circuit (one with a higher metric) if all active lower metric circuit paths fail.

You must establish a metric for failover circuits. If no metric is provided, circuit data will be sent through both ports, and the load will be balanced. Circuits have a default metric of 0. A metric of 1 is required for a standby (passive) circuit.

Active-active Configuration

The following example shows an active-active configuration in which two circuits are configured with the same metric, one circuit going over xge0 and the other circuit going over the crossport using xge1 as the external port. The metric values of both the circuits are the same (default value), so both circuits send data. The load is balanced across these circuits. The effective bandwidth of the tunnel in this example is 2 Gb/s.

1. Configure an IP address on interface xge0.

```
portcfg ipif 8/xge0 create 192.168.11.20/24 mtu 1500
```

2. Configure an IP address on crossport interface xge1.

```
portcfg ipif 8/xge1 create 192.168.10.10/24 mtu 1500 -x
```

3. Create a tunnel with one circuit going over xge0.

```
portcfg fciptunnel 8/22 create --remote-ip 192.168.11.20 --local-ip 192.168.11.21 -b 2750000 -B
2750000
```

4. Add another circuit, going over crossport xge1, to the tunnel.

```
portcfg fcipcircuit 8/22 create 1 --remote-ip 192.168.10.10 --local-ip 192.168.10.11 -b 1000000 -B
1000000
```

5. Display local and crossport interface details for xge0.

```
portshow ipif 8/xge0
```

NOTE
If the source and destination addresses are on different subnets, you must configure IP routes to the destination addresses. See Configuring IP Route on page 104.

Active-passive Configuration

The following example shows an active-passive configuration in which two circuits are configured with different metrics, one circuit going over xge0 and the other circuit going over the crossport using xge1 as the external port. In this example, circuit 1 is a failover circuit because it has a higher metric. When circuit 0 goes down, the traffic is failed over to circuit 1. The effective bandwidth of the tunnel in this example is 1 Gb/s.

1. Configure an IP address on interface xge0.

```
portcfg ipif 8/xge0 create 192.168.11.20/24  mtu 1500
```

2. Configure an IP address on crossport interface xge1.

```
portcfg ipif 8/xge1 create 192.168.10.10/24 mtu 1500 -x
```

3. Create a tunnel with one circuit going over xge0.

```
portcfg fciptunnel 8/22 create --remote-ip 192.168.11.21 --local-ip 192.168.11.20 -b 2750000 -B
2750000 --metric 0
```

4. Add another circuit, going over crossport xge1, to the tunnel.

```
portcfg fcipcircuit 8/22 create 1 --remote-ip 192.168.10.10 --local-ip 192.168.10.11 -b 1000000 -B
1000000 --metric 1
```

5. Display local and crossport interface details for xge0.

```
portshow ipif 8/xge0
```

NOTE
If the source and destination addresses are on different subnets, you must configure IP routes to the destination addresses. See Configuring IP Route on page 104.

Configuring Failover Groups

With circuit failover groups, you can better control which metric 1 circuits will be activated if a metric 0 circuit fails. To create circuit failover groups, you define a set of metric 0 and metric 1 circuits that are part of the same failover group. When all metric 0 circuits in the group fail, metric 1 circuits will take over data transfer, even when there are metric 0 circuits still active in other failover groups.

Typically, you would only define one metric 0 circuit in the group so that a specific metric 1 circuit will take over data transfer when the metric 0 circuit fails. This configuration prevents the problem of the tunnel operating in a degraded mode, with fewer than the defined circuits, before multiple metric 0 circuits fail.

Use the `portcfg fciptunnel` or `portcfg fcipcircuit` commands to configure the circuit metric and the failover group ID for a circuit. The metric value can be 0 or 1. The failover group ID is a vaule between 0 and 9.

The following steps modify an existing tunnel and create new circuits. The result is two failover groups for tunnel 4/16 that contain two circuits each. Note that circuit 0 is typically created automatically when the tunnel is created.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portshow ipif` command to verify IP addresses.

   The following example shows the IP addresses configured on slot 4, port GE15, DP0 for a Brocade SX6 Blade. This is the local switch.

```
Local_switch:admin> portshow ipif 4/ge15.dp0

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 4/ge15.dp0   10.1.15.20                    / 24   1500  0     U R M I
 4/ge15.dp0   10.1.15.21                    / 24   1500  0     U R M I
 4/ge15.dp0   10.1.15.22                    / 24   1500  0     U R M I
 4/ge15.dp0   10.1.15.23                    / 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 IP addresses configured on slot 4, port GE15, DP0 for a Brocade SX6 Blade. This is the remote switch.

```
Remote_switch:admin>  portshow ipif ge15.dp0

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge15.dp0     10.1.15.10                    / 24   1500  0     U R M I
 ge15.dp0     10.1.15.11                    / 24   1500  0     U R M I
 ge15.dp0     10.1.15.12                    / 24   1500  0     U R M I
 ge15.dp0     10.1.15.13                    / 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` command on the switch to create four circuits and configure the metric and failover-group values.

   The following example modifies circuit 0 and creates circuits 1 through 3, and assigns metric and failover values on the local switch.

```
Local_switch:admin> portcfg fciptunnel 4/16 create -S 10.1.15.20 -D 10.1.15.10 -b 1000000 -B 1000000
Local_switch:admin> portcfg fcipcircuit 4/16 modify 0 --metric 0 --failover-group 0

!!!! 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

Local_switch:admin> portcfg fcipcircuit 4/16 create 1 -S 10.1.15.21 -D 10.1.15.11 -x 0 -g 1 -b
1000000 -B 1000000
Operation Succeeded
Local_switch:admin> portcfg fcipcircuit 4/16 create 2 -S 10.1.15.22 -D 10.1.15.12 -x 1 -g 0 -b
1000000 -B 1000000
Operation Succeeded
Local_switch:admin> portcfg fcipcircuit 4/16 create 3 -S 10.1.15.23 -D 10.1.15.13 -x 1 -g 1 -b
1000000 -B 1000000
Operation Succeeded
```

   The following example modifies circuit 0 and creates circuits 1 through 3, and assigns metric and failover values on the remote switch. Notice that the source and destination circuit endpoints on the remote switch are the reverse of the local switch. The failover group ID must match at each end of the circuit.

```
Remote_switch:admin> portcfg fciptunnel 24 create -S 10.1.15.10 -D 10.1.15.20 -b 1000000 -B 1000000
Local_switch:admin> portcfg fcipcircuit 24 modify 0 --metric 0 --failover-group 0

!!!! 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

Remote_switch:admin> portcfg fcipcircuit 24 create 1 -S 10.1.15.11 -D 10.1.15.21 -x 0 -g 1 -b
1000000 -B 1000000
Operation Succeeded
Remote_switch:admin> portcfg fcipcircuit 24 create 2 -S 10.1.15.12 -D 10.1.15.22 -x 1 -g 0 -b
1000000 -B 1000000
Operation Succeeded
Remote_switch:admin> portcfg fcipcircuit 24 create 3 -S 10.1.15.13 -D 10.1.15.23 -x 1 -g 1 -b
1000000 -B 1000000
Operation Succeeded
```

4. Use the `portcfg fciptunnel` command to verify that the metric and failover values are configured correctly at each end of the circuit.

   The following example shows the `portshow fciptunnel all -c -h` command and that all circuits on the local switch are up, indicating that both sides of the tunnel are communicating.

```
Local_switch:admin> portshow fciptunnel all -c -h

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
-------------------------------------------------------------------------------
 4/16  -         Up      -M-------   55m47s    0.00    0.00  17      -       -
 4/16  0 4/ge15  Up      ----a---4    44m2s    0.00    0.00  17  1000/1000  0/-
 4/16  1 4/ge15  Up      ----a---4    40m3s    0.00    0.00   1  1000/1000  0/1
 4/16  2 4/ge15  Up      ----a---4      51s    0.00    0.00   5  1000/1000  1/-
 4/16  3 4/ge15  Up      ----a---4      14s    0.00    0.00   2  1000/1000  1/1
-------------------------------------------------------------------------------
 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
```

Configuring Spillover

To configure spillover, use the `portcfg` command. If you do not configure spillover, the default action is to use failover.

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

```
switch:admin> portcfg fciptunnel 24 modify -L spillover
Operation Succeeded
```

   The `-L` option is the short form for `--load-leveling`.

3. Use the `portshow fciptunnel` command to display the current load leveling algorithm. Because the value is negotiated from end-to-end, the configured value, peer value, and actual (negotiated) value are displayed. (The output is truncated.)

```
switch:admin> portshow fciptunnel 24

 Tunnel: VE-Port:24 (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 of the tunnel is configured with failover. (The output is truncated.)

```
switch:admin > portshow fciptunnel 24
Tunnel: VE-Port:24 (idx:0, DP0)
 ================================================
  Oper State           : Online Warning
  ...
  Load-Level (Cfg/Peer): Failover (Failover / Spillover)
  ...
  Configuration Warnings:
     Load Leveling (failover|spillover)
```

5. The current PDU/Byte counts that are maintained for each circuit can be used to verify spillover usage of metric 1 circuits. The `--qos/-q` option provides a more detailed view as to which QoS the spillover is occurring on. The following examples are of commonly used commands, with some providing information about PDU/byte counts and others showing throughput. Only the commands are shown because output will vary depending on the configuration and traffic conditions.

```
switch:admin> portshow fciptunnel -c
switch:admin> portshow fciptunnel -c -q
switch:admin> portshow fciptunnel 24 -c
switch:admin> portshow fcipcircuit 24 -c
switch:admin> portshow fcipcircuit 24 -c -q
```

Example of Spillover Circuits at Various Rates

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

```
switch:admin> portshow fciptunnel -c

Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
-------------------------------------------------------------------------------
24    -         Up      --------I   17m52s  309.05    0.00   2      -       -
24    0 ge2     Up      ----a---4   17m51s  309.05    0.00   2  3000/3000  0/-
24    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
-------------------------------------------------------------------------------
24    -         Up      --------I   20m40s  306.29    0.00   2      -       -
24    0 ge2     Up      ----a---4   20m40s  285.33    0.00   2  2400/2400  0/-
24    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
-------------------------------------------------------------------------------
24    -         Up      --------I   23m30s  305.17    0.00   2      -       -
24    0 ge2     Up      ----a---4   23m30s  249.70    0.00   2  2100/2100  0/-
24    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
-------------------------------------------------------------------------------
24    -         Up      --------I   24m52s  304.78    0.00   2      -       -
24    0 ge2     Up      ----a---4   24m51s  178.11    0.00   2  1500/1500  0/-
24    1 ge3     Up      ----a---4   24m52s  126.67    0.00   2  3000/3000  1/-
-------------------------------------------------------------------------------
```

Configuring VE_Ports to Persistently Enable

NOTE

The Brocade 7810 Switch does not support VE operating mode (i.e., --ve-mode [10VE|20VE]).

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

2. Use the `portshow` command to show the status of a port.

   This example shows the status of VE_Port 24. You can see the status of persistent disable in the output. (The "<<<<<" is not part of the actual output.)

```
switch:admin> portshow 24
portIndex:  24
portName: 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: No
```

3. Use the `portcfgpersistentenable` command to disable any VE_Ports that you will use in the tunnel configuration.

   This example enables VE_Port 24.

```
switch:admin> portcfgpersistentenable 24
switch:admin>
```

   There is no additional response or confirmation after entering the command.

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

   This example shows the status of VE_Port 24. You can see the port is online and no longer disabled. (The "<<<<<" is not part of the actual output.)

```
switch:admin>  portshow 24
portIndex:  24
portName: port24
portHealth: Not Monitored

Authentication: None
portDisableReason: None                      <<<<<
portCFlags: 0x1
portFlags: 0x4903       PRESENT ACTIVE VIRTUAL E_PORT G_PORT U_PORT LOGICAL_ONLINE LOGIN LED
LocalSwcFlags: 0x0
portType:  12.0
portState: 1    Online                       <<<<<
Protocol: FC
portPhys:  255  N/A     portScn:   16   E_Port
port generation number:    24
state transition count:    21

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: No
```

5. If the switch is in FMS mode, use the `portenable` command to enable a port.

   The following example enables port 24, a VE_Port on a Brocade 7840 switch.

```
switch:admin> portenable 24
switch:admin>
```

   There is no additional response or confirmation after entering the command.

Verifying Tunnel Configuration

After you have created local and remote configurations, verify that the tunnel and circuit parameters are correct using the `portshow fciptunnel` and the `portshow fcipcircuit` commands. Refer to the *Brocade Fabric OS Command Reference* for a description of the command syntax and output.

By comparing the command output at each end of the tunnel, you can usually identify configuration issues. The display will often indicate if there is a mismatch between parameters.

1. Use the `portshow fciptunnel` command to display information about the tunnel. The examples show various uses of the command.

   This example shows basic use of `portshow fciptunnel`. The tunnel is up, but has a warning.

```
switch:admin> portshow fciptunnel

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
-------------------------------------------------------------------------------
 4/16  -         UpWarn  ---------    3d19m    0.00    0.00  17      -       -
-------------------------------------------------------------------------------
 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
```

2. Use the `portshow fciptunnel` command to show additional information about a specific tunnel (or VE_Port).

   This example shows the `portshow fciptunnel` *`slot/port`* command, which expands the basic `portshow` command information. The tunnel is up, but has a warning. (The "<<<<<" points out the warning; it does not appear in the actual display.)

```
switch:admin> portshow fciptunnel 4/16

 Tunnel: VE-Port:4/16 (idx:0, DP0)
 ====================================================
  Oper State           : Online Warning                  <<<<<
  TID                  : 80
  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            : 10:00:00:05:33:e7:c5:10
  Peer WWN             : 10:00:00:05:33:65:83:48
  RemWWN (config)      : 00:00:00:00:00:00:00:00
  cfgmask              : 0x0000001f 0x40000208
  Flow Status          : 0
  ConCount/Duration    : 17 / 7d6h50m
  Uptime               : 3d19m
  Stats Duration       : 3d19m
  Receiver Stats       : 2084 bytes / 14 pkts /    0.00 Bps Avg
  Sender Stats         : 1748 bytes / 13 pkts /    0.00 Bps Avg
  TCP Bytes In/Out     : 1883857860 / 1848343280
  ReTx/OOO/SloSt/DupAck: 1040 / 0 / 64 / 0
  RTT (min/avg/max)    : 1 / 1 / 1 ms
  Wan Util             : 0.0%
  TxQ Util             : 0.0%
```

3. Use the `portshow fcipcircuit` command to display specific circuit information.

   This example shows the `portshow fciptunnel` *slot/port* command with a specific circuit specified, 4.16/1. The configuration warning indicates a mismatch in the ARL settings at each end of the circuit. (The "<<<<<" points out the warning, it does not appear in the actual display.)

```
switch:admin> portshow fcipcircuit 4/16 1

 Circuit 4/16.1 (DP0)
 ====================================================
  Admin/Oper State     : Enabled / Online Warning             <<<<<
  Flags                : 0x00000000
  IP Addr (L/R)        : 10.1.15.21 4/ge15 <-> 10.1.15.11
  HA IP Addr (L/R)     : 0.0.0.0  <-> 0.0.0.0
  Configured Comm Rates: 1000000 / 1000000 kbps
  Peer Comm Rates      : 1000000 / 1000000 kbps
  Actual Comm Rates    : 1000000 / 1000000 kbps
  Keepalive (Cfg/Peer) : 6000 (6000 / 6000) ms
  Metric               : 0
  Connection Type      : Default
  ARL-Type             : Auto
  PMTU                 : Disabled
  SLA                  : (none)
  Failover Group       : 1
  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 0x00000c6f
  Configuration Warnings:                                  <<<<<
    ARL Type
  Flow Status          : 0
  ConCount/Duration    : 1 / 3d38m
  Uptime               : 3d13m
  Stats Duration       : 3d13m
  Receiver Stats       : 0 bytes / 0 pkts /    0.00 Bps Avg
  Sender Stats         : 0 bytes / 0 pkts /    0.00 Bps Avg
  TCP Bytes In/Out     : 777180488 / 758776120
  ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
  RTT (min/avg/max)    : 1 / 1 / 1 ms
  Wan Util             : 0.0%
```

4. Use the `portshow fcipcircuit` command to verify that no problems exist.

   This example shows the `portshow fciptunnel` *`slot/port`* command with a specific circuit specified, 4.16/0. The output indicates that the circuit is functioning and no configuration errors exist, by lack of any warning indicators.

```
switch:admin>  portshow fcipcircuit 4/16 0

 Circuit 4/16.0 (DP0)
 ====================================================
  Admin/Oper State     : Enabled / Online
  Flags                : 0x00000000
  IP Addr (L/R)        : 10.1.15.20 4/ge15 <-> 10.1.15.10
  HA IP Addr (L/R)     : 0.0.0.0  <-> 0.0.0.0
  Configured Comm Rates: 1000000 / 1000000 kbps
  Peer Comm Rates      : 1000000 / 1000000 kbps
  Actual Comm Rates    : 1000000 / 1000000 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 0x00000c0f
  Flow Status          : 0
  ConCount/Duration    : 17 / 7d7h7m
  Uptime               : 3d24m
  Stats Duration       : 3d24m
  Receiver Stats       : 0 bytes / 0 pkts /    0.00 Bps Avg
  Sender Stats         : 0 bytes / 0 pkts /    0.00 Bps Avg
  TCP Bytes In/Out     : 1917176420 / 1899892860
  ReTx/OOO/SloSt/DupAck: 1040 / 0 / 64 / 0
  RTT (min/avg/max)    : 1 / 1 / 1 ms
  Wan Util             : 0.0%
```

Configuring Extension Hot Code Load

Extension Hot Code Load (eHCL) is a feature that allows firmware to be upgraded without disrupting traffic. It is supported on the Brocade 7840 Switch and the Brocade 7840 Switch. The Brocade 7810 Switch does not support eHCL. eHCL is supported in both FCIP and IPEX.

eHCL uses the dual DP complexes, DP0 and DP1, on the Brocade 7840 Switch or Brocade SX6 Blade to provide uninterrupted traffic flow when the switch firmware is upgraded. The primary circuit, or main tunnel, is the circuit that is created between the DP0 on the local switch and DP0 on the remote switch. The backup tunnel is a circuit created from the local switch DP1 to the remote switch DP0. When a firmware upgrade occurs, the local DP0 is brought down gracefully and its traffic processing is moved to the local DP1. The local DP1 uses the backup tunnel to the remote DP0 to continue the uninterrupted traffic flow. When the local DP0 comes back up, traffic processing is moved from the local DP1 backup tunnel to the DP0 main tunnel.

NOTE
For releases before Fabric OS 8.1.0, eHCL is disruptive to any IP Extension traffic.

Additional considerations for eHCL configuration are as follows:

- The primary circuit is referred to as the main tunnel (MT). The backup circuit from local switch to remote switch is the local backup tunnel (LBT). The tunnel that points from the remote switch back to the local switch is the remote backup tunnel (RBT).

- The IPIFs for the MT and LBT side of eHCL must be on different DPs.
- When you create or modify a circuit, only those circuits that specify a local and remote HA IP address are protected.
- If a circuit is not specifically configured for eHCL, it will go down during the firmware upgrade process.
- When the switch is in 10VE mode, you can configure up to five tunnel on each DP for eHCL. This allows a total of 15 addresses on the DP to use for eHCL: five each to use as endpoints for the main tunnel, local backup tunnel, and remote backup tunnel.
- When the switch is in 20VE mode, you can configure up to 10 tunnels on each DP for eHCL, for a total of 30 addresses to use for eHCL.
- Whether you use DP0 or DP1 as the main tunnel, the backup tunnel must be configured on the other DP.
- The tunnel and circuit parameters of the main tunnel are replicated on the LBT and RBT, including the circuit properties such as QoS markings, FastWrite, and FICON Acceleration.

To configure eHCL, perform the following steps.

1. Connect to the local switch and log in using an account assigned to the admin role.
2. Use the `portshow ipif` command to confirm the IPIF addresses for the main tunnel and backup tunnel are configured on DP0 and DP1.

   The following example shows IPIFs on the local switch. The address 192.168.5.2 on DP0 will be used for the main tunnel, and 192.168.5.12 will be used for the backup tunnel.

```
local_switch:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.5.2                   / 24   1500  0     U R M I
 ge2.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
```

   The following example shows IPIFs on the remote switch. The address 192.168.1.2 on DP0 will be used for the main tunnel, and 192.168.1.12 on DP1 will be used for the backup tunnel.

```
remote_switch:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.1.2                   / 24   1500  0     U R M I
 ge2.dp1      192.168.1.12                  / 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
```

3. Use the `portcfg fciptunnel` command to configure the tunnel and circuit 0. On each switch (local switch, remote switch), the destination address of both the main tunnel and backup tunnel is on DP0 of the other switch

   The following example configures tunnel 24 on the local switch with a main tunnel and backup tunnel.

```
local_switch:admin> portcfg fciptunnel 24 create -S 192.168.5.2 -D 192.168.1.2  --local-ha-ip
192.168.5.12 --remote-ha-ip 192.168.1.12 -b 5000000 -B 5000000
local_switch:admin>
```

   The following example configures tunnel 24 on the remote switch with a main tunnel and backup tunnel.

```
remote_switch:admin> portcfg fciptunnel 24 create -S 192.168.1.2 -D 192.168.5.2  --local-ha-ip
192.168.1.12 --remote-ha-ip 192.168.5.12 -b 5000000 -B 5000000
remote_switch:admin>
```

4. When the switch is in Hybrid mode (FCIP and IP Extension), use the `portcfg fciptunnel` command with the `--ipext enable` option.

   The following example configures tunnel 24 on the local switch with a main tunnel and backup tunnel in Hybrid mode.

```
local_switch:admin> portcfg fciptunnel 24 create -S 192.168.5.2 -D 192.168.1.2  --local-ha-ip
192.168.5.12 --remote-ha-ip 192.168.1.12 -b 5000000 -B 5000000 --ipext enable
local_switch:admin>
```

   The following example configures tunnel 24 on the remote switch with a main tunnel and backup tunnel in Hybird mode.

```
remote_switch:admin> portcfg fciptunnel 24 create -S 192.168.1.2 -D 192.168.5.2  --local-ha-ip
192.168.1.12 --remote-ha-ip 192.168.5.12 -b 5000000 -B 5000000 --ipext enable
remote_switch:admin>
```

5. Use the `portshow fcipcircuit` command to display the configured values for eHCL.

   The following example shows the high-availability (HA) configuration values. You can verify the endpoints for the main tunnel and backup tunnel on the local switch.

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

 Tunnel Circuit   AdminSt Flags   Local IP             Remote Ip
--------------------------------------------------------------------------------
 24    -         Enabled -M-------
 24    0 ge2     Enabled ----ah--4 192.168.5.2         192.168.1.2
 24    -         Enabled -R-------
 24    0 ge2     Enabled ----ah--4 192.168.5.2         192.168.1.12
 24    -         Enabled -L-------
 24    0 ge2     Enabled ----ah--4 192.168.5.12        192.168.1.2
--------------------------------------------------------------------------------
 Flags (tunnel): 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 endpoints for the main tunnel and backup tunnel that are configured on the remote switch. (The command uses short notation for the options.)

```
remote_switch:admin>  portshow fciptunnel -chC

 Tunnel Circuit   AdminSt Flags   Local IP             Remote Ip
--------------------------------------------------------------------------------
 24    -         Enabled -M-------
 24    0 ge2     Enabled ----ah--4 192.168.1.2         192.168.5.2
 24    -         Enabled -R-------
 24    0 ge2     Enabled ----ah--4 192.168.1.2         192.168.5.12
 24    -         Enabled -L-------
 24    0 ge2     Enabled ----ah--4 192.168.1.12        192.168.5.2
--------------------------------------------------------------------------------
 Flags (tunnel): 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 main tunnel. The R flag identifies the remote backup tunnel. It shows the local IP on DP0 as the endpoint of the remote backup tunnel. The L flag identifies the local backup tunnel. It shows the remote DP0 IP as the local backup tunnel endpoint.

6. Use the `portshow fciptunnel` command to monitor eHCL status.

   The following example shows the status of a correctly configured eHCL when no firmware is being upgraded.

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

Checking FCIP Tunnel HA Status.

Current Status         : Ready
CP Version             : v8.1.0
4/DP0 Status:
   State               : Online - Inactive
   Version             : v8.1.0
   Current FC HA Stage : IDLE
   Current IP HA Stage : IDLE
   IP SVI Swapped      : NO
   DP COMM Status      : N/A
4/DP1 Status:
   State               : Online - Inactive
   Version             : v8.1.0
   Current FC HA Stage : IDLE
   Current IP HA Stage : IDLE
   IP SVI Swapped      : NO
   DP COMM Status      : N/A

 Tunnel 24 HA configured and HA Online. Traffic will not be disrupted.
```

Configuring DP Complexes and eHCL Tunnels

When you configure interfaces and tunnels for Extension HCL, the main tunnel (MT) interface cannot be on the same data processor (DP) complex as the local backup tunnel (LBT). Because the Brocade 7840 switch shares Ethernet ports, you can assign an Ethernet port to a specific DP using the `portcfg ipif` command with the `ge_port.dp_num` option.

In the following figure, the Brocade 7840-A switch and the Brocade 7840-B switch both use their DP0 complexes as the endpoints for the main tunnel. Each switch uses its DP1 for the local backup tunnel. On each side, the local backup tunnel on DP1 points to the remote backup tunnel on DP0 on the remote switch. In the simplest eHCL configuration, a single circuit is configured between the local and remote switches and the circuit consists of the main tunnel, local backup tunnel, and remote backup tunnel endpoints.

FIGURE 20 Extension HCL Configuration with MT, LBT, and RBT

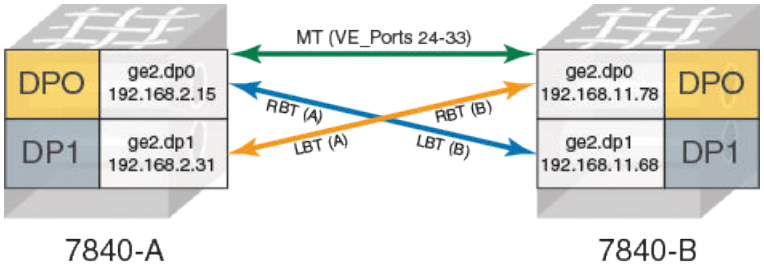

The DP complex creates the following TCP connections between 7840-A and 7840-B:

- 192.168.2.15 to 192.168.11.78: The main TCP connections for the MT between the switches
- 192.168.2.31 to 192.168.11.78: The connection from 7840-A LBT to 7840-B RBT
- 192.168.2.15 to 192.168.11.68: The connection between 7840-A RBT and 7840-B LBT

The following table shows the possible options for configuring MTs and LBTs on a Brocade 7840 Switch. The options show ports available in 20VE mode and 10VE mode. In 20VE mode, up to 10 HCL tunnels can be configured on each DP. In 10VE mode, the

limit is 5 eHCL tunnels on each DP. The VE_Ports used on the local and remote switches need not be bidirectional. That is, you can use one of the VE_Ports 24-33 on the local switch and one of the VE_Ports 34-43 on the remote switch for the main tunnel. GbE ports 0 and 1 (ge0, ge1) are 40GbE ports; the remaining ports are 10GbE ports.

TABLE 34 Brocade 7840 Extension HCL Considerations for VE_Port Configuration

7840-A Available VE_Ports in 20VE Mode / 10 VE Mode	MT Local IP Addresses Defined on	LBT Local IP Addresses Defined on	7840-B Available VE_Ports in 20VE Mode / 10VE Mode	RBT IP Addresses Defined on
24-33 / 24-28	One of ge0-ge17 on dp0 (such as ge2.dp0)	One of ge0-ge17 on dp1 (such as ge2.dp1)	24-33 / 24-28	One of ge0-ge17 on dp1 (such as ge2.dp1)
24-33 / 24-28	One of ge0-ge17 on dp0 (such as ge2.dp0)	One of ge0-ge17 on dp1 (such as ge2.dp1)	34-43 / 34-38	One of ge0-ge17 on dp0 (such as ge2.dp0)
34-43 / 34-38	One of ge0-ge17 on dp1 (such as ge2.dp1)	One of ge0-ge17 on dp0 (such as ge2.dp0)	34-43 / 34-38	One of ge0-ge17 on dp0 (such as ge2.dp0)
34-43 / 34-38	One of ge0-ge17 on dp1 (such as ge0.dp1)	One of ge0-ge17 on dp0 (such as ge0.dp0)	24-33 / 24-28	One of ge0-ge17 on dp1 (such as ge0.dp1)

Between two Brocade SX6 blades, the options for a similar SX6-A and SX6-B setup are shown in the following table.

TABLE 35 Brocade SX6 Extension HCL Considerations for VE_Port Configuration

SX6-A Available VE_Ports in 20VE Mode / 10VE Mode	MT Local IP Addresses Defined on	LBT Local IP addresses Defined on	SX6-B Available VE_Ports in 20VE Mode / 10VE Mode	RBT IP Addresses Defined on
16-25 / 16-20	One of ge0-ge17 on dp0 (such as ge2.dp0)	One of ge0-ge17 on dp1 (such as ge2.dp1)	16-25 / 16-20	One of ge0-ge17 on dp1 (such as ge2.dp1)
16-25 / 16-20	One of ge0-ge17 on dp0 (such as ge2.dp0)	One of ge0-ge17 on dp1 (such as ge2.dp1)	26-35 / 26-30	One of ge0-ge17 on dp0 (such as ge2.dp0)
26-35 / 26-30	One of ge0-ge17 on dp1 (such as ge2.dp1)	One of ge0-ge17 on dp0 (such as ge2.dp0)	26-35 / 26-30	One of ge0-ge17 on dp0 (such as ge2.dp0)
26-35 / 26-30	One of ge0-ge17 on dp1 (such as ge0.dp1)	One of ge0-ge17 on dp0 (such as ge0.dp0)	16-25 / 16-20	One of ge0-ge17 on dp1 (such as ge0.dp1)

Configuring IP Extension

IP Extension (IPEX) features are available when you configure the Brocade 7840 Switch, the Brocade 7810 Switch or the Brocade SX6 Blade to operate in Hybrid mode where you can configure most tunnel features associated with tunnels, and additionally configure IP Extension features. Most tunnel or circuit configurations performed with a switch not in Hybrid mode are carried over to Hybrid mode.

NOTE
The Brocade 7810 Switch operates only in Hybrid mode.

Before you begin to configure IP Extension features, consider the following:

- Follow the recommendations in Configuration Prerequisites on page 88 to prepare the WAN support. WAN support must be completed prior to configuring IP Extension features on the Brocade 7840 Switch, the Brocade 7810 Switch or the Brocade SX6 Blade.
- For the Brocade 7840 Switch or the Brocade SX6 Blade, you must configure the VE_Port operating mode for 10VE mode (the default).
- Determine the VE_Port numbers you want to use. The VE_Port numbers serve as tunnel IDs.
- FIPS is not supported when the Brocade 7840 Switch or Brocade SX6 Blade operate in Hybrid mode.

NOTE
Once enabled, FIPS cannot be disabled.

- Determine which Ethernet ports will be used for LAN connectivity. Ensure these ports are in the default switch and configured to LAN mode.

  NOTE
  The LAN ports do not participate in Layer 2 Ethernet switching on the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade. All traffic is terminated and sent over the tunnel, based on the configured traffic control list (TCL).

  NOTE
  The Brocade 7840 Switch and the Brocade SX6 Blade support a maximum of 1024 defined and 128 active TCLs. The Brocade 7810 Switch supports a maximum of 256 defined and 32 active TCLs.

- For the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade, the IP MTU size must be at least 1280. If the supported maximum IP MTU size in the network exceeds 9216, the IP MTU should be 9216.
- Identify the subnets that will be extended through the Brocade 784 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade. Assign IP addresses, subnet masks, and MTUs to be used as LAN gateways to the IPEX switch or blade.
- With the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade are operating in Hybrid mode and IP Extension features enabled, you must configure traffic control lists (TCLs).
- Determine how many additional circuits you want to create. You will need the source and destination IP addresses for each circuit, and the minimum and maximum rates for ARL, or the committed rate if not using ARL. You will need to know if you intend to assign metrics to circuits so that lower metric circuits fail over to circuits with higher metrics. For all circuits except circuit 0, these values are set by the **portCfg fcipcircuit create** command.
- You must consider the maximum limit of LAN connections allowed for each DP, which is 512 accelerated TCP connections per DP and 64 UDP connections for the Brocade 7840 switch and the Brocade SX6 blade; and 128 TCP connections per DP and 32 UDP connections for the Brocade 7810 switch.
- When planning for VE_Port and bandwidth usage, consider which VE_Port is hosted by which DP complex. This affects load balancing between each DP complex.

Configuration Steps for IP Extension Features

Perform the following major steps for configuring IP Extension features on extension switches. The WAN configuration steps are presented first, followed by LAN configuration steps.

IP Extension WAN Configuration

The major steps that apply to WAN configuration are as follows:

1. Configure the operating mode on the Brocade 7840 switch or the Brocade SX6 blade for 10VE mode (in Hybrid mode, 10VE is the default and 20VE is disallowed). This step is necessary if the switch or blade has an existing FCIP configuration.
2. Configure the Brocade 7840 switch or the Brocade SX6 blade to operate in *hybrid mode* using the `extncfg --app-mode hybrid` command. (Recall that the Brocade 7810 switch operates only in Hybrid mode.)

   NOTE
   Configuring the switch for Hybrid mode is disruptive. On a Brocade 7840 switch, it reboots and loads the hybrid image. On a Brocade SX6 blade, it reboots and the chassis is unaffected.

3. If an existing configuration is present, persistently disable the existing VE_Ports. See Configuring VE_Ports to Persistently Disable on page 120.
4. Create an IP interface (IPIF) for each circuit that you want on a port by assigning an IP address, netmask, and an IP MTU size to an Ethernet port using the `portCfg ipif` command. See Configuring IPIF on page 102. Note that this step applies to overall WAN configuration and is not specific to IP Extension LAN configuration.
5. Create one or more IP routes to a port, if required, using the `portCfg iproute` command. See Configuring IP Route on page 104. Note that this step applies to overall WAN configuration and is not specific to IP Extension LAN configuration.
6. Test the WAN IP connection using the `portCmd --ping` command. Note that this step applies to overall WAN configuration and is not specific to IP Extension LAN configuration. See Verifying IP Connectivity on page 107 for details.

   NOTE
   If a VLAN is present in the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade, it requires configuration on only the IPIF. Furthermore, the VLANs must match only the configured local Ethernet hop. See Configuring VLANs on page 106 for details.

   NOTE
   Stacked VLAN tagging is not supported on the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade. Stacked VLAN tagging is defined in IEEE 802.1ad.

7. Create IPsec policies before creating extension tunnels. Though optional, this step is recommended for security. See Configuring IPsec on page 111. Note that this step applies to overall WAN configuration and is not specific to IP Extension LAN configuration.
8. Create extension tunnels using the `portCfg fciptunnel` command. See Configuring Extension Tunnels for FCIP on page 119. Note that this step applies to overall WAN configuration and is not specific to IP Extension configuration.
9. Configure bandwidth distribution for IP Extension using the `portCfg fciptunnel` command. The bandwidth distribution determines how QoS traffic is distributed between FC and IP bandwidth.
10. Configure the compression mode for IP Extension using the `portCfg fciptunnel` command. Note that IP traffic compression cannot be set to fast deflate.
11. Create FCIP circuits (after circuit 0) and enable or disable features using the `portCfg fcipcircuit` command. See Configuring a Tunnel and Circuits for IP Extension on page 160 for information. Use the `portshow fciptunnel` and `portCfg fcipcircuit` commands to verify tunnel and circuit status before proceeding. See Verifying Tunnel Configuration on page 141 for additional information. Note that this step applies to overall WAN configuration and is not specific to IP Extension configuration.

IP Extension LAN configuration

The major steps that apply to LAN configuration for IP Extension are as follows:

1. Identify the subnets that will be extended through the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade. Assign IP addresses, subnet masks, and MTUs to be used as LAN gateways to the switch or blade.

   NOTE
   The local-side LANs that contain the end devices must be on different subnets. For example, you cannot have end-devices on DC-A subnet 10.0.0.0/24 communicating with end-devices DC-B subnet 10.0.0.0/24. Those are the same subnets on each side.

2. Configure a GE port for LAN mode using the `portCfgGe` command.

3. Configure the GE port for link aggregation group (LAG) operation using the `portcfg lag` command. (This step is optional.)
4. Configure a switch virtual interface (SVI) to provide IP access for LAN devices using the `portcfg ipif` command.
5. Configure the traffic control list (TCL) for the port using the `portcfg tcl` command. A TCL consists of a rule name, a priority, an input filter, and a target for the rule. See IP Extension and Traffic Control Lists on page 66 for details.
6. Persistently enable the VE_Ports.

Configuring Hybrid Mode for IP Extension Features

Configure the Brocade 7840 Switch or Brocade SX6 Blade to operate in Hybrid mode, which supports FCIP tunnel features and IP Extension features.

The Brocade 7840 Switch DP or the Brocade SX6 Blade DP cannot be configured in FIPS mode, which Hybrid mode does not support.

The Brocade 7840 Switch or Brocade SX6 Blade must be in 10VE mode before you can enable Hybrid mode. Changing VE modes is disruptive as it requires rebooting the switch. If you plan to use IP Extension, configure Hybrid mode during initial deployment. Because 20VE mode is unavailable in Hybrid mode, it is normal for the 20VE_Ports to be disabled in 10VE mode. The following table shows the 20VE and 10VE ports on the two platforms.

TABLE 36 20VE and 10VE Mode Ports

Product	20VE Mode VE_Ports	10VE Mode VE_Ports
Brocade 7840 switch	DP0 - 24 through 33 DP1 - 34 through 43	DP0 - 24 through 28 DP1 - 34 through 38
Brocade SX6 blade	DP0 - 16 through 25 DP1 - 26 through 35	DP0 - 16 through 20 DP1 - 26 through 30

NOTE
If you configured tunnels and circuits while in FCIP mode, that configuration is carried over to Hybrid mode provided it is compatible with Hybrid mode. If the configuration is incompatible with Hybrid and 10VE mode, you are prompted to make changes.

The following steps are required to configure the Brocade 7840 Switch or Brocade SX6 Blade to operate in Hybrid mode.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `extncfg --app-mode` command to enable Hybrid mode. When prompted, confirm the command action.

```
switch:admin> extncfg --app-mode hybrid

This action will configure the system for Hybrid (FCIP/IPExt) mode.

WARNING: This is a disruptive operation that requires a reboot to take effect. Would you like to
continue (Y,y,N,n): [ n]   y

Operation succeeded. Rebooting the system...
```

3. Use the `extncfg --show` command to confirm Hybrid mode is enabled.

```
switch:admin> extncfg --show
APP Mode is Hybrid (FCIP with IPEXT)
VE-Mode: configured for 10VE mode.
switch:admin>
```

Configuring an IP interface for IP Extension

A circuit endpoint terminates at an IP interface (IPIF), which is defined by an IP address, subnet mask, MTU, and VLAN ID. The default MTU value is 1500 bytes, and the MTU and VLAN ID are optional. If no VLAN is specified, the IPIF will be untagged.

A VLAN ID is needed only when the traffic from the WAN core router/switch is VLAN tagged. The core router/switch to which the circuit connects must be doing VLAN tagging on the traffic to the IP extension platform (such as a Brocade 7840 Switch or Brocade SX6 Blade) before VLANs apply. In most cases, the port on the data center switch is a member of a particular VLAN, but is not set up to do VLAN tagging, which makes VLAN tagging on the IP extension platform unnecessary. Communication with the data center LAN switch will not work if VLAN usage is mismatched. VLANs are used on the Brocade 7840 Switch or Brocade SX6 Blade when multiple circuits pass through the same Ethernet interface, and the connected data center LAN switch directs those circuits to various paths based on the VLAN tag.

For additional information and examples, see Configuring IPIF on page 102.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg ipif` command to create a circuit endpoint IP address.

```
switch:admin> portcfg ipif ge2.dp0 create 10.0.1.3/24
switch:admin> portcfg ipif ge3.dp0 create 10.0.1.4/24
```

   In the example, the endpoint IPIFs of two different circuits on the same Brocade 7840 are created and assigned to GE2 (10.0.1.3) and GE3 (10.0.1.4) respectively. A 24-bit subnet mask is being used (255.255.255.0 = /24). Both IP addresses are processed by DP0, and later when the tunnel circuits are created these local IP addresses (10.0.1.3 and 10.0.1.4) will both be assigned to the same VE_Port as members of an extension trunk. There is no requirement that extension trunk circuits be on the same subnet.

Configuring a WAN IP Route for IP Extension

The IP route is based on the destination IP address to be used by the extension circuit. If the WAN IP network has a different subnet on each end, the network is a Layer 3 network and requires an IP route to be configured on the IP extension platform (like the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade). You can define up 128 routes per Ethernet interface on the Brocade 7840 Switch or Brocade SX6 Blade and up to 120 routes per interface on the Brocade 7810 switch. You can have at most 120 routes per DP. When you create the route, you specify the destination IP subnet or IP address, subnet mask, and gateway IP address for the router that forwards packets to the specified IP subnet or IP address.

You cannot use the `portCfg iproute` *`interface-number.dpnumber`* `modify` command to modify the destination network address. To make changes to the IP route definition, you must delete the IP route, and then recreate it. Also, you cannot use this command to change the prefix length or network mask. This command is intended to change the gateway IP address only.

NOTE

When you create an IP route for IP Extension, the route is created on a LAN interface of DP0 or DP1. This differs from creating IP routes for FCIP. For more information on configuring IP routes for FCIP, see Configuring IP Route on page 104.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Verify that the switch or blade is in Hybrid mode. You must be in Hybrid mode to configure a LAN IPIF.

```
FCIP_Local:admin> extncfg --show
APP Mode is HYBRID (FCIP with IPEXT)
VE-Mode: configured for 10VE mode.
```

3. Use the `portCfg iproute lan.dp-number create` command to create the IP route from the LAN IPIF gateway to the destination network address.

```
switch:admin> portcfg iproute lan.dp0 create 192.168.12.100/32 10.0.1.1
switch:admin> portcfg iproute lan.dp0 create 192.168.12.0/24 10.0.1.1
```

   The first route is a host-based route, meaning that it is a single IP address route to a host, which is indicated by the full netmask, /32.

   The second route is a subnet-based route, meaning that it is a route for an entire subnet, which is indicated by a partial netmask, in this case /24. By specifying a subnet route, you can save configuration time if there are multiple WAN circuits in the same subnet. In a subnet you need not add an IP route for every circuit IP address. The same LAN IPIF gateway is used for each route.

4. Use the `portshow iproute` command to display IP routes configured on the interface. IP routes for FCIP are configured on a GE interface, whereas IP routes for IP Extension are configured on a LAN interface.

```
switch:admin> portshow iproute
Port         IP Address               / Pfx  Gateway               Flags
-------------------------------------------------------------------------------
 ge10.dp0     10.10.0.0                / 16   *                      U C
 ge10.dp0     192.168.12.0             / 24   *                      U C
 ge10.dp0     192.168.12.8             / 32   *                      U H L
 ge11.dp0     192.168.86.0             / 24   *                      U C
 ge11.dp0     192.168.86.14            / 32   *                      U H L
 ge11.dp0     192.168.86.114           / 32   *                      U H L
 ge11.dp1     192.168.86.0             / 24   *                      U C
 ge11.dp1     192.168.86.14            / 32   *                      U H L
 ge12.dp0     192.168.16.0             / 24   *                      U C
 ge12.dp0     192.168.16.15            / 32   *                      U H L
 lan.dp0      192.168.53.128           / 26   *                      U C
 lan.dp0      192.168.53.130           / 32   *                      U H L
 lan.dp0      192.168.53.192           / 26   *                      U C
-------------------------------------------------------------------------------
 Flags: U=Usable G=Gateway H=Host C=Created(Interface)
        S=Static L=LinkLayer X=Crossport
```

   The following example shows IP route information where WAN gateways and LAN gateways are configured. The WAN gateway is configured on the ge10.dp1 interface, and the LAN gateway is configured on the lan.dp0 interface.

```
switch:adming> portshow iproute
Port         IP Address               / Pfx  Gateway               Flags
-------------------------------------------------------------------------------
7/ge10.dp1   192.168.1.0              / 32   *                      U C
7/ge10.dp1   192.168.2.0              / 24   192.168.1.1            U G S
7/lan.dp0    192.168.23.2             / 32   *                      U H L
7/lan.dp0    192.168.40.194           / 32   3.3.5.0                U G S
7/lan.dp0    192.168.53.64            / 26   *                      U C
-------------------------------------------------------------------------------
 Flags: U=Usable G=Gateway H=Host C=Created(Interface)
        S=Static L=LinkLayer X=Crossport
```

   NOTE

   The `portshow iproute` command may display more routes than you have configured. Some routes are added by default. Do not remove or alter these other routes.

5. Use the `portCfg iproute lan.dp-number delete` command to delete IP routes configured on a LAN IPIF gateway.

```
switch:admin> portcfg iproute lan.dp0 delete 10.0.0.3/24
```

Configuring a Tunnel to Support IP Extension

Configure a tunnel to support IP Extension features.

The Brocade 7840 Switch or Brocade SX6 Blade must be in Hybrid mode. GE ports must be configured for LAN operation. Optionally, LAGs are defined.

To use IP Extension features, you enable IP Extension capability on a tunnel. You can create an IP Extension capable tunnel or you can modify an existing tunnel to support IP Extension.

At least one circuit is required for each tunnel before the tunnel can become operational. If the circuit parameters are included in the tunnel creation, circuit 0 is automatically created when the tunnel is created.

A recommended practice is to stage your tunnels and circuits. When you stage your tunnels and circuits, you create a basic tunnel and use the `portcfg fciptunnel ve-port modify` command to add, remove, or change values. When your tunnel is configured, you can add circuits with the `portcfg fcipcircuit create` and `portcfg fcipcircuit modify` commands.

NOTE

When circuit options are specified on the `portcfg fciptunnel ve-port create` command and the `portcfg fciptunnel ve-port modify` command, they apply only to circuit 0 because these are tunnel operands and not `fciptunnel` operands. When additional circuits are added, circuit options must be applied per circuit using the `portcfg fcipcircuit create` or `portcfg fcipcircuit modify` commands.

ATTENTION

When you configure a tunnel, it is strongly recommended that you apply IPsec policies. IPsec is used for data in-flight across the WAN-side circuits. IPsec is applied to the entire tunnel and all member circuits. You cannot apply IPsec to just one individual circuit of a tunnel. Brocade IPsec has been implemented in HW and operates at line rate, adding negligible propagation delay. See Configuring IPsec on the Brocade 7810, the Brocade 7840, and the Brocade SX6 on page 112 for additional information about IPsec implementation.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg fciptunnel ve-port create` command to create an IP Extension capable tunnel.

   This example creates and enables an IP Extension tunnel on VE_Port 24. Notice that protocol and QoS distribution are reset to default values.

```
switch:admin> portcfg fciptunnel 24 create --ipext enable

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]
```

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

4. Use the `portcfg fciptunnel ve-port modify` command to modify an existing tunnel to be IP Extension capable.

   This example modifies an existing tunnel on VE_Port 24 to disable IP Extension.

```
switch:admin> portcfg fciptunnel 24 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]
```

Configuring Bandwidth Distribution

Tunnel traffic is distributed between a Fibre Channel traffic group and an IP Extension traffic group. You can configure bandwidth ratios between the two groups when a tunnel is configured to support IP Extension. QoS priority is distributed within the bandwidth allocations.

When you configure bandwidth distribution, you can change the default allocation of 50/50 between Fibre Channel (FC) and IP Extension (IP) traffic in a tunnel. All bandwidth allocations are expressed as percentages.

NOTE
The minimum percentage allowed for a QoS priority or a distribution group cannot go below 10 percent.

After configuring the FC and IP bandwidth distribution you can configure QoS high, medium, and low priorities in each of the FC and IP distributions. The default priority values for QoS are 50/30/20 for high, medium, and low.

NOTE
Creating or modifying distribution bandwidth can disrupt traffic on the specified tunnel for a brief period of time. The tunnel is brought down before the new configuration is applied, then the tunnel is brought up.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg fciptunnel create --distribution protocol` command to create protocol bandwidth distribution.

   This command creates protocol bandwidth distribution at 60 percent for FC and 40 percent for IP traffic.

```
switch:admin> portcfg fciptunnel 24 create --distribution protocol,60,40
Operation Succeeded
```

3. Use the `portcfg fciptunnel modify --distribution` command to change traffic protocol distribution. The first value applies to FC traffic and the second value applies to IP traffic.

   This command modifies protocol bandwidth values to 40 percent for FC traffic and 60 percent for IP traffic.

```
switch:admin> portcfg fciptunnel 24 modify --distribution protocol,40,60

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
```

4. Use the `portcfg fciptunnel modify` command to change the QoS priority bandwidth allocations in the FC traffic group and the IP traffic group.
   a) Use the `portcfg fciptunnel modify --fc-qos-ratio` command to configure the FC QoS priorities.

      This command modifies the FC QoS priority bandwidth values for high to 30 percent, medium to 50 percent, and low to 20 percent.

      ```
      switch:admin> portcfg fciptunnel 24 modify --fc-qos-ratio 30,50,20
      Operation Succeeded
      ```

   b) Use the `portcfg fciptunnel modify --ip-qos-ratio` command to configure the IP QoS priorities.

      This command modifies IP QoS priority bandwidth values for high to 60 percent, medium to 30 percent, and low to 10 percent.

      ```
      switch:admin> portcfg fciptunnel 24 modify --ip-qos-ratio 60,30,10
      Operation Succeeded
      ```

5. Use the `portshow fciptunnel -d` command to display the distribution values.

   The following example displays the results of QoS protocol distribution to 40/60 percent, and the results of changing the IP QoS and FC QoS bandwidth values. (The output is truncated.)

   ```
   switch:admin> portshow fciptunnel -d

    Tunnel: VE-Port:24 (idx:0, DP0)
    ==================================================
     Oper State           : In Progress
     TID                  : 24
     Flags                : 0x00000000
     IP-Extension         : Enabled
     Compression          : None
     FC-Compression       : None (Inherited)
     IP-Compression       : None (Inherited)
     QoS Distribution     : Protocol (FC:40% / IP:60%)
     FC QoS BW Ratio      : 30% / 50% / 20%
     IP QoS BW Ratio      : 60% / 30% / 10%
    [...]
   ```

   Remember that the distribution and bandwidth values should be configured the same at both ends of the tunnel.

Configuring a LAN IP Route for IP Extension and Policy-Based Routing

IP Extension supports policy-based routing (PBR) connections to the LAN interfaces on the Brocade 7840 Switch, the Brocade 7810 Switch, and the Brocade SX6 Blade. This allows Layer 3 (routing) connections, in addition to the Layer 2 (direct) connections, to the extension platform.

Most configuration steps for PBR are done on the router. You configure a next-hop gateway address on the router that identifies the extension platform interface as a next-hop, and you configure a policy, such as an access control list, that determines what traffic is sent to the extension platform.

On the extension platform (the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade), you create an IP route to the next-hop gateway. Steps for configuring the IP route for PBR are similar to configuring an IP route for WAN.

When you configure an address for the next-hop gateway, the extension platform learns the MAC address of the next-hop gateway through an ARP message. Thereafter, any packets to the next-hop gateway are forwarded to the learned MAC address through the LAN FE port.

The following illustration shows an example network topology where the Brocade 7840 Switch functions as a next-hop gateway for policy-based routing traffic from the DC router. Traffic that is not diverted to the Brocade 7840 Switch by the PBR in the DC router passes through to the core router.

FIGURE 21 PBR Example Network

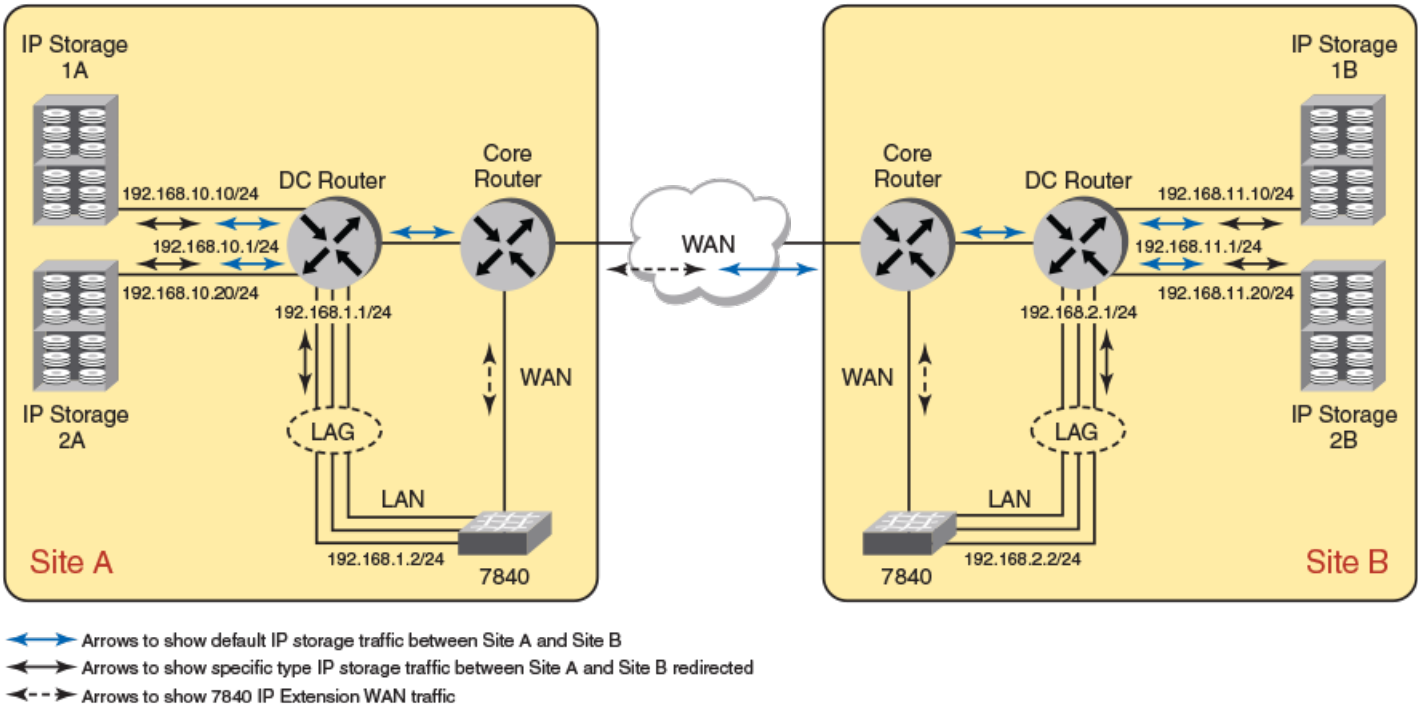

The policies are implemented in the data center (DC) router at each site. In the configuration steps, the router configurations and the IP storage configurations are generic. You must refer to the documentation specific to your equipment for the detailed steps.

NOTE

It is assumed that the WAN tunnel configuration across the core router is complete. It is also assumed that DC routers are configured and can route across core routers using the path between sites A and B.

Configuring the Brocade 7840 Switch at Each Site

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg ipif` command to configure an IP LAN interface. You perform this action on the Brocade 7840 at each site.

```
admin:switch_siteA> portcfg ipif lan.dp0 create 192.168.1.2 netmask 255.255.255.0
admin:switch_siteB> portcfg ipif lan.dp0 create 192.168.2.2 netmask 255.255.255.0
```

This command creates a LAN interface on the Brocade 7840 Switch.

3. Use the `portcfg iproute` command to specify the next-hop gateway address. You perform this action on the Brocade 7840 at each site.

```
admin:switch_siteA> portcfg iproute lan.dp0 create 192.168.10.0 netmask 255.255.255.0 192.168.1.1

admin:switch_siteB> portcfg iproute lan.dp0 create 192.168.11.0 netmask 255.255.255.0 192.168.2.1
```

These commands configure a LAN route with the destination subnet of the local storage using the DC router LAG interface IP as the gateway.

Configuring the Routers at Each Site

NOTE
You must refer to the documentation specific to your router equipment for the detailed steps. The following steps provide generic information.

4. On the DC Router at Site A, configure the router interface for the LAG that is connected to the Brocade 7840 Switch.

```
DC Router Site A:
IP: 192.168.1.1 mask 255.255.255.0
```

5. On the DC Router at Site B, configure the router interface for the LAG that is connected to the Brocade 7840 Switch.

```
DC Router Site B:
IP: 192.168.2.1 mask 255.255.255.0
```

Configuring the router policy at each site

NOTE
You must refer to the documentation specific to your router equipment for the detailed steps. The following steps provide generic information.

6. Configure DC Router at Site A with a policy-based route to direct a specific IP traffic type or ACL, destined to the remote site, to use the Brocade 7840 Switch LAN interface as the gateway or next hop.

```
DC Router Site A:
Destination subnet: 192.168.11.0 mask 255.255.255.0
Source subnet: 192.168.10.0 mask 255.255.255.0
Type: <specific protocol or well-known port or ACL>
Gateway: 192.168.1.2
```

7. Configure DC Router at Site B with policy-based route to direct a specific IP traffic type or ACL, destined to the remote site, to use the Brocade 7840 Switch LAN interface as the gateway or next hop.

```
DC Router Site B:
Destination subnet: 192.168.10.0 mask 255.255.255.0
Source subnet: 192.168.11.0 mask 255.255.255.0
Type: <specific protocol or well-known port or ACL>
Gateway: 192.168.2.2
```

Configuring the IP storage devices at each site

NOTE
You must refer to the documentation specific to your storage devices for the detailed steps. The following steps provide generic information.

8. Configure the IP interfaces on the storage devices at Site A, which are shown in the illustration as **IP Storage 1A** and **IP Storage 2A**.

```
IP Storage 1A:
IP: 192.168.10.10 mask 255.255.255.0

IP Storage 2A:
IP: 192.168.10.20 mask 255.255.255.0
```

9. Configure the IP interfaces on the storage devices at Site B, which are shown in the illustration as **IP Storage 1B** and **IP Storage 2B**.

```
IP Storage 1B:
IP: 192.168.11.10 mask 255.255.255.0

IP Storage 2B:
IP: 192.168.11.20 mask 255.255.255.0
```

10. Configure the destination routes for peer-site subnet using the DC Router at each site as a gateway. The Site A IP storage devices have a destination route to Site B IP storage devices, and the Site B devices to Site A devices.

```
IP Storage 1A destination route to 1B:
Destination route: 192.168.11.0 mask 255.255.255.0 gateway 192.168.10.1

IP Storage 2A destination route to 2B:
Destination route: 192.168.11.0 mask 255.255.255.0 gateway 192.168.10.1

IP Storage 1B destination route to 1A:
Destination route: 192.168.10.0 mask 255.255.255.0 gateway 192.168.11.1

IP Storage 2B destination route to 2A:
Destination route: 192.168.10.0 mask 255.255.255.0 gateway 192.168.11.1
```

Configuring Tunnel Compression

Configure tunnel compression per protocol on the tunnel. You can override inherited compression values.

The Brocade 7840 Switch or Brocade SX6 Blade must be in Hybrid mode. GE ports must be configured for LAN operation and a tunnel must be configured to support IP Extension.

The enhancements for IP Extension allow you to configure compression on the tunnel at a protocol level. The protocol compression options override the main tunnel compression level and set the compression for the specified protocol to the desired mode. The available modes depend on the protocol, whether FC or IP.

The IP priorities do not support fast-deflate compression, so only the deflate and aggressive deflate options are allowed with the `--ip-compression` option. If the main tunnel compression is set to fast deflate, the IP priorities are set to none. The protocol-based compression modes can be set to default, which causes the protocol compression to inherit the configuration from the main tunnel compression setting.

The configuration steps show how to set compression for a tunnel and how to set compression overrides for traffic in that tunnel.

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

2. Use the `portcfg fciptunnel modify --compression` command to configure the fast-deflate compression level for tunnel compression.
   a) Configure fast-deflate tunnel compression.

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

   b) Verify that fast-deflate compression is not supported for the IP protocol.

```
switch:admin> portshow fciptunnel 24

 Tunnel: VE-Port:24 (idx:0)
 ==================================================
  ...
  Compression          : Fast Deflate
  FC-Compression       : Fast Deflate (Inherited)
  IP-Compression       : None (Inherited)
```

3. Use the `portcfg fciptunnel modify --ip-compression` command to change the IP compression to deflate.
   a) Configure deflate compression for the IP protocol.

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

   b) Verify that deflate compression is configured for IP protocol.

```
switch:admin> portshow fciptunnel 24

 Tunnel: VE-Port:24 (idx:0)
 ==================================================
  ...
  Compression          : Fast Deflate
  FC-Compression       : Fast Deflate (Inherited)
  IP-Compression       : Deflate (Override)
  ...
```

4. You can use the `portcfg fciptunnel modify` command to return compression values to their default, inherited values.

```
switch:admin> portcfg fciptunnel 24 modify --fc-compr default --ip-compr default
```

Configuring a Tunnel and Circuits for IP Extension

After you configure a tunnel on a VE_Port, create circuits on the VE_Port that belong to this tunnel. Start with circuit 0 and add circuits incrementing by 1. Circuit 0 is created by default when you create a tunnel. You must create a circuit for each path you want through the IP infrastructure. A VE_Port on the Brocade 7840 switch or Brocade SX6 blade can have up to 10 circuits assigned to it. Circuits require identical settings on each end of the circuit to come online.

When you configure a circuit, carefully consider the keep-alive timeout value (KATOV). Keepalive Time-Out Value (KATOV) is important for proper error recovery. The KATOV should be based on application requirements. Check with your IP storage application provider to determine the appropriate KATOV for your application. The sum of all circuit KATOVs in a tunnel should be slightly less than the I/O timeout value. As an example, a mirroring application has a 6-second I/O timeout. There are three circuits belonging to a VE_Port (3 circuit members in the tunnel). Set the KATOV to 2 seconds on each circuit. This will allow maximum retries over all available circuits before an I/O is timed out by the IP storage application. In many cases, a 2- to 3-second KATOV is optimal for IP Extension and should be strongly considered.

See Keep-alive Timeout Values for Different FC Protocols on page 123 for more information.

NOTE

There is a 2-second KATOV boundary for support. When KATOV is set for 2 seconds and above, the supported round-trip time (RTT) is 250 ms and a 1-percent packet loss. When KATOV is less than 2 seconds, the supported RTT is 200 ms and 0.1-percent packet loss.

1. Use the `portcfg fcipcircuit` command to create and modify circuits. Use the `portcfg fciptunnel` command to create a tunnel with no circuits. A tunnel must be in place before you can add circuits.

```
switch:admin> portcfg fciptunnel 24 create
switch:admin> portcfg fcipcircuit 24 create 0
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 0 --max-comm-rate 1000000 --min-comm-rate 5000000
switch:admin> portcfg fcipcircuit 24 modify 0 --keepalive 1000

switch:admin> portcfg fcipcircuit 24 create 1
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 1 --max-comm-rate 1000000 --min-comm-rate 5000000
switch:admin> portcfg fcipcircuit 24 modify 1 --keepalive 1000
```

   The example shows adding two circuits (0 and 1) to VE_Port 24. The keepalive is changed to 1 second, which decreases the link loss recovery time. The ARL value for the circuit is modified with the ARL floor at 5 Gb/s and the ceiling at 10 Gb/s. When configuring the remote side, the commands are nearly the same except for the local and remote IP addresses are switched.

2. Use the `portshow fcipcircuit` command to display circuit values. This command is useful for keeping track of the staging activity while you modify circuit values. You can also see when circuit configuration has all the required values and is considered complete.

```
switch:admin> portshow fcipcircuit 24 0

 Circuit 24.0 (DP0)
 ==================================================
  Admin/Oper State     : Enabled / Configuration Incomplete
  Flags                : 0x00000000
  IP Addr (L/R)        : 172.16.1.3 <-> 172.16.2.3
  HA IP Addr (L/R)     : 0.0.0.0 <-> 0.0.0.0
  Configured Comm Rates: 1000000 / 5000000 kbps
  Peer Comm Rates      : 0 / 0 kbps
  Actual Comm Rates    : 0 / 0 kbps
  Keepalive (Cfg/Peer) : 1000 (6000 / 0) ms
```

   The example shows output for VE_Port 24 and circuit 0 before any values are configured, such as the local and remote IP addresses, bandwidth values, or KATOV.

Configuring Ethernet Interfaces (GbE Port) for IP Extension LAN Features

Configure the GbE port for IP Extension LAN features.

The Brocade 7840 Switch or Brocade SX6 Blade must be in Hybrid mode. (The Brocade 7810 Switch operates only in Hybrid mode.)

To use the IP Extension features, you must configure a GbE port to operate in LAN mode. IP Extension features cannot function without a LAN interface. The end-device storage on an IP network communicates with IP Extension by means of the LAN IPIF gateway, a function provided through the GbE port when it is configured as a LAN interface. If the interface is not configured as LAN, it defaults to a WAN interface and incoming traffic is treated as WAN traffic and dropped.

Any existing IP configuration must be removed before changing a GbE port to LAN mode. Once a port is configured as a LAN port, it cannot be used as a WAN port for any circuit definitions.

Only the 1GbE/10GbE ports, or Ethernet interfaces, can be configured as LAN ports. The 16 available Ethernet interfaces are numbered from GE2 to GE17. Up to eight 10GE interfaces can be configured as LAN ports. Speeds of 1G or 10G are supported. The 40GbE ports, GE0 and GE1, cannot be configured as LAN ports.

NOTE

The Brocade 7810 switch provides a maximum of six Ethernet interfaces (6 * 1/10GbE optical ports or 4 x 1/10GbE optical ports and 2 x 1GbE RJ-45 copper ports). All ports support LAN and WAN configuration and can be used as per installed license.

The following steps are required to configure a GbE port to operate in LAN mode.

1. Use the `portcfgge` command to configure a GbE port for LAN operation.

   This example puts port GE10 in LAN mode.

```
switch:admin> portcfgge ge10 --set -lan
Operation Succeeded.
```

2. Optionally, use the `portcfgge` command to change the interface speed.

   The example sets port speed for port GE10 to 1 Gb/s.

```
switch:admin> portcfgge ge10 --set -speed 1G
Operation Succeeded.
```

3. Use the `portcfgge --show` command to verify that the GbE port is in LAN mode and to show the port speed.

   (The example output is truncated.)

```
switch:admin> portcfgge --show

 Port         Speed     Flags     LAG-ID
-----------------------------------------
 ge0          40G       -----     -
 ge1          40G       -----     -
 ge2          10G       -----     -
 ge3          10G       -----     -
 ge4          10G       -----     -
 ge5          10G       -----     -
 ge6          10G       -----     -
 ge7          1G        -----     -
 ge8          10G       -----     -
 ge9          10G       -----     -
 ge10         1G        --L--     -
[ . . . ]
-----------------------------------------
Flags: A:Auto-Negotiation Enabled C:Copper Media Type
L:LAN Port G=LAG Member
```

4. Alternatively, use the `switchshow` command to verify that the GbE port is in LAN mode.

   (The example output is truncated.)

```
switch:admin> switchshow
Index  Port  Address Media Speed  State   Proto
==================================================
         ge0         id    40G    Online      FCIP
         ge1         id    40G    Online      FCIP
         ge2         id    10G    Online      FCIP
         ge3         id    10G    Online      FCIP
         ge4         id    10G    Online      FCIP
         ge5         id    10G    Online      FCIP
         ge6         id    10G    Online      LAN
         ge7         id    10G    Online      LAN
         ge8         id    10G    Online      LAN
         ge9         id    10G    Online      FCIP
         ge10        id     1G    Online      LAN
```

Configuring a LAN Gateway (SVI) for IP Extension

The switch virtual interface (SVI) IP interface (IPIF) acts as the gateway address for the endpoint device(s) being redirected through the IP extension platform. The gateway address on the LAN interface acts as a virtual switch and is often the next-hop address for IP routes from the endpoint device.

The IP extension platform (such as the Brocade 7840 Switch or Brocade SX6 Blade) must be configured for Hybrid mode. (The Brocade 7810 Switch operates only in Hybrid mode.) GE ports must be configured for LAN operation.

There is only one SVI interface per data processor (DP) complex, meaning one LAN-side Ethernet device and MAC per DP. The Brocade 7840 Switch and the Brocade SX6 Blade can have at most 8 SVI IP addresses defined per DP, but they all use the same single SVI interface. The Brocade 7810 Switch can have a maximum of 4 SVI IPIFs.

NOTE
For this discussion, the SVI IPIF is referred to as the LAN gateway, because that is the function it provides for directly connected devices.

You must configure a LAN gateway (SVI IPIF) for each different subnet that an end device has an interface on. On the end device, for each different IP subnet used on interfaces, you must add an IP route on the end device that points to the LAN gateway. These IP routes are required before the end device can send traffic to the IP extension platform.

Local end-device IP addresses can be on the same subnet, different subnets, or a combination of subnets. The local and remote end-device IP addresses must be on different subnets.

In the following situations, separate LAN gateways must be configured:

NOTE
Multiple LAN gateways belonging to the same IP subnet cannot be configured on the same DP. A separate IPIF gateway is needed for each VLAN on the LAG.

If the LAG is not configured to use VLAN tagging (802.1Q) on the data center LAN switch, do not associate a VLAN ID with the LAN IPIF because that will prevent communications with the data center LAN switch. Tagged traffic can only talk to interfaces that are configured to recognize tagged traffic.

When there are multiple VLANs, the LAG is a VLAN trunk and the traffic is tagged. In this instance, multiple logical ISLs pass over the physical LAG and the different VLAN traffic is tagged with the appropriate VLAN ID. For the tagged traffic to be properly forwarded to the corresponding LAN IPIF gateway, the gateway must be configured with the matching VLAN ID for each VLAN. If the LAG to the LAN switch is tagged, a LAN IPIF gateway is created for each VLAN.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use `portcfg ipif create` to configure a LAN gateway port on a DP.

   This example shows a LAN gateway on DP0 configured on VLAN 100.

   ```
   switch:admin> portcfg ipif lan.dp0 create 10.0.0.1/24 vlan 100
   ```

   This example shows a tagged LAN gateway on DP0 configured for the maximum MTU of 9216 on VLAN 200.

   NOTE
   PMTU auto-discovery is not supported on the LAN gateway interface (SVI IPIF).

   ```
   switch:admin> portcfg ipif lan.dp0 create 10.0.1.1/24 vlan 200 mtu 9216
   ```

   This example shows an untagged LAN gateway on DP1.

   ```
   switch:admin> portcfg ipif lan.dp1 create 10.0.2.1/24
   ```

3. Use `portcfg ipif delete` to delete a LAN gateway port on a DP.

   The example deletes a LAN gateway port on DP0.

   ```
   switch:admin> portcfg ipif lan.dp0 delete 10.0.0.3
   ```

   A LAN IPIF can be deleted while still in use. Clean-up enforcement checks are not done on a LAN IPIF. Inadvertently deleting a LAN IPIF will disrupt all IP extension flows using that LAN gateway.

4. Use `portshow ipif` to display the IPIF ports.

   The example shows two LAN ipif gateways (10.0.0.1 on VLAN 100 with MTU 1500, and 10.0.1.1 on VLAN 200 with MTU 9216) configured on DP0, as well as two WAN circuit endpoint IPIFs (192.168.60.20 on GE4 and 192.168.10.107 on GE17) configured on DP0.

   ```
   switch:admin> portshow ipif

    Port         IP Address                    / Pfx  MTU   VLAN  Flags
   ------------------------------------------------------------------------------
    ge4.dp0      192.168.60.20                 / 24   1500  0     U R M
    ge17.dp0     192.168.10.107                / 24   1500  0     U R M
    lan.dp0      10.0.0.1                      / 24   1500  100   U R M
    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
          N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
   ```

Configuring Traffic Control Lists for IP Extension

A traffic control list (TCL) defines how LAN traffic is mapped to specific tunnels. Multiple TCL rules, arranged by priority, provide a high level of control over the LAN traffic flow through a particular DP. For additional details about TCL rules and how they operate, see IP Extension and Traffic Control Lists on page 66.

NOTE

Brocade 7840 Switch and Brocade SX6 Blade support a maximum of 1024 defined and 128 active TCLs whereas the Brocade 7810 Switch supports a maximum of 256 defined and 32 active TCLs.

NOTE

TCL rules must be configured. One default rule exists, which is to deny all traffic. This default rule cannot be removed or modified. It is the lowest priority rule, 65535, so it will be the last rule enforced. To have traffic over an IP extension tunnel, you must configure one or more rules that allow traffic to pass through.

Each TCL rule is identified by a name assigned to the rule. Each name must be unique within the IP extension platform (such as a Brocade 7840 Switch, a Brocade 7810 Switch, or a Brocade SX6 Blade). Rules are local to each platform and are not shared across platforms.

The TCL priority number provides an order of precedence to the TCL rule within the overall TCL list. The priority value must be unique across all active TCL rules within an IP extension platform.

The TCL input filter inspects a set of parameters to help identify the input traffic. The TCL input filter identifies a particular host, device, or application by means of the Layer 4 protocol encapsulated within IP, Layer 4 destination ports, Layer 3 source or destination IP addresses and subnets, DSCP/802.1P QoS values, or VLAN tagging.

When defining a TCL rule, the TCL action (allow or deny) and TCL target will determine the behavior with regard to the DPs.

Use the `portshow tcl --help` command to view the options available for displaying TCL information.

The configuration steps show how to create a TCL named FromSubnetA, enable it, identify a target, source address, and set the rule priority all with a single command. The source IP address identifies a subnet.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg tcl` command to create a TCL.

```
switch:admin> portcfg tcl FromSubnetA create --admin enable --target 24 --src-addr 10.0.0.0/8 --
priority 10
Operation Succeeded
```

3. Use the `portshow tcl` command to display TCL information.

```
switch:admin> portshow tcl

 Pri    Name             Flgs      Target    L2COS  VLAN DSCP Proto  Port  Hit
                         Src-Addr                  Dst-Addr
-------------------------------------------------------------------------------
 *10    FromSubnetA      AI---     24-Med    ANY    ANY  ANY  ANY    ANY   0
                         10.0.0.0/8                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.
```

Configuring TCL for Multiple Tunnels

You can configure TCL rules that allow traffic to go to a specific tunnel when more than one tunnel is needed for different destinations.

Multiple tunnels are typically used to go to multiple data centers. Normally, only one tunnel is used per data center between two IP extension platforms (i.e., a pair of Brocade 7840 Switches, Brocade 7810 Switches, or Brocade SX6 Blades). A single tunnel can leverage Extension Trunking for aggregated bandwidth using multiple circuits, failover and failback, lossless link loss (LLL), and other benefits.

Refer to the following figure. It 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 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 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 figure). Priority rule 20 is evaluated next. That rule allows IP address. 10.10.10.0 with 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 65535. The final rule, which cannot be altered or removed, is to deny all traffic. Any traffic processed by this final priority is dropped.

FIGURE 22 TCL Rules and Multiple Tunnels

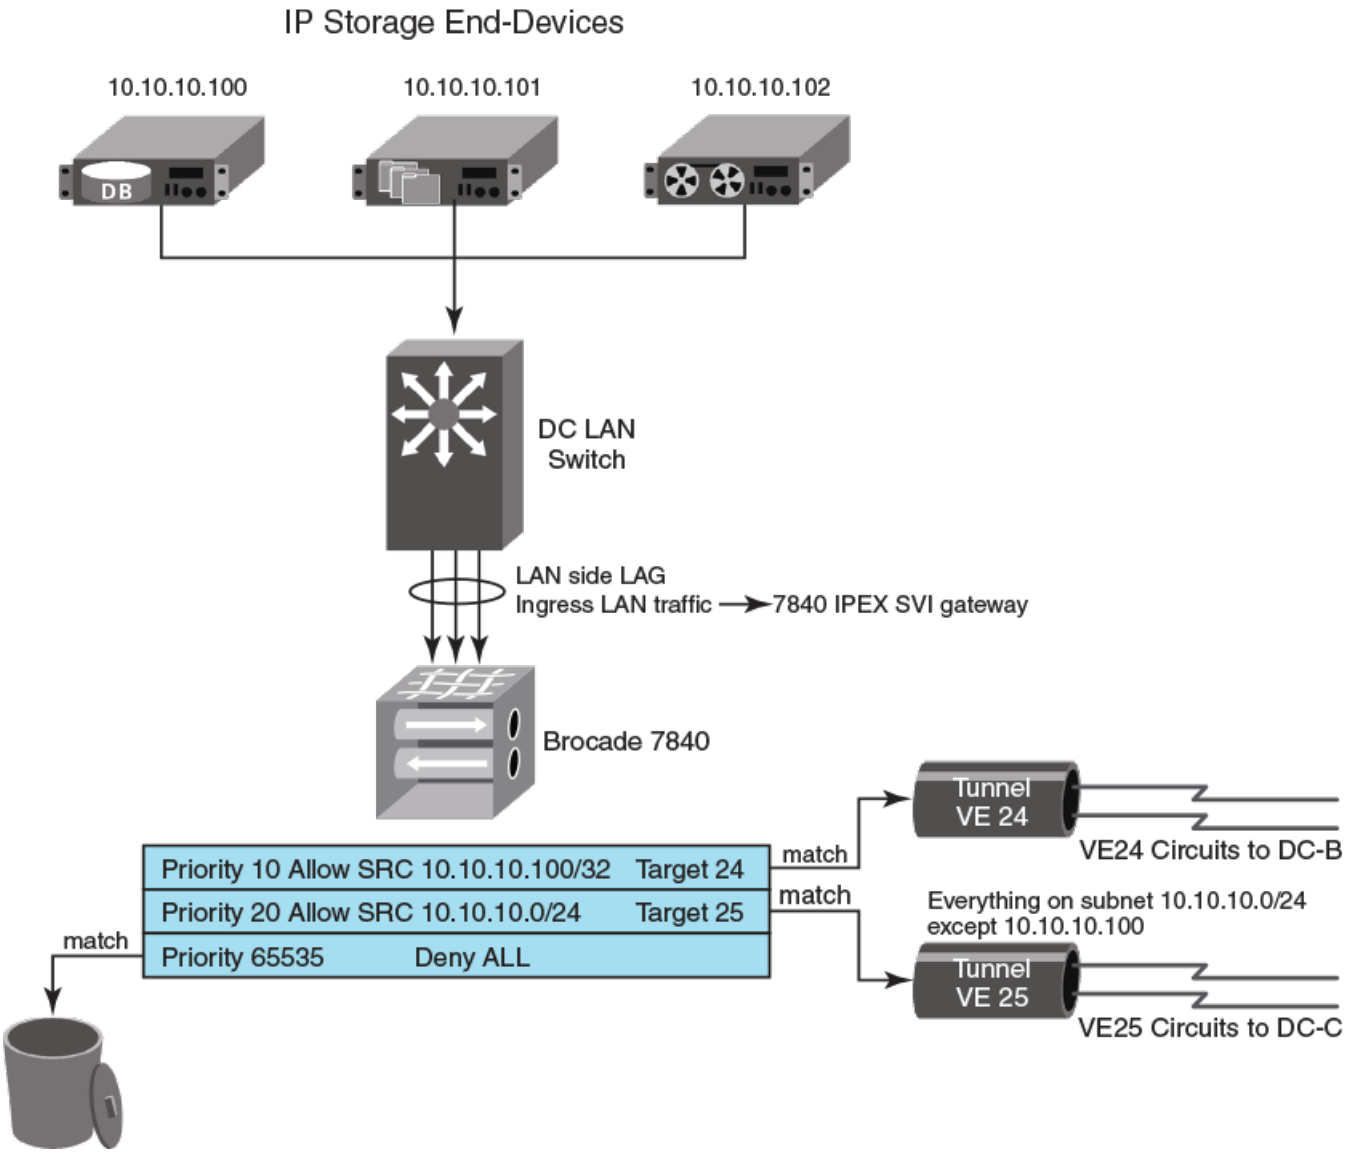

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg tcl` command to create the TCL rules.

```
switch:admin> portcfg tcl DCB create --admin enable --target 24 --src-addr 10.10.10.100/32 --
priority 10
switch:admin> portcfg tcl DCC create --admin enable --target 25 --src-addr 10.10.10.0/24 --priority
20
```

3. Use the `portshow tcl` command to display the configured TCL rules.

```
switch:admin> portshow tcl

 Pri    Name          Flgs      Target   L2COS  VLAN  DSCP  Proto   Port   Hit
                      Src-Addr                  Dst-Addr
-------------------------------------------------------------------------------
 *10    DCB           AI---      24-Med   ANY   ANY   ANY   ANY     ANY    0
                      10.10.10.100/32           ANY

 *20    DCC           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.
```

Configuring Traffic Control Lists for a Single Tunnel

You can create traffic control lists (TCLs) to filter traffic based on different criteria. For example, use source addresses and destination addresses to determine whether traffic is passed to a tunnel. All traffic filtered by a TCL is from the local LAN. Traffic from a WAN extension tunnel is not filtered.

The following figure shows traffic is filtered based on source address, destination address, and VLAN tagging:

- The first rule (priority 10) is processed to deny traffic from IP source address 10.10.10.123. Because the address prefix is /32 (netmask 255.255.255.255) all bits of the address must match. This address is a single interface and not a subnet address.
- The next rule (priority 20) is processed to allow traffic from IP source address 10.10.10.0, and is configured as a subnet address. The address prefix is /24 (netmask 255.255.255.0). The traffic target is the tunnel at VE_Port 24.
- The next rule (priority 30) allows traffic destined for a single interface at IP address 192.168.1.50. The address prefix /32 identifies it as a single interface, not a subnet. None of the traffic for this destination address is from subnet 10.10.10.0, because that traffic was filtered by the previous rule (priority 20). Instead, this traffic must be from a different subnet, or interface, and destined for 192.168.1.50.
- The next rule (priority 40) allows all traffic tagged with VLAN 22 (802.1Q). Any traffic that matches this rule will not have a source address from subnet 10.10.10.0, and will not be destined for 192.168.1.50. The priority 20 and priority rule 30 have filtered that traffic.
- The final rule is the default rule (priority 65535), which denies all traffic that is not specifically allowed.

FIGURE 23 Example TCL for Target 24

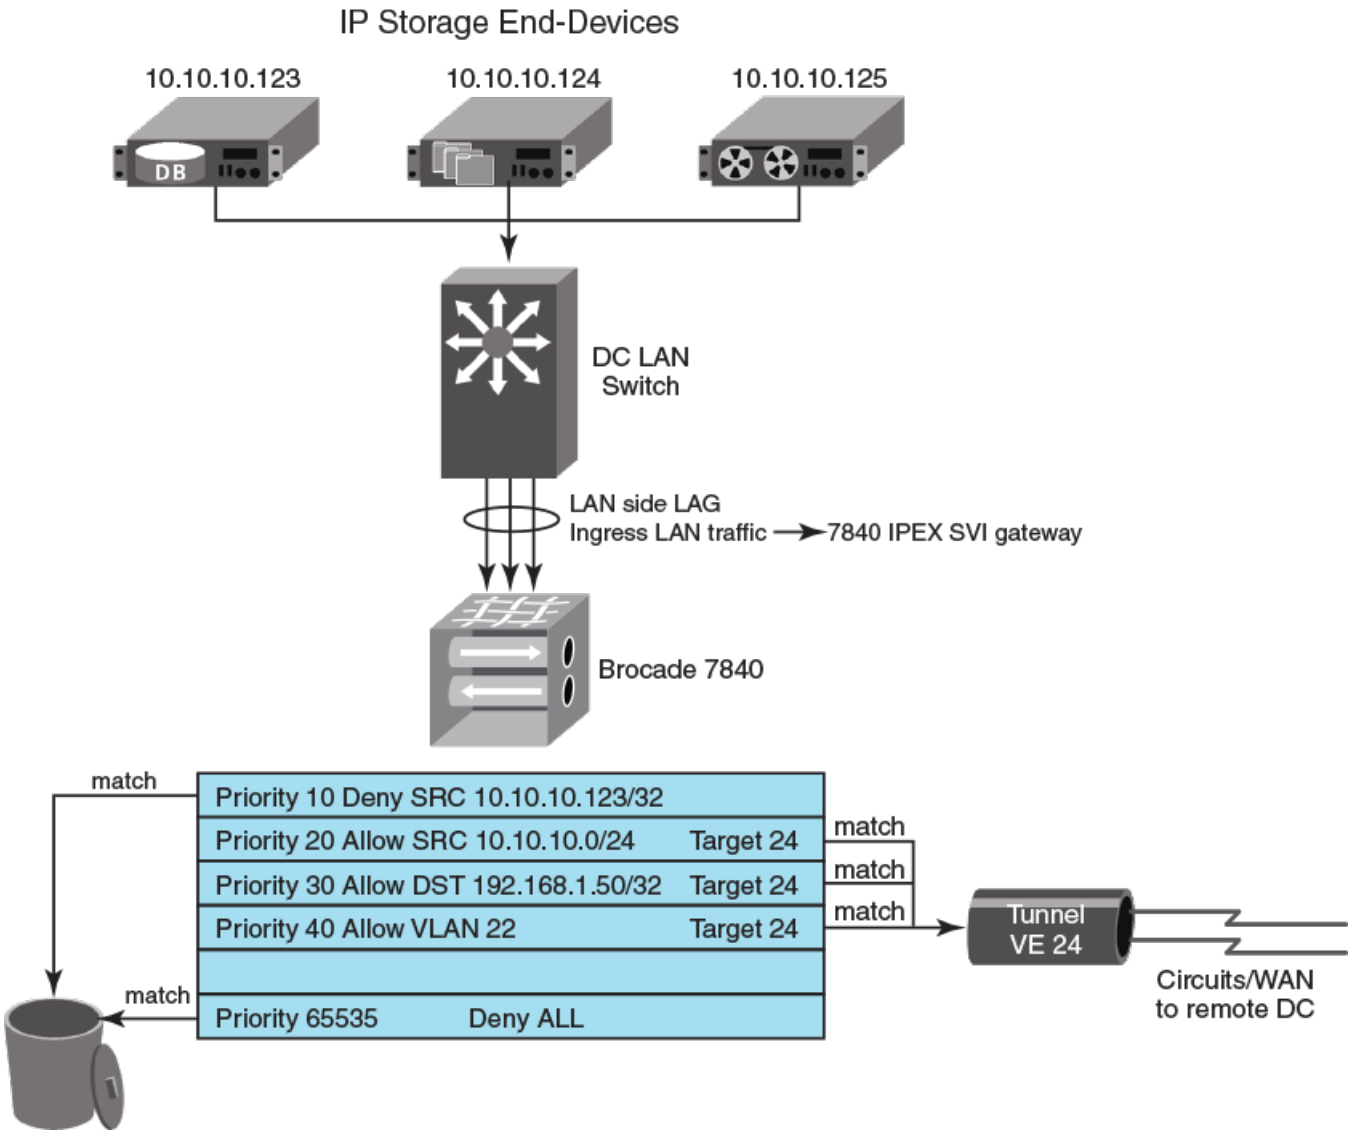

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg tcl` command to configure the TCL rules. Each rule must have a unique name and priority.

```
switch:admin> portcfg tcl MYTCL10 create --admin enable --action deny --target 24 \
              --src-addr 10.10.10.123/32 --priority 10
switch:admin> portcfg tcl MYTCL20 create --admin enable --target 24 \
              --src-addr 10.10.10.0/24 --priority 20
switch:admin> portcfg tcl MYTCL30 create --admin enable --target 24 \
              --dst-addr 192.168.1.50/32 --priority 30
switch:admin> portcfg tcl MYTCL40 create --admin enable --target 24 \
              --vlan 22 --priority 40
```

Configuring Traffic Control Lists for non-terminated traffic

You can create a traffic control list (TCL) rule for traffic that does not terminate at the Brocade 7810 switch, the Brocade 7840 switch, or the Brocade SX6 blade. The non-terminated TCL option allows TCP traffic to be sent as-is to the other endpoint over a tunnel.

As a best practice, configure the non-terminated TCLs with a higher priority than the terminated TCL. For example, a priority value of 10 is higher than a priority value of 100. By doing so, lookup overhead for the non-terminated flows is reduced.

The following steps are required to configure TCL for non-terminated traffic. By default this option is disabled.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg tcl create` command with the `--non-terminated enable` option to create a TCL rule and enable it.

```
switch:admin> portcfg tcl control_1 create --priority 500 --proto-port 3400-3500 --non-terminated
enable --admin-status enable  --target 25-High
```

3. To disable a non-terminated TCL feature, Use the `portcfg tcl modify` command with the `--non-terminated disable` option.

```
switch:admin> portcfg tcl control_1 modify --non-terminated disable
```

4. Use the `portcfgshow 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     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.
```

5. Use the `portcfgshow tcl`command with the `--detail` option to display 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:               789
 Target:                 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:              0
```

Configuring an App-type for a Traffic Control List

You can create an app-type to use for additional TCL filtering.

You create an app-type when you have a custom application that uses a port, or port range, that needs to be provided in a traffic control list (TCL) so that the traffic for that port can be allowed or denied.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Use the `portcfg app-type` command to configure the app name and port or port range.

```
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,

```
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 an app-type.

```
switch:admin> portcfg app-type SpecialApp delete
Operation Succeeded

switch:admin> portcfg app-type SSH delete
Unable to modify or delete default app-types
```

You cannot delete or modify a default app-type.

Example of an IP Extension Configuration

The configuration example shows IP Extension between two data center LANs that are connected through a WAN.

The configuration example is based on the following figure. The examples use “left>” to indicate the data center as shown on the left side of the figure, “right>” for the data center on the right side, and “both>” for actions affecting both sides of the data center.

FIGURE 24 IP Extension Configuration Example

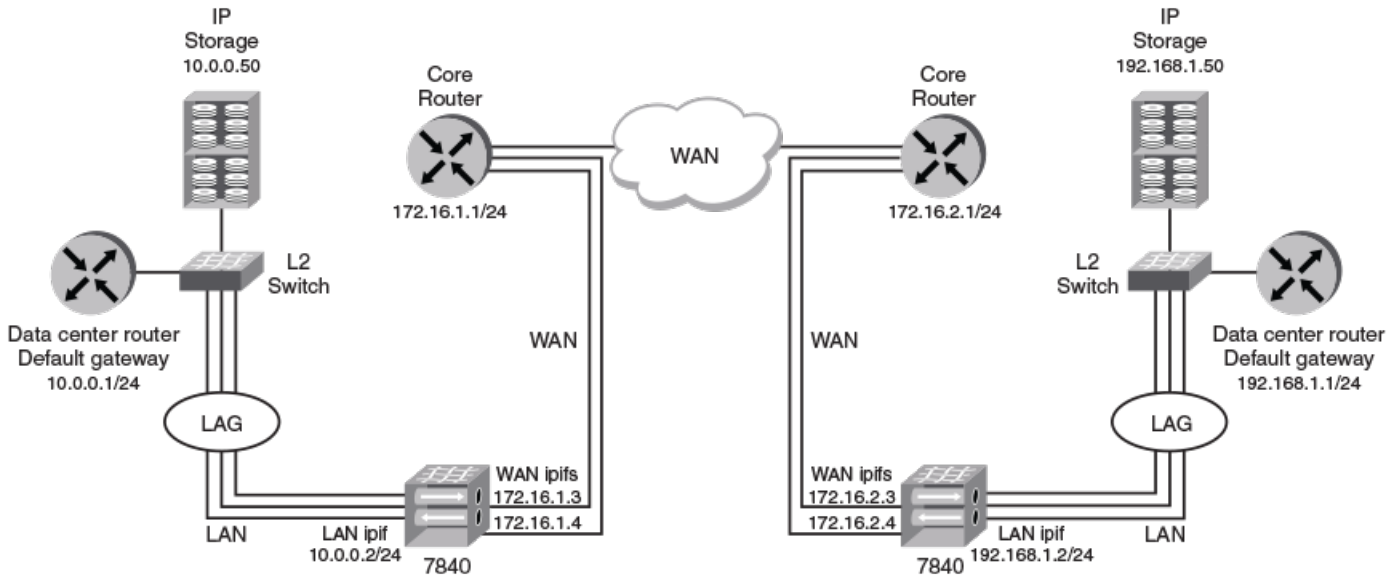

Configuring Hybrid Mode

The following table shows the command used to configure both Brocade 7840 switches for Hybrid mode. The switches reboot when the `app-mode` is changed.

TABLE 37 Brocade 7840 Switch, Hybrid Configuration

Action / Command on Left>	Description	Action / Command on Right>	Description
`both> extncfg --app-mode hybrid`	Places switch in Hybrid mode to support IP Extension	`both> extncfg --app-mode hybrid`	Places switch in Hybrid mode to support IP Extension

Configuring WAN IPs

The following table shows the commands to configure WAN IP for the switches. The commands assume the following:

- The Brocade 7840 switch on the left uses **ge2** and **ge3** for WAN load balanced circuits.
- The core router WAN gateway on the left has IP address **172.16.1.1**
- The Brocade 7840 switch on the right uses **ge2** and **ge3** for WAN load balanced circuits.
- The core router WAN gateway on the right has IP address **172.16.2.1**

TABLE 38 Brocade 7840 Switch, WAN IP Configuration

Action / Command on Left>	Description	Action / Command on Right>	Description

`left> portcfg ipif ge2.dp0 create 172.16.1.3/24`
 	
`left> portcfg ipif ge3.dp0 create 172.16.1.4/24`
	Create IPIFs for ge2 and ge3. IPIFs are on a subnet.	`right> portcfg ipif ge2.dp0 create 172.16.2.3/24`
		
		`right> portcfg ipif ge3.dp0 create 172.16.2.4/24`
			Create IPIFs for ge2 and ge3. IPIFs are on a subnet.

TABLE 38 Brocade 7840 Switch, WAN IP Configuration (continued)

Action / Command on Left>	Description	Action / Command on Right>	Description

`left> portcfg iproute ge2.dp0 create 172.16.2.0/24 172.16.1.1`
 	
`left> portcfg iproute ge3.dp0 create 172.16.2.0/24 172.16.1.1`
	Create IP routes to right-side data center, including subnet mask.	`right> portcfg iproute ge2.dp0 create 172.16.1.0/24 172.16.2.1`
		
		`right> portcfg iproute ge3.dp0 create 172.16.1.0/24 172.16.2.1`
			Create IP routes to left-side data center, including subnet mask.

Configuring IPsec

It is recommended that IPsec should be configured for traffic that goes across the WAN. The following table shows IPsec configuration.

TABLE 39 Brocade 7840 Switch, Optional IPsec Configuration

Action / Command on Left>	Description	Action / Command on Right>	Description
`both> portcfg ipsec-policy siteAtoB create --preshared-key ipsecPresharedKey`	Create IPsec policy and provide a shared key.	`both> portcfg ipsec-policy siteAtoB create --preshared-key ipsecPresharedKey`	Create IPsec policy and provide a shared key.

Configuring a Service-level Agreement

You can configure an service-level agreement (SLA) that will test the circuit before it is enabled. The SLA is assigned to a circuit when you create tunnel and circuit configurations. The following table shows an SLA configuration.

TABLE 40 Brocade 7840 Switch, Optional SLA Configuration

Action / Command on Left>	Description	Action / Command on Right>	Description
`both> portcfg sla primaryNet create --loss 0.2 --runtime 10 --timeout 30`	Create SLA at each end of the circuit.	`both> portcfg sla primaryNet create --loss 0.2 --runtime 10 --timeout 30`	Create an SLA at each end of the circuit.

Verifying WAN Connectivity with ping and WAN Tool

You can use the `portcmd --ping` command and the `portcmd --wtool` commands to verify WAN connectivity. If you have configured an SLA, it is not necessary to use the `portcmd --wtool` commands. The following table shows the `portcmd --ping` and `portcmd --wtool` commands.

TABLE 41 Brocade 7840 Switch, Verifying WAN Connectivity with ping and WAN Tool

Action / Command on Left>	Description	Action / Command on Right>	Description

`left> portcmd --ping ge2.dp0 -s 172.16.1.3 -d 172.16.2.3`
 	
`left> portcmd --ping ge3.dp0 -s 172.16.1.4 -d 172.16.2.4`
	Ping ge2 to destination.
	
	Ping ge3 to destination.
		No action required.	

`left> portcmd --wtool 0 create --src 172.16.1.3 --dst 172.16.2.3 --rate 2500000 --admin enable --time 10 --bi-directional --ipsec siteAtoB`	Configure WAN Tool (skip if SLA is configured). Use IPsec if it is configured.	`right> portcmd --wtool 0 create --src 172.16.2.3 --dst 172.16.1.3 --rate 2500000 --admin enable --time 10 --bi-directional --ipsec siteAtoB`	Configure matching WAN Tool (skip if SLA is configured). Use IPsec if it is configured.

TABLE 41 Brocade 7840 Switch, Verifying WAN Connectivity with ping and WAN Tool (continued)

Action / Command on Left>	Description	Action / Command on Right>	Description
`left> portcmd --wtool 0 start`	Start WAN Tool.	No action required.	
`left> portcmd --wtool show --detail`	Display details and verify performance.	No action required.	
`left> portcmd --wtool 0 delete`	Remove WAN Tool for this pair of IP address.	`right> portcmd --wtool 0 delete`	Remove WAN Tool for this pair of IP address
`left> portcmd --wtool 0 create --src 172.16.1.4 --dst 172.16.2.4 --rate 2500000 --admin enable --time 10 --bi-directional --ipsec siteAtoB`	Set up and repeat for second pair of IP addresses.	`right> portcmd --wtool 0 create --src 172.16.2.4 --dst 172.16.1.4 --rate 2500000 --admin enable --time 10 --bi-directional --ipsec siteAtoB`	Set up and repeat for second pair of IP addresses

Creating Tunnel and Circuit Configurations

After verifying WAN connectivity, create the tunnel and circuit configurations between the left-side and right-side locations.

TABLE 42 Brocade 7840 Switch, Tunnel and Circuit Configuration

Action / Command on Left>	Description	Action / Command on Right>	Description

`left> portcfg fciptunnel 24 create --ipsec siteAtoB --sla primaryNet --local-ip 172.16.1.3 --remote-ip 172.16.2.3 --min-comm-rate 2500000 --max-comm-rate 5000000`
 	
`left> portcfg fcipcircuit 24 create 1 --sla primaryNet --local-ip 172.16.1.4 --remote-ip 172.16.2.4 --min-comm-rate 2500000 --max-comm-rate 5000000`
	Create the tunnel and circuit from left side to right side.
	
	IPsec is applied to the tunnel. SLA is applied to both tunnel and circuit. (IPsec and SLA are optional)
		`right> portcfg fciptunnel 24 create --ipsec siteAtoB --sla primaryNet --local-ip 172.16.2.3 --remote-ip 172.16.1.3 --min-comm-rate 2500000 --max-comm-rate 5000000`
		
		`right> portcfg fcipcircuit 24 create 1 --sla primaryNet --local-ip 172.16.2.4 --remote-ip 172.16.1.4 --min-comm-rate 2500000 --max-comm-rate 5000000`
			Create the tunnel and circuit from right side to left side.
			
			IPsec is applied to the tunnel. SLA is applied to both tunnel and circuit. (IPsec and SLA are optional.)

`both> portshow fciptunnel --circuit`	Verify the tunnel and circuit configurations. The tunnel should show Up. If the optional SLA is configured, the tunnel will show Testing until the SLA completes.	`both> portshow fciptunnel --circuit`	Verify the tunnel and circuit configurations. The tunnel should show Up. If the optional SLA is configured, the tunnel will show Testing until the SLA completes.

Configuring LAN Ports

The LAN ports and IPIFs are configured with a link aggregation group (LAG) for the ports.

TABLE 43 Brocade 7840 Switch, LAN Port Configuration

Action / Command on Left>	Description	Action / Command on Right>	Description

`both> portcfgge ge6 --set -lan`
 	
`both> portcfgge ge7 --set -lan`
 	
`both> portcfgge ge8 --set -lan`
 	
`both> portchannel --create LAN1 -type static`
 	
`both> portchannel --add LAN1 -port ge6-ge8`
 	
NOTE
With FOS 8.2.1, you can represent a port range with comma-separated ports (i.e., **portchannel --add slag101 -port ge16, ge 17**.
	On both sides, ports ge6, ge7, and ge8 are configured in a static LAG.	`both> portcfgge ge6 --set -lan`
		
		`both> portcfgge ge7 --set -lan`
		
		`both> portcfgge ge8 --set -lan`
		
		`both> portchannel --create LAN1 -type static`
		
		`both> portchannel --add LAN1 -port ge6-ge8`
			On both sides, ports ge6, ge7, and ge8 are configured in a LAG.

`left> portcfg ipif lan.dp0 create 10.0.0.2/24`	An IPIF address is assigned to the LAN port and set up as a subnet address.	`right> portcfg ipif lan.dp0 create 192.168.1.2/24`	An IPIF address is assigned to the LAN port and set up as a subnet address.

`both> portenable ge6`
 	
`both> portenable ge7`
 	
`both> portenable ge8`
	Enable the ports.	`both> portenable ge6`
		
		`both> portenable ge7`
		
		`both> portenable ge8`
			Enable the ports.

Configuring Traffic Control Lists

Traffic Control Lists (TCLs) must be configured. If none is configured, no traffic will pass through the ports.

In this example, TCLs are configured to allow traffic that matches the IP address of the storage arrays on the left side and the right side. No other traffic will pass through the Brocade 7840 Switch LAN interfaces. After the TCL configuration is complete and the ports are enabled, the Brocade 7840 Switch configuration is complete for the network topology shown in the preceding figure.

TABLE 44 Brocade 7840 Switch, TCL Configuration

Action / Command on Left>	Description	Action / Command on Right>	Description
`left> portcfg tcl storageArray1 create --priority 10 --admin enable --target 24 --src-addr 10.0.0.50/32 --dst-addr 192.168.1.50/32 --l4proto TCP`	The /32 prefix corresponds to a netmask of 255.255.255.255, so the entire address must match.	`right> portcfg tcl storageArray1 create --priority 10 --admin enable --target 24 --src-addr 192.168.1.50/32 --dst-addr 10.0.0.50/32 --l4proto TCP`	The /32 prefix corresponds to a netmask of 255.255.255.255, so the entire address must match.

Configuring Vendor IP Storage Addresses and Routes

Each vendor uses equipment-specific commands to configure network addresses, IP routes, and other information. Refer to the vendor documentation for specific instructions on how to configure the equipment. Generic information is provided in the example.

TABLE 45 Vendor IP Addresses and Route, Generic Configuration

Action / Command on Left>	Description	Action / Command on Right>	Description

`left> route add default gateway 10.0.0.1`
 	
`left> route add 192.168.1.50 netmask 255.255.255.255 gateway 10.0.0.2`
	Generic commands to add a default gateway for the IPIF interface and IP route information to the right side. Refer to vendor documentation for specific commands.	`right> route add default gateway 192.168.1.1`
		
		`right> route add 10.0.0.50 netmask 255.255.255.255 gateway 192.168.1.2`
			Generic commands to add a default gateway for the IPIF interface and IP route information to the left side. Refer to vendor documentation for specific commands.

Configuring Brocade FX8-24 Crossport Features

A crossport is the non-local DP XGE port. The Brocade FX8-24 Blade provides two data processing (DP) complexes identified as DP0 and DP1. Each DP has a local 10Gb/s XGE port, xge0 and xge1 that corresponds to DP0 and DP1. You can configure a DP to use its non-local XGE port, which is done to provide an alternate traffic path if the local XGE port fails for some reason. Crossports are supported only on the Brocade FX8-24 Blade and are available when the blade is configured for 10Gb/s mode.

For DP0 and its local xge0 port, the crossport is xge1. Likewise, for DP1 and its local xge1 port, the crossport is xge0.

Typically, IP interface addresses (IPIFs) used by ge0 through ge9 and xge1 are used for any circuits that use VE_Ports 12 through 21. The xge1 port is the local XGE interface for VE_Ports 12 through 21. Likewise, IP addresses configured for xge0 are used by circuits for VE_Ports 22 through 31.

Configure a crossport by assigning an IP address to a remote XGE port that can be used by the local XGE port. For example, assigning an IP address to xge0 as a crossport makes the address available on the remote xge0 for VE_Ports 12 through 21 on the local xge1.

You can also assign IP routes (iproutes) used by the local port, VLAN tagging, and circuits with metrics to the remote XGE port to allow failover to the crossports.

Crossports contain the IP interface addresses (IPIFs) and IP routes (iproutes) that belong to the remote interface. To use crossports, both XGE ports must be configured in 10Gb/s mode.

Configuring Crossports on the Brocade FX8-24 Blade

Configure crossport XGE port addresses using the `--crossport` or `-x`(shorthand) options for the `portcfg ipif` command, as shown in the following example. Note that in this example, IP address 192.168.11.20 is made available on xge1 for circuits on VE_Ports 12 through 21 on local port xge1.

1. Configure an interface for the local XGE port (xge1).

```
switch:admin> portcfg ipif 8/xge1 create 192.168.10.20 netmask
255.255.255.0 mtu 1500
Operation Succeeded
```

2. Configure interface 192.168.11.20 on remote port xge1 to be available for VE_Ports 12 through 21.

```
switch:admin> portcfg ipif 8/xge1 create 192.168.11.20 netmask 255.255.255.0
mtu 1500 --crossport
```

or

```
switch:admin> portcfg ipif 8/xge1 create 192.168.11.20 netmask 255.255.255.0
mtu 1500 -x
```

The output from `portshow ipif` for xge1 shows the crossport tag.

```
switch:admin> portshow ipif 8/xge1
Port         IP Address            / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------
 8/xge1      192.168.10.20         / 24   1500  0     U R M
 8/xge1      192.168.11.20         / 24   1500  0     U R M X
```

Delete the crossport address using the `delete` option instead of the `create` option for the `portcfg ipif` command.

```
switch:admin>portcfg ipif 8/xge1 delete 192.168.11.20 netmask 255.255.255.0
mtu 1500 -x
```

When deleted, output from `portshow ipif` for xge1 will not show the crossport.

```
switch:admin> portshow ipif 8/xge1
Port         IP Address            / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------
 8/xge1      192.168.10.20         / 24   1500  0     U R M
```

NOTE

If the `crossport` or `-x` option is not specified and the address is on the crossport, the command will fail with an unknown IP address. The command will also fail if the `crossport` option is specified and the address is not on the crossport.

Display local and crossport interface configuration details for a specific XGE port using the`portshow ipif` *`slot/xgeport`* command. Use the `portshow ipif` command to display details for all interfaces.

```
portshow ipif 8/xge0
portshow ipif
```

Crossports and Failover

A cross-port is the non-local DP XGE port. The Brocade FX8-24 Blade provides two data processing (DP) complexes identified as DP0 and DP1. Each DP has a local 10Gb/s XGE port, xge0 and xge1 that corresponds to DP0 and DP1. You can configure a DP to use its non-local XGE port, which is done to provide an alternate traffic path if the local XGE port fails for some reason. Cross-ports are supported only on the Brocade FX8-24 Blade and are available when the blade is configured for 10Gb/s mode.

For DP0 and its local xge0 port, the cross-port is xge1. Likewise, for DP1 and its local xge1 port, the cross-port is xge0.

Typically, IP interface addresses (IPIFs) used by ge0 through ge9 and xge1 are used for any circuits that use VE_Ports 12 through 21. The xge1 port is the local XGE interface for VE_Ports 12 through 21. Likewise, IP addresses configured for xge0 are used by circuits for VE_Ports 22 through 31.

Configure a cross-port by assigning an IP address to the remote XGE port that can be used by the local XGE port. For example, assigning an IP address to xge0 as a cross-port makes the address available on the remote xge0 for VE_Ports 12 through 21 on the local xge1.

You can also assign IP routes (iproutes) used by the local port, VLAN tagging, and circuits with metrics to the remote XGE port to allow failover to the cross-ports.

Cross-ports contain the IPIFs and IP routes that belong to the remote interface. To use crossports, both XGE ports must be configured in 10Gb/s mode.

Configuring Lossless Failover with Crossports on a Brocade FX8-24 Blade

There are two types of configurations supported:

- Active-active: Data will be sent on both 10GbE ports to initiate weighted balancing of the batches across the trunk circuits.
- Active-passive: Data fails over using LLL to a passive circuit (one with a higher metric) if all active lower metric circuit paths fail.

You must establish a metric for failover circuits. If no metric is provided, circuit data will be sent through both ports and the load will be balanced. Circuits have a default metric of 0. A metric of 1 is required for a standby (passive) circuit.

Active-active Configuration

The following example shows an active-active configuration in which two circuits are configured with the same metric, one circuit going over xge0 and the other circuit going over the crossport using xge1 as the external port. The metric values of both the circuits are the same (default value), so both circuits send data. The load is balanced across these circuits. The effective bandwidth of the tunnel in this example is 2 Gb/s.

1. Configure an IP address on interface xge0.

```
portcfg ipif 8/xge0 create 192.168.11.20/24 mtu 1500
```

2. Configure an IP address on crossport interface xge1.

```
portcfg ipif 8/xge1 create 192.168.10.10/24 mtu 1500 -x
```

3. Create a tunnel with one circuit going over xge0.

```
portcfg fciptunnel 8/22 create --remote-ip 192.168.11.20 --local-ip 192.168.11.21 -b 2750000 -B
2750000
```

4. Add another circuit, going over crossport xge1, to the tunnel.

```
portcfg fcipcircuit 8/22 create 1 --remote-ip 192.168.10.10 --local-ip 192.168.10.11 -b 1000000 -B
1000000
```

5. Display local and crossport interface details for xge0.

```
portshow ipif 8/xge0
```

> NOTE
> If the source and destination addresses are on different subnets, you must configure IP routes to the destination addresses. See Configuring IP Route on page 104.

Active-passive Configuration

The following example shows an active-passive configuration in which two circuits are configured with different metrics, one circuit going over xge0 and the other circuit going over the crossport using xge1 as the external port. In this example, circuit 1 is a failover circuit because it has a higher metric. When circuit 0 goes down, the traffic is failed over to circuit 1. The effective bandwidth of the tunnel in this example is 1 Gb/s.

1. Configure an IP address on interface xge0.

```
portcfg ipif 8/xge0 create 192.168.11.20/24  mtu 1500
```

2. Configure an IP address on crossport interface xge1.

```
portcfg ipif 8/xge1 create 192.168.10.10/24 mtu 1500 -x
```

3. Create a tunnel with one circuit going over xge0.

```
portcfg fciptunnel 8/22 create --remote-ip 192.168.11.21 --local-ip 192.168.11.20 -b 2750000 -B
2750000 --metric 0
```

4. Add another circuit, going over crossport xge1, to the tunnel.

```
portcfg fcipcircuit 8/22 create 1 --remote-ip 192.168.10.10 --local-ip 192.168.10.11 -b 1000000 -B
1000000 --metric 1
```

5. Display local and crossport interface details for xge0.

```
portshow ipif 8/xge0
```

NOTE

If the source and destination addresses are on different subnets, you must configure IP routes to the destination addresses. See Configuring IP Route on page 104.

Configuring IP Routes with Crossports

You can configure IP routes with crossport addresses using the `portcfg iproute [slot/port] create` command, as in the following example. In the example, the route will be available for tunnel circuits using VE_ports 12 through 21.

```
portcfg iproute 8/xge0 create 1.1.1.0 netmask 255.255.255.0 192.168.11.250 --crossport
```

or

```
portcfg iproute 8/xge0 create 1.1.1.0 netmask 255.255.255.0 192.168.11.250 -x
```

Delete the route using the `delete` option instead of the `create` option for the `portcfg iproute` command.

```
portcfg iproute 8/xge0 delete 1.1.1.0 netmask 255.255.255.0 -x
```

NOTE

If the `crossport` or `-x` options are not specified and the address is on the crossport, the command will fail with an unknown IP address. The command will also fail if the `crossport` option is specified and the address is not on the crossport.

Display the static IP routes for the local interface and crossport using the `portshow iproute` command:

```
portshow iproute 1/xge0
```

Display the IP interface configured for the local interface and crossport using the `portshow ipif` command.

```
portshow ipif 1/xge0
```

For more information on configuring an IP route, see Configuring IP Route on page 104.

NOTE

If an XGE port has both regular and crossport addresses configured on it, and they use the same IP route, then two routes must be configured: a regular route and an identical route on the crossport.

Configuring VLAN Tags with Crossports

Add entries with crossport addresses to the VLAN tag table using the `portcfg vlantag [slot/port] add` command, as in the following example. This example allows VE_ports 12 through 21 to use the configured local IP interface with this VLAN tag.

```
portcfg vlantag 8/xge0 add 192.168.11.20 200 1 --crossport
```

or

```
portcfg vlantag 8/xge0 add 192.168.11.20 200 1 -x
```

Delete the VLAN tag using the `delete` option instead of the `add` option for the `portcfg vlantag` command.

```
portcfg vlantag 8/xge0 delete 192.168.11.20 200 1 -x
```

Display the VLAN tag configuration using the `portshow vlantag` command.

> NOTE
> To tag Class F traffic or data path traffic, use the `-v` or `-vlan-tagging` options for the `fcipcircuit create` or `fcipcircuit modify` command. The `portcfg vlantag` command is primarily used for ping and traceroute operations and not for tunnels and circuits.

For more information on managing VLAN tags, see Configuring VLANs on page 106.

Displaying VLAN the Tag Configuration Using the portshow vlantag Command

Following is an example of displaying VLAN tagging information for port 0 on blade 8.

```
portshow vlantag 8/xge0
```

For more information on managing VLAN tags, see Configuring VLANs on page 106.

For more information on using Fabric OS commands, optional arguments, and command output, refer to the *Brocade Fabric OS Command Reference*.

Using ping with Crossports

You can ping crossport addresses, as in the following example. Note that if the `crossport` or `x` options are not specified and the address is on the crossport, the `portCmd` command will fail with an unknown IP address.

```
portcmd --ping 8/xge0 -s 192.168.11.20 -d 1.1.1.1 --crossport
```

or

```
portcmd --ping 8/xge0 -s 192.168.11.20 -d 1.1.1.1 -x
```

When using VLANs, VLAN tagging ensures that test traffic traverses the same path as real traffic. A VLAN tag entry for both the local and remote sides of the route must exist before using the `portCmd --ping` command. See Configuring VLANs on page 106 for details.

Using traceroute with Crossports

You can trace a route to a crossport address, as in the following example. Note that if the `crossport` or `x` options are not specified and the address is on the crossport, the `portCmd` command will fail with an unknown IP address. The command will also fail if the `x` option is specified and the address is not on the crossport.

```
portcmd --traceroute 8/xge0 -s 192.168.11.20 -d 1.1.1.1 --crossport
```

or

```
portcmd --traceroute 8/xge0 -s 192.168.11.20 -d 1.1.1.1 -x
```

When using VLANs, VLAN tagging ensures that test traffic traverses the same path as real traffic. A VLAN tag entry for both the local and remote sides of the route must exist before you can use the `portCmd --traceroute` command. Refer to Configuring VLANs on page 106 for details.

For more information on using traceroute, see #unique_187.

Using Logical Switches

Configuring tunnels and other components in switches enabled for Virtual Fabrics is somewhat different than on switches not enabled for Virtual Fabrics. The information that follows provides a brief overview of common logical switch concepts and terminology followed by the specifics of configuring logical switches.

> NOTE
> The Brocade 7810 switch does not support logical switches.

Logical Switch Overview

The logical switch feature allows you to divide a physical chassis into multiple fabric elements. Each of these fabric elements is referred to as a logical switch. Each logical switch functions as an independent self-contained FC switch. Each chassis can have multiple logical switches. In Fabric OS 8.1.0 and later, up to a total of 16 logical switches can be supported in a single Gen6 chassis or on a single Brocade SX6 blade. The Brocade 7840 switch can support up to four logical switches.

Default Logical Switch

Virtual Fabrics allow Ethernet ports in the default switch to be shared among VE_Ports in any logical switch. To use the Virtual Fabrics features, you must first enable Virtual Fabrics on the switch. Enabling Virtual Fabrics creates a single logical switch in the physical chassis. This logical switch is called the default logical switch, and it initially contains all of the ports in the physical chassis. After you enable Virtual Fabrics, you can create additional logical switches. The number of logical switches that you can create depends on the switch model.

After you create logical switches, the chassis appears as multiple independent logical switches. All of the ports continue to belong to the default logical switch until you explicitly move them to other logical switches. The default logical switch always exists. You can add and delete other logical switches, but you cannot delete the default logical switch unless you disable Virtual Fabrics.

Creating Logical Switches

To create logical switches and logical fabrics, you must perform the following steps.

1. Enable Virtual Fabrics mode on the switch using instructions in the "Managing Virtual Fabrics" chapter of the *Brocade Fabric OS Administration Guide*.

2. Configure logical switches to use basic configuration values using instructions in the "Managing Virtual Fabrics" chapter of the *Brocade Fabric OS Administration Guide*.
3. Create logical switches using instructions for creating a logical switch or base switch in the "Managing Virtual Fabrics" chapter of the *Brocade Fabric OS Administration Guide*.

Port Assignment

Initially, all ports belong to the default logical switch. When you create additional logical switches, they are empty and you can assign ports to those logical switches. As you assign ports to a logical switch, the ports are moved from the default logical switch to the newly created logical switch. Following are some requirements for assigning ports:

- A given port can be in only one logical switch.
- You can move ports from one logical switch to another.
- A logical switch can have as many ports as are available in the chassis.
- Ports with defined configuration settings in a logical switch or the default switch cannot be moved to another logical switch without first deleting the current settings. For example, you cannot move a VE_ Port with a defined tunnel in the default switch or a logical switch to a different logical switch until you delete the circuits and the tunnel in the logical switch currently containing the port that you want to move. Similarly, you cannot move a GE_Port between logical switches until all IP routes and IP interfaces have been deleted in the logical switch currently containing the port that you want to move.

Use the`lsCfg -config` *`slot/ge_port`* command to move ports from one logical switch to a different logical switch. The FID is the fabric ID of the logical switch to where you want to move the ports. The ports are automatically removed from the logical switch where they are currently assigned.

As a recommended best practice, leave Ethernet interfaces in the default logical switch and do not move them to another logical switch. There is no reason to move them because of the Ethernet Port Sharing (EPS) feature. A VE_Port in any logical switch context can use an Ethernet interface in the default switch. In addition, by moving a physical port from the default switch to a logical switch, it will not be available to tunnels configured in other logical switches. See Ethernet Port Sharing on page 182 for details.

Logical Switches and Fabric IDs

When you create a logical switch, you must assign it a fabric ID (FID). The fabric ID uniquely identifies the logical switch within a chassis and indicates the fabric to which the logical switch belongs. You cannot define multiple logical switches with the same fabric ID within the chassis. A logical switch in one chassis can communicate with a logical switch in another chassis (or to a switch not enabled for logical switches) only if the switches have the same fabric ID (FID). The default logical switch is initially assigned FID 128, which can be changed.

Only logical switches with the same FID can form a logical fabric. If you connect two logical switches with different FIDs, the link between the switches segments.

Create logical switches 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 to the `lsCfg` command in the *Brocade Fabric OS Command Reference*.

Logical Switch Contexts

You can configure features or perform other tasks on a specific logical switch as you would any Fibre Channel switch by entering commands while in the "context" of that logical switch, which is the FID of the logical switch. Note that "128" is sometimes referred to the context for the default switch because that is the initial FID of the default switch when you enable Virtual Fabrics. However, this FID can be changed.

There are two methods for changing to the context of a specific logical switch so that you can perform configuration or other tasks:

- Use the `setcontext` *`fabricID`* command. This changes the context to a specific logical switch and changes the command line prompt to reflect the new FID. Any commands entered at this prompt are initiated on the logical switch with that FID.
- Use the `fosexec --fid` *`FID`* `-cmd "`*`command-string`*`"` to initiate a specific command on a specific logical switch, where *command-string* is the command string you want to perform.

Using the `fosexec` command to issue the command string on any logical switch runs the specified FOS command string on the specified logical switch, whereas using `setcontext` command issues the command string in only the current logical switch.

Connecting Logical Switches

A logical fabric is a fabric that contains at least one logical switch. You can connect logical switches to non-virtual fabrics switches and to other logical switches using two methods:

- Through inter-switch links (ISLs). For extension traffic, the ISL connection is through a tunnel.
- Through base switches and extended ISLs (XISLs). This is supported by the Brocade SX6 blade, the Brocade 7840 switch, and the Brocade FX8-24 blade. See Enabling XISL for VE_Ports (Brocade FX8-24 Blade/7840 Switch/SX6 Blade) on page 190.

For More Information on Virtual Fabrics

For more detail on managing and configuring virtual fabrics, refer to the chapter on managing Virtual Fabrics in the *Brocade Fabric OS Administration Guide*.

Considerations for Logical Switches

Before creating IPIFs, IP routes, tunnels, and circuits, follow procedures for creating logical switches as outlined in Logical Switch Overview on page 180 and as detailed in the "Managing Virtual Fabrics" chapter in the *Brocade Fabric OS Administration Guide*. Use the following information and instructions for creating tunnels and other components on logical switches.

Ethernet Port Sharing

In Fabric OS 7.0 and later, VE_Ports in different logical switches can share a single Ethernet port (1 GbE, 10 GbE, or 40 GbE) located on the default switch. As a best practice, leave Ethernet interfaces in the default switch even if you will only use a single virtual fabric logical switch. If new VF logical switches are added and need to use the Ethernet interface, then the Ethernet interface does not have to be moved back to the default switch.

NOTE
For Fabric OS 7.4.0 and later versions that support IP Extension, an Ethernet port that is used as a LAN port must be in the default switch.

NOTE
For Fabric OS versions before Fabric OS 7.0, in order to use a Ethernet port for a tunnel, that port must be in the same logical switch as the tunnel's VE_Port.

With Ethernet port sharing, you can have the following configuration, as an example:

- Default switch has port GbE0.
- Logical switch 1 has VE_Port 24 (or tunnel 24), which has a circuit over GbE0.
- Logical switch 2 has VE_Port 25 (or tunnel 25), which also has a circuit over GbE0.
- There are no LAN ports associated with GbE0.

All of the committed-rate restrictions and bandwidth sharing of the Ethernet ports for ARL remain the same for shared ports in the logical switches. VE_Ports created from shared Ethernet ports initiate as regular VE ISLs in their respective logical switches.

When IPIFs are created for physical ports (including crossports) located in the default switch, these IP interfaces can be used by circuits assigned to tunnels created in other logical switches. This means that multiple VE_Ports in multiple logical switches can use the same Ethernet port. Although multiple circuits can use the same Ethernet port, these circuits can be differentiated in the IP network using VLAN tags or access control lists (ACLs) set for the source and destination IP addresses in the circuit. See Configuring VLANs on page 106 for more information on using VLAN tagging for Extension features.

Limitations of Ethernet Port Sharing

Note the following limitations of port sharing:

- Only Ethernet ports in the default switch can be shared by VE_Ports in different logical switches. A Ethernet port in a non-default switch can only be used by VE_Ports in that same logical switch.
- The GbE ports in other logical switches or ports on the base switch cannot be shared by ports in different logical switches.
- Tunnels created on Brocade FX8-24 blades with a mix of dedicated ports (ports within the same logical switch) and shared ports (ports in the default switch) are not supported.
- When using shared Ethernet interfaces between the default switch and other logical switches, if the default switch is disabled, the Ethernet ports in the default switch will also be disabled. This will impact all tunnels in the other logical switches using the Ethernet interfaces.

Port Sharing Example

This section illustrates an example of port sharing on a Brocade FX8-24 blade. The following output for the `portshow ipif all` command illustrates IP interfaces, IP routes, and crossports configured for ports in the default logical switch and tunnels and circuits on two different logical switches that use these configurations.

Note the following about the configuration detailed in the output:

- This example is for port sharing configuration on a Brocade FX8-24 blade.
- There are three logical switches:
  - LS 0 has FID 128 and is the default switch.
  - LS 2 has FID 50.
  - LS 4 has FID 70.
- IP interfaces and IP routes for these IPIFs were created for xge0 and xge1. The `portcfg --ipif` and `portcfg --iproute` commands were issued in the default logical switch context where the ports reside. See Configuring IP Interfaces and IP Routes on page 185 for more information.
- Crossports were configured for both xge0 and xge1 on the default switch. See Crossports and Failover on page 59 for more information.
- A tunnel with VE_Port 22 and circuits was created on LS 2. VE_Port 22 was first moved to LS 2, and the `portcfg fciptunnel` commands to configure the tunnel and circuits were issued in the context for LS 2 (FID 50). See Moving Ports Between Logical Switches on page 188 and Configuring Tunnels and Circuits on page 185 for more information.
- A tunnel with VE_Port 12 and circuits was created on LS 4. VE_Port 12 was first moved to LS 4, and the `portcfg fciptunnel` commands to configure the tunnel and circuits were issued in the context for LS 4 (FID 70). See Moving Ports Between Logical Switches on page 188 and Configuring Tunnels and Circuits on page 185 for more information.

```
CURRENT CONTEXT -- LS: 0, FID: 128 *NOTE this
is the default switch.*
switch:admin> portshow ipif
```

```
Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
8/xge0       10.108.0.90                   / 24   1500  n/a   U R M
8/xge0       10.108.0.91                   / 24   1500  n/a   U R M
8/xge0       10.108.0.92                   / 24   1500  n/a   U R M X
8/xge0       10.108.0.93                   / 24   1500  n/a   U R M X
8/xge1       10.108.1.90                   / 24   1320  n/a   U R M
8/xge1       10.108.1.91                   / 24   1320  n/a   U R M
8/xge1       10.108.1.92                   / 24   1320  n/a   U R M X
8/xge1       10.108.1.93                   / 24   1320  n/a   U R M X
-------------------------------------------------------------------------------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport

switch:admin> portshow iproute

Port         IP Address               / Pfx  Gateway               Flags
-------------------------------------------------------------------------------
8/xge0        10.108.0.0               / 24   *                     U C
8/xge0        10.108.0.91              / 32   *                     U C
8/xge0        10.108.100.0             / 24   10.108.0.250          U G S
8/xge0        10.108.0.0               / 24   *                     U C X
8/xge0        10.108.0.93              / 32   *                     U C X
8/xge0        10.108.100.0             / 24   10.108.0.250          U G S X
8/xge1        10.108.1.0               / 24   *                     U C
8/xge1        10.108.1.91              / 32   *                     U C
8/xge1        10.108.101.0             / 24   10.108.1.250          U G S
8/xge1        10.108.1.0               / 24   *                     U C X
8/xge1        10.108.1.93              / 32   *                     U C X
8/xge1        10.108.101.0             / 24   10.108.1.250          U G S X
-------------------------------------------------------------------------------
Flags: U=Usable G=Gateway H=Host C=Created(Interface)
        S=Static L=LinkLayer X=Crossport

CURRENT CONTEXT -- LS: 2, FID: 50 *Note that this is one of the logical
switches (not the default switch).*
portshow fciptunnel all -c:
-------------------------------------------------------------------------------
Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt  Met
-------------------------------------------------------------------------------
1/22   -         Up      cft----   14d18h  226.60    2.73    5      -      -
1/22   0 1/xge0  Up      ---4v-s 7d17h34m   64.80    0.78    7  1000/3000  0
1/22   1 1/xge0  Up      ---4v-s  7d5h24m   48.59    0.59    7  1000/2000  0
1/22   2 1/xge1  Up      ---4vxs 7d17h34m   64.60    0.78    7  1000/3000  0
1/22   3 1/xge1  Up      ---4vxs  7d5h24m   48.60    0.58    7  1000/2000  0
-------------------------------------------------------------------------------
Flags (tunnel): M=MainTunnel  L=LocalBackup  R=RemoteBackup
                i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                A=AdvCompr L=LZCompr d=DeflateCompr  D=AggrDeflateCompr
     (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU 4=IPv4 6=IPv6
                ARL a=Auto r=Reset s=StepDown t=TimedStepDown


CURRENT CONTEXT -- LS: 4, FID: 70 *Note that this is a different logical switch
(and not the default switch).*
portshow fciptunnel all -c     :
-------------------------------------------------------------------------------
Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt  Met
-------------------------------------------------------------------------------
1/12   -         Up      c--F---   19d15h    0.00    0.00    1      -      -
1/12   0 1/xge0  Up      ---4vxs 7d17h34m    0.00    0.00    3  1000/3000  0
1/12   1 1/xge0  Up      ---4vxs  7d5h24m    0.00    0.00    4  1000/2000  1
1/12   2 1/xge1  Up      ---4v-s 7d17h34m    0.00    0.00    3  1000/3000  0
1/12   3 1/xge1  Up      ---4v-s  7d5h24m    0.00    0.00    4  1000/2000  1
-------------------------------------------------------------------------------
Flags:  tunnel: c=compression m=moderate compression a=aggressive compression
                 A=Auto compression f=fastwrite t=Tapepipelining F=FICON
                 T=TPerf i=IPSec l=IPSec Legacy
Flags: circuit: s=sack v=VLAN Tagged x=crossport 4=IPv4 6=IPv6
                 L=Listener I=Initiator
```

Configuring IP Interfaces and IP Routes

The following example configures IP interfaces (IPIF) and IP routes (iproutes) for ports that reside on the default switch and creating tunnels and circuits on a different logical switch that use these IP interfaces.

You must issue the `portcfg ipif` and `portcfg iproute` commands in the logical switch context where the Ethernet port resides. If the Ethernet port is in the default switch, then the commands must be entered from the default switch context. If the Ethernet ports are in a logical switch other than the default switch, you must issue the commands in that context. In the latter case, the Ethernet ports cannot be used by tunnels created in any other logical switch in the chassis.

In the following example, port ge2.dp0 is on a Brocade 7840 (the default switch with a "default" switch FID of 128).

1. If you are in a different logical switch context than the default switch, set the context to 128 using the `setcontext 128` command.

```
switch:admin> setcontext 128
```

2. Enter the `portcfg ipif` command to create the interface on port ge2 on DP0.

```
switch:FID128:admin> portcfg ipif ge2.dp0 create 192.168.2.2 netmask 255.255.255.0 mtu 1320
```

3. Configure an IP route using the `portcfg iproute` command in the FID 128 context.

   The following command creates an IP route to destination network 192.168.2.0 for port ge2.dp0. The route is through local gateway 192.168.12.1.

```
switch:FID128:admin> portcfg iproute ge2.dp0 create 192.168.12.0 netmask 255.255.255.0 192.168.2.1
```

4. Use the `portshow ipif` and `portshow iproute` commands in the FID 128 context to display IPFI and IP route information.

```
switch:FID128:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.2.2                   / 24   1320  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:FID128:admin> portshow iproute

 Port         IP Address               / Pfx  Gateway               Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.2.0              / 24   *                     U C
 ge2.dp0      192.168.2.1              / 32   *                     U H L
 ge2.dp0      192.168.12.0             / 24   192.168.2.1           U G S
-------------------------------------------------------------------------------
 Flags: U=Usable G=Gateway H=Host C=Created(Interface)
        S=Static L=LinkLayer X=Crossport
```

Other than issuing commands for IP interfaces and IP routes from the correct logical switch context, other aspects of the commands used in this procedure are the same as for any switch. For more information, see Configuring IPIF on page 102.

Configuring Tunnels and Circuits

To configure a tunnel on a logical switch other than the default switch, you first must move the VE_Port to the logical switch from the default switch and then create the tunnel and circuits in that logical switch context. You issue the `portcfg fciptunnel` command in the context of the logical switch where the VE_Port resides. In the following example, the VE_Port resides on the default switch with FID 128. A tunnel has not been configured yet using this VE_Port.

The following example steps show how to configure a tunnel between two logical switches. Other than issuing commands to move VE_Ports and to create tunnels and circuits from the correct logical switch context, other aspects of configuring tunnels and circuits are the same for any switch. For more information, see Configuring Extension Tunnels for FCIP on page 119.

1. Use the `lscfg` command to create a logical switch 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.
```

2. Move the VE_Port from the default switch to the logical switch with 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.
```

3. Set the context to the logical switch with FID 60 using the following command.

```
switch:FID128:admin> setcontext 60
```

4. Create a tunnel endpoint on the new logical switch using the IP interface created for port ge2.dp0 on the default logical switch for the blade and a destination address for a remote endpoint. In the following example, the local address (192.168.2.2) is specified first, followed by the remote address (192.168.12.2). ARL minimum (-b) and maximum (-B) committed rates are specified for circuit 0, which is the default circuit created automatically when you configure a tunnel. The example command also enables IP extension and configures spillover.

```
switch_60:FID60:admin> portenable 24
switch_60:FID60:admin> portcfg fciptunnel 24 create --local-ip 192.168.2.2 --remote-ip 192.168.12.2 -
b 4000000 -B 4000000  --ipext enable  -L spillover  -k 1000
Operation Succeeded

switch_60:FID60:admin> portshow fciptunnel all -c

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
-------------------------------------------------------------------------------
 24    -         Up       --------I       9s    0.00    0.00   1      -       -
 24    0 ge2     Up       ----a---4       9s    0.00    0.00   1  4000/4000  0/-
-------------------------------------------------------------------------------
 Flags (tunnel): 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
```

5. Set the context to FID 128 and show the IPIF and IP route information.

```
switch_60:FID60:admin> setcontext 128

switch:FID128:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.2.2                   / 24   1320  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:FID128:admin> portshow iproute

 Port         IP Address               / Pfx  Gateway               Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.2.0              / 24   *                     U C
 ge2.dp0      192.168.2.1              / 32   *                     U H L
 ge2.dp0      192.168.12.0             / 24   192.168.2.1           U G S
-------------------------------------------------------------------------------
 Flags: U=Usable G=Gateway H=Host C=Created(Interface)
        S=Static L=LinkLayer X=Crossport
```

6. Configure an IPIF on ge4.dp0 of the default switch.

```
switch:FID128:admin> portcfg ipif ge4.dp0 create 192.168.4.2 netmask 255.255.255.0 mtu 1360
```

7. Configure an IP route on ge4.dp0 of the default switch.

```
switch:FID128:admin> portcfg iproute ge4.dp0 create 192.168.14.0 netmask 255.255.255.0 192.168.4.1
```

8. Use the `portshow` command to display the created IPIF and IP route information.

```
switch:FID128:admin> portshow ipif

  Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.2.2                   / 24   1320  0     U R M I
 ge4.dp0      192.168.4.2                   / 24   1360  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:FID128:admin> portshow iproute

 Port         IP Address               / Pfx  Gateway               Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.2.0              / 24   *                     U C
 ge2.dp0      192.168.2.1              / 32   *                     U H L
 ge2.dp0      192.168.12.0             / 24   192.168.2.1           U G S
 ge4.dp0      192.168.4.0              / 24   *                     U C
 ge4.dp0      192.168.4.1              / 32   *                     U H L
 ge4.dp0      192.168.14.0             / 24   192.168.4.1           U G S
-------------------------------------------------------------------------------
 Flags: U=Usable G=Gateway H=Host C=Created(Interface)
        S=Static L=LinkLayer X=Crossport
```

9. Set the context to FID 60.

```
switch:FID128:admin> setcontext 60
```

10. Use the `portshow fciptunnel` command to display existing tunnel and circuit information.

```
switch_60:FID60:admin> portshow fciptunnel all -c

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
--------------------------------------------------------------------------------
 24    -         Up      --------I    5m18s    0.00    0.00   1      -       -
 24    0 ge2     Up      ----a---4    5m19s    0.00    0.00   1  4000/4000  0/-
--------------------------------------------------------------------------------
 Flags (tunnel): 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
```

11. Use the `portcfg fcipcircuit` command to create a circuit.

```
switch_60:FID60:admin> portcfg fcipcircuit 24 create 2 --local-ip 192.168.4.2 --remote-ip
192.168.14.2 -b 4000000 -B 4000000 -x 1 -k 1000
```

12. Use the `portshow fciptunnel` command to confirm the new tunnel and circuit information.

```
switch_60:FID60:admin> portshow fciptunnel all -c

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
--------------------------------------------------------------------------------
 24    -         Up      --------I    5m36s    0.00    0.00   1      -       -
 24    0 ge2     Up      ----a---4    5m37s    0.00    0.00   1  4000/4000  0/-
 24    2 ge4     Up      ----a---4       9s    0.00    0.00   1  4000/4000  1/-
--------------------------------------------------------------------------------
 Flags (tunnel): 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
```

Moving Ports Between Logical Switches

To move ports between logical switches, use the following command:

`lscfg --config` *`FID`* `port` *`slot/port`*

- The *FID* variable is the Fabric ID of the logical switch where port is moving to.
- The *slot* number is required for the Brocade FX8-24 Blade and the Brocade SX6 Blade. It is omitted on the Brocade 7840 Switch.
- The *port* number is the FC, VE, or GE port number. For the Brocade FX8-24 Blade, XGE (10-GbE) ports are xge0 or xge1, and GbE ports are ge0-ge9. For the Brocade 7840 Switch and Brocade SX6 Blade, 40-GbE ports are ge0-1, and 10-GbE ports are ge2 through ge17.

The following are considerations for moving ports between logical switches:

- The 1GbE ports (Brocade FX8-24 Blade), 10GbE ports (Brocade 7840 Switch, Brocade SX6 Blade, and Brocade X8-24 Blade), 40GbE ports (Brocade 7840 Switch or Brocade SX6 Blade), and VE_Ports can be part of any logical switch. They can be moved between any two logical switches unless they are members of a circuit configuration.
- Because Ethernet ports and VE_Ports are independent of each other, both must be moved in independent steps. You must delete the configuration on VE_Ports and Ethernet ports before moving them between logical switches.
- You must move a VE_Port from the logical switch where it resides to a new logical switch in order to create a tunnel for the new logical switch.

Displaying Logical Switch Configurations

You can display the logical switch configuration for a switch and the Ethernet ports located in each logical switch using the `lscfg --show -ge` command. The following output shows that all Ethernet ports are located in the default switch (FID 128).

```
switch:admin> lscfg --show -ge

Created switches FIDs(Domain IDs):  128(ds)(80)  60(1)

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 |
```

You can display the logical switch configuration and the VE_Ports assigned to each logical switch using the `lscfg --show` command. The following output shows that besides the default switch with FID 128, other default switches have been created with FID 80 and 60.

Note that some of the VE_Ports have been moved from the default switch to other logical switches. In this example, VE_Port 24 is in logical switch 60 (FID60).

```
switch:admin> lscfg --show

Created switches FIDs(Domain IDs):  128(ds)(80)  60(1)

Port      0     1     2     3     4     5     6     7     8     9
-------------------------------------------------------------------
FID     128 | 128 | 128 | 128 | 128 | 128 | 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 |  60 | 128 | 128 | 128 | 128 | 128 |

Port     30    31    32    33    34    35    36    37    38    39
-------------------------------------------------------------------
FID     128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 |

Port     40    41    42    43
-------------------------------
FID     128 | 128 | 128 | 128 |
```

Brocade 7840 Switch, and Brocade SX6 Considerations and Limitations

Following are considerations and limitations for the Brocade 7840 Switch, Brocade 7810 Switch, and the Brocade SX6 Blade configured to support Virtual Fabrics:

- For Brocade 7840 Switch, or Brocade SX6 Blade, you can make the logical switch a "base" switch if you plan to use an extended inter-switch link (XISL) connection between such switches rather than use separate ISL connections from logical switches.
- Up to four logical switches will support FICON CUP on a Brocade 7840 Switch and up to 16 logical switches will support FICON CUP on a Brocade Gen6 chassis; however, refer to your system qualification letter-specific limits.
- A tunnel from a Brocade 7840 Switch or Brocade SX6 Blade requires a VE_Port in the logical switch and shared Ethernet port either in the default switch or in the same logical switch. A VE_Port in a logical switch and Ethernet port in the default switch is considered best practice.
- On the Brocade SX6 Blade, the GbE ports and VE_Ports must be on the same blade. You cannot mix GbE ports and VE_Ports from different blades into the same logical switch.

When in 10VE mode, all unused VE_Ports must be in the default switch. If the unused VE_Ports are not in the default switch, the VE-Mode cannot be set to 10VE mode. Unused VE_Ports cannot be moved to other logical switches while in 10VE mode. Refer to the following table for available ports in 10VE and 20VE mode.

TABLE 46 20VE and 10VE Mode Ports

Product	20VE Mode VE_Ports	10VE Mode VE_Ports
Brocade 7840 Switch	DP0 - 24 through 33 DP1 - 34 through 43	DP0 - 24 through 28 DP1 - 34 through 38
Brocade SX6 Blade	DP0 - 16 through 25 DP1 - 26 through 35	DP0 - 16 through 20 DP1 - 26 through 30

Brocade FX8-24 Blade Considerations and Limitations

The following are considerations and limitations of Brocade FX8-24 Blades configured to support Virtual Fabrics:

- The number of logical switches that you can create and the limits on logical switch support for FICON CUP depends on the chassis where the Brocade FX8-24 Blade is installed. For example, up to eight logical switches can be configured on a blade installed on DCX 8510 platforms. Refer to your chassis specifications for details.
- For the Brocade FX8-24 Blade, you can make the logical switch a base switch if you are planning on using an extended interswitch link (XISL) connection between base switches instead of using separate ISL connections from logical switches.

Enabling XISL for VE_Ports (Brocade FX8-24 Blade/7840 Switch/SX6 Blade)

Another way to connect logical switches is to use extended interswitch links (XISLs) and base switches. When you divide a chassis or fixed-port switch into logical switches, you can designate one of the switches to be a base switch. A base switch is a special logical switch that is used for interconnecting the physical chassis.

NOTE
The Brocade 7810 Switch does not support "base" switches.

An XISL connection can be created between base switches, instead of using separate ISLs. The base fabric provides the physical connectivity across which logical connectivity will be established. The XISL can carry combined traffic for multiple logical fabrics while maintaining traffic separation for each fabric.

Because of the expense of long-distance links, this feature has particular benefit for the Brocade extension platforms. This feature is supported as follows:

- On tunnels between Brocade FX8-24 Blades running Fabric OS 7.0 and later.
- On tunnels between Brocade 7840 Switches running Fabric OS 7.4.0 and later.
- On tunnels between Brocade SX6 Blades, or Brocade 7840 and 7810 Switches running Fabric OS 8.0.1 and later.
- FICON Emulation features are not supported on a Brocade FX8-24 Blade VE XISL port. FICON Emulation features are supported on a Brocade 7840 Switch or Brocade SX6 Blade VE XISL port.

To create a base switch on the Brocade FX8-24 Blade, use the `-base` option for the `lsCfg` command when creating a logical switch. To use XISL, refer to instructions for configuring a logical switch to use XISLs in the *Brocade Fabric OS Administration Guide*.

For the Brocade FX8-24 Blade, if an XISL is enabled, it is recommended that you do not configure VE_Ports on both the logical switch and the base switch because tunnels support only a maximum of two hops.

Traffic Isolation Zoning

Traffic Isolation Zoning allows you to control the flow of inter-switch traffic by creating a dedicated path for traffic flowing from a specific set of source ports (N_Ports).

You might use Traffic Isolation Zoning for the following scenarios:

- To dedicate an ISL to high-priority, host-to-target traffic.
- To force high volume, low-priority traffic onto a given ISL to limit the effect on the fabric of this high-traffic pattern.
- To ensure that requests and responses of FCIP-based applications such as tape pipelining use the same VE_Port tunnel across a metaSAN.

Traffic isolation is implemented using a special zone, called a Traffic Isolation zone (TI zone). A TI zone indicates the set of N_Ports, E_Ports, and VE_Ports to be used for a specific traffic flow. When a TI zone is activated, the fabric attempts to isolate all inter-switch traffic entering from a member of the zone to only those E_Ports or VE_Ports that have been included in the zone. The fabric also attempts to exclude traffic not in the TI zone from using E_Ports or VE_Ports within that TI zone.

NOTE
Traffic Isolation Zoning is an advanced feature, but does not require a license. A strong understanding of Fabric Shortest Path First (FSPF) is required.

NOTE
Use of Traffic Isolation Zones is not considered best practice in FCIP Extension configurations. Instead, you should consider using Logical Switches and multiple circuit tunnels between connected domains.

For more information and details about configuring TI Zoning, refer to the "Traffic Isolation Zoning" chapter in the *Brocade Fabric OS Administration Guide*.

Zoning

A recommended best practice is to use zoning in your network. Zoning is a fabric-based service that enables you to partition your storage area network (SAN) into logical groups of devices that can access each other.

For example, you can partition your SAN into two zones, “winzone” and “unixzone”, so that your Windows servers and storage do not interact with your UNIX servers and storage. You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions; for example, you can create a temporary zone to back up nonmember devices.

A device in a zone can communicate only with other devices connected to the fabric within the same zone. A device not included in the zone is not available to members of that zone. When zoning is enabled, devices that are not included in *any* zone configuration are inaccessible to all other devices in the fabric.

By contrast, if no zones are defined, each device in the fabric can access every other device in the fabric. In a small network with few devices, this might not be a problem. In large networks with many devices, having all devices accessible can present security issues. That 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 on configuring zones, refer to the “Administering Advanced Zoning" chapter in the *Brocade Fabric OS Administration Guide*.

IP Extension Flow Monitor Overview

IP Extension Flow Monitor allows you to monitor and view LAN-side statistics. By viewing 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 feature separate from the Brocade Flow Vision and Flow Monitor features. IP Extension Flow Monitor is accessed through Fabric OS CLI commands on the Brocade X6 and Brocade 7840 platforms running Fabric OS 8.2.0 or later. The following table lists the types of IP traffic information that you can monitor.

NOTE

Because of a reduction in DP memory relative to the Brocade 7840 Switch and SX6 Blade platforms, the maximum upper limit for VE port-based monitoring of FC on a Brocade 7810 Switch will be lowered (the maximum learned VE flows and static VE flows will be 4 (1 per VE port) and 256, respectively). However, a Brocade 7810 Switch will provide similar support for IPEX flow monitoring; the maximum supported flows (TCP connections; Brocade 7840 Switch and Brocade SX6 Blade (4K flow per DP for current and historical stats) and max supported IP-PAIR (32 IP pairs) will be unchanged.

TABLE 47 IP Extension Flow Monitor IP Traffic

Layer	Monitored information
L2	• VLAN ID Non-IP traffic is not monitored.
L3	• IP source and IP destination address Non-TCP traffic is not monitored. UDP traffic monitor is not supported.
L4	• TCP port information based on TCP port number • TCP port information based on TCP port range • TCP port information based on TCP application type Non-TCP traffic is not monitored. UDP traffic monitor is not supported.

Monitoring Traffic Flows

The primary use of IP Extension Flow monitor is to display statistical information about the flow of L2, L3, and L4 traffic on the ports or addresses. The type of traffic information displayed by 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 or VE port
- Slot (in a multi-slot chassis)
- DP, either DP0 or DP1
- 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 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 the 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 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 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 historical IP pair information as a total of the past 24 hours that is 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 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 resets IP pair statistics and results in loss of current and historical information:

- Firmware upgrade or downgrade
- HCL operations
- DP reset

See Using IP Extension Flow Monitor on page 194 for usage and configuration information.

Using IP Extension Flow Monitor

When you want to view LAN traffic statistics, use the IP Extension Flow Monitor feature.

Before you can use IP Extension Flow Monitor, you must configure IP Extension on the platform. Refer to Configuring IP Extension on page 148 for information. The IP Extension Flow monitor is useful only when 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 in using an account assigned to the admin role. For information on all command options, see *Brocade Fabric OS Command Reference*.

NOTE

NTTCP 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 single platform, for example Brocade X6 Directors, Brocade 7840 Extension Switch, or Brocade 7810 Extension Switch. For additional information on how to configure flow views, refer to the following:

- Configuring a Port-based Flow on page 195
- Configuring an IP Address Flow on page 198
- Configuring a TCP Port Flow on page 200
- Configuring a Flow Using Logical Operators on page 201
- Displaying Historical Flow Statistics on page 202

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. For information on how to modify the IP pair display, refer to the following:

- Displaying IP Pair Detail on page 206
- Displaying IP Pair History on page 207
- Resetting IP Pair Statistics on page 208

The following steps show how to display the default output for flows and IP pairs.

1. To display 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
 DP1      0       0       0       0
-----------------------------------------------------
 ActTCP=Active TCP Conns  ExdTCP=Exceeded TCP Conn Cnt
```

   The default display will show 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 display 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. You can 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, see *Brocade Fabric OS Command Reference*.

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. The following steps show how to name and create port-based flows.

NOTE

The flow name is case-sensitive. “GE3” is not the same as “ge3.”

1. To create a port-based flow, use the `portcfg lan-stats --flow create` command. This step shows how to create 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

 portcfgshow fciptunnel

 Tunnel Circuit   AdminSt Flags
-------------------------------------------------------------------------------
 25    -         Enabled --------I
 35    -         Enabled -------DI
-------------------------------------------------------------------------------
 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, use the `portcfg lan-stats --flow create` command.

```
switch:admin> portcfg lan-stats --flow test create --port 25
Operation Succeeded.
```

3. To view the flow names that you created on the platform, use the `portshow lan-stats --flow` command.

```
switch:admin> portshow lan-stats --flow

 Name                Flow Info
-------------------------------------------------------------------------------
 APP-flow            (TCP:3225-3226)
 host223             (port:25)
 test                (port:25)
-------------------------------------------------------------------------------
```

   The Flow Info column shows the flow definition associated with the flow name.

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)
-------------------------------------------------------------------------------
 DP0    20   192.168.10.76    192.168.20.38    53023 59780 TCP   19.6m   0
 DP0    23   192.168.10.76    192.168.20.38    53012 59780 TCP   19.7m   0
 DP0    21   192.168.20.38    192.168.10.76    59779 52994 TCP   0       0
 DP0    25   192.168.10.76    192.168.20.38    53016 59780 TCP   19.7m   0
 DP0    26   192.168.10.76    192.168.20.38    53025 59780 TCP   19.7m   0
 DP0    30   192.168.10.76    192.168.20.38    53021 59780 TCP   19.9m   0
 DP0    27   192.168.10.76    192.168.20.38    53020 59780 TCP   19.8m   0
 DP0    13   192.168.10.76    192.168.20.38    53018 59780 TCP   19.7m   0
 DP0    8    192.168.10.76    192.168.20.38    53003 59780 TCP   19.7m   0
 DP0    15   192.168.10.76    192.168.20.38    53001 59780 TCP   19.6m   0
 DP0    12   192.168.10.76    192.168.20.38    52997 59780 TCP   19.7m   0
 DP0    18   192.168.10.76    192.168.20.38    53008 59780 TCP   19.8m   0
 DP0    16   192.168.10.76    192.168.20.38    53024 59780 TCP   19.7m   0
 DP0    10   192.168.10.76    192.168.20.38    53006 59780 TCP   19.7m   0
 DP0    22   192.168.10.76    192.168.20.38    53011 59780 TCP   19.6m   0
 DP0    14   192.168.10.76    192.168.20.38    53007 59780 TCP   19.3m   0
 DP0    7    192.168.10.76    192.168.20.38    53013 59780 TCP   19.7m   0
 DP0    6    192.168.10.76    192.168.20.38    53002 59780 TCP   19.7m   0
 DP0    4    192.168.10.76    192.168.20.38    52999 59780 TCP   19.6m   0
 DP0    1    192.168.10.76    192.168.20.38    53004 59780 TCP   19.7m   0
 DP0    3    192.168.10.76    192.168.20.38    52998 59780 TCP   19.8m   0
 DP0    28   192.168.10.76    192.168.20.38    53010 59780 TCP   19.4m   0
 DP0    0    192.168.10.76    192.168.20.38    53022 59780 TCP   19.8m   0
 DP0    9    192.168.10.76    192.168.20.38    53017 59780 TCP   19.8m   0
 DP0    17   192.168.10.76    192.168.20.38    53014 59780 TCP   19.7m   0
 DP0    29   192.168.10.76    192.168.20.38    53026 59780 TCP   19.7m   0
 DP0    11   192.168.10.76    192.168.20.38    53009 59780 TCP   19.8m   0
 DP0    5    192.168.10.76    192.168.20.38    53005 59780 TCP   19.7m   0
 DP0    2    192.168.10.76    192.168.20.38    53000 59780 TCP   19.7m   0
 DP0    24   192.168.10.76    192.168.20.38    53019 59780 TCP   19.3m   0
 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
-----------------------------------------------------------
 DP0      31      0       0       0
 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 associated with a TCP connection. The following steps show how to name and create an IP address flow.

1. To display IP addresses associated with the TCP connections, use the `portshow lan-stats --per-flow` command. This step is optional.

```
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)
-------------------------------------------------------------------------------
 DP0    9    192.168.10.76    192.168.20.38    53017 59780 TCP   19.9m   0
 DP0    12   192.168.10.76    192.168.20.38    52997 59780 TCP   19.9m   0
 DP0    27   192.168.10.76    192.168.20.38    53020 59780 TCP   19.9m   0
 DP0    2    192.168.10.76    192.168.20.38    53000 59780 TCP   19.8m   0
 DP0    19   192.168.10.76    192.168.20.38    53015 59780 TCP   19.8m   0

[Output is truncated.]
```

2. To create a flow using IP address `192.168.20.38`, use the `portcfg lan-stats --flow create` command.

```
switch:admin> portcfg lan-stats --flow host create --ipaddr 192.168.10.76
Operation Succeeded.
```

3. To view the flow names that you created on the platform, use the `portshow lan-stats --flow` command.

```
switch:admin> portshow lan-stats --flow

 Name                Flow Info
-------------------------------------------------------------------------------
 APP-flow            (TCP:3225-3226)
 host                (ipaddr:192.168.10.76/32)
 host223             (port:25)
 test                (port:25)
-------------------------------------------------------------------------------
```

   The Flow Info column shows the flow definition associated with the flow name.

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 host
 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)
-------------------------------------------------------------------------------
 DP0    20   192.168.10.76    192.168.20.38    53023 59780 TCP   19.6m   0
 DP0    23   192.168.10.76    192.168.20.38    53012 59780 TCP   19.7m   0
 DP0    21   192.168.20.38    192.168.10.76    59779 52994 TCP   0       0
 DP0    25   192.168.10.76    192.168.20.38    53016 59780 TCP   19.7m   0
 DP0    26   192.168.10.76    192.168.20.38    53025 59780 TCP   19.8m   0
[Output is truncated.]
 DP0    19   192.168.10.76    192.168.20.38    53015 59780 TCP   19.9m   0
-------------------------------------------------------------------------------
 Sport=Source-Port Dport=Destination-Port Pro=Protocol


 DP       ActTCP  ExdTCP  TCLDeny TCLFail
-----------------------------------------------------------
 DP0      31      0       0       0
 DP1      21      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 following steps show how to create and name a TCP port flow.

1. To see the TCP ports that are in use, 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
 DP1      0       0       0       0
-----------------------------------------------------------
 ActTCP=Active TCP Conns  ExdTCP=Exceeded TCP Conn Cnt
```

   The `Sport` and `Dport` columns show the TCP port numbers.

2. To display a list that shows TCP ports associated with known application ports, use the `portshow lan-stats --known-apps` command. This step is optional.

```
switch:admin> portshow lan-stats --known-apps

 App                    Port-Id(s)
-------------------------------------------------------------------------------

 CIFS                   139,445
 FCIP                   3225-3226
 FTP                    20-21,989-990,115
 HTTP                   80,8080,8000-8001,3128
 HTTPS                  443
 iSCSI                  3260
 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
 Data-Domain            2051

-------------------------------------------------------------------------------
```

   You can use the port IDs to create a TCP port flow. You can also use the application information to create an application-based flow.

3. To create a TCP flow using port address `49883`, use the `portcfg lan-stats --flow create` command.

```
switch:admin> portcfg lan-stats --flow TCP-flow create --tcp 49883
Operation Succeeded.
```

4. To view the flow names that you created on the platform, use the `portshow lan-stats --flow` command.

```
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)
--------------------------------------------------------------------------------
```

   The Flow Info column shows the flow definition associated with the flow name.

5. 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 TCP-flow
No Active Per-flow stats to Display
```

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.

The following steps show how to create flows that use logical operators.

1. To combine a specific DP with an IP address and VE_Port using the AND operator, use the `portcfg lan-stats --flow create` command.

```
switch:admin>  portcfg lan-stats --flow dp1_225 create --dp dp1 --and --ipaddr 192.168.10.76/24 --
port 35
Operation Succeeded.
```

   In this example, DP1 is selected along with an IP address on a specific port. Traffic flow statistics that meet all criteria will be displayed for the `dp1_225` flow.

2. To combine a DP with an IP address and VE_Port using the OR operator, use the `portcfg lan-stats --flow create` command.

```
switch:admin> portcfg lan-stats --flow dp0_ve35 create --dp dp0 --or --ipaddr 192.168.10.76/24 --
port 35
Operation Succeeded.
```

   In this example, the statistics for all TCP traffic on DP0 will be displayed, along with all traffic on VE_Port 35 or IP address `192.168.10.76`.

3. To view the flow names that you created on the platform, use the `portshow lan-stats --flow` command.

```
admin> portshow lan-stats --flow

 Name               Flow Info
--------------------------------------------------------------------------------
 dp0_ve35           (dp:DP0 or ipaddr:192.168.10.76/24 or port:35)
 dp1_225            (dp:DP1 and ipaddr:192.168.10.76/24 and port:35)
--------------------------------------------------------------------------------
```

The Flow Info column shows the flow definition associated with the flow name.

Displaying Historical Flow Statistics

You can display historical statistics for LAN traffic in the IP Extension Flow Monitor. The historical display is a snapshot of information for all closed connections that is 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.

You can apply options to the command to filter specific information. For example, you can filter on a DP or you can define additional filters.

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, it is recommended that you first freeze the historical statistics.

1. To display the default view of historical statistics, use the `portshow lan-stats --hist-stats` command.

```
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:
 32   192.168.20.38       192.168.10.76       49374  52996  TCP  231      139
 25   192.168.20.38       192.168.10.76       59767  52973  TCP  1.0g     14.5m
 13   192.168.20.38       192.168.10.76       59767  52961  TCP  752.1m   10.5m
 30   192.168.20.38       192.168.10.76       59767  52978  TCP  755.1m   10.6m
 21   192.168.20.38       192.168.10.76       59767  52969  TCP  754.8m   10.6m
Last 5 Connections:
 12   192.168.20.38       192.168.10.76       59767  52960  TCP  1.0g     14.5m
 3    192.168.20.38       192.168.10.76       59767  52951  TCP  755.3m   10.6m
 1    192.168.20.38       192.168.10.76       59766  52947  TCP  220      168
 11   192.168.20.38       192.168.10.76       59767  52959  TCP  755.0m   10.6m
 31   192.168.20.38       192.168.10.76       59767  52979  TCP  1.0g     14.5m
--------------------------------------------------------------------------------

Total Connection count:   33
Oldest Entry:             10/17/2017 09:45:16 PDT
Newest Entry:             10/17/2017 09:46:21 PDT
Close RX/TX FIN:          1 / 2
Close RX/TX RST:          30 / 0
Total TX Errors
 Slow Starts:             3
 FastRetrans/RetransTO:   0 / 3
Total RX Errors
 Out of Orders/Dup Ack:   0 / 0

DP1: No historical stats available.
```

The snapshot display shows the first and last five entries and summary information for each DP.

2. To freeze the historical flow data, use the `portshow lan-stats --hist-stats -freeze` command.

```
switch:admin> portshow lan-stats --hist-stats -freeze
Operation Succeeded.
DP0 Hist Status: Frozen
DP1 Hist Status: Frozen
```

3. To thaw the historical flow data, use the `portshow lan-stats --hist-stats -thaw` command.

```
switch:admin> portshow lan-stats --hist-stats -thaw
Operation Succeeded.
DP0 Hist Status: Thawed
DP1 Hist Status: Thawed
```

4. To display a detailed view, use the `portshow lan-stats --hist-stats -detail` command. This command displays details for the first and last five entries. When 10 or fewer entries are stored, all connection details are displayed. Use the -all option to display details for all entries stored in the DP instead of the first and last five entries.

```
switch:admin>  portshow lan-stats --hist-stats -detail

Warning: It is recommended to freeze the table when using detailed stats.
DP0 Connection Summary: (Thawed)
-------------------------------------------------------------------------------
 Idx  Src-Address         Dst-Address         Sport  Dport  Pro  Tx(B)    Rx(B)
-------------------------------------------------------------------------------
First 5 Connections:
-------------------------------------------------------------------------------
 32   192.168.20.38       192.168.10.76       49374  52996  TCP  231      139
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
 25   192.168.20.38       192.168.10.76       59767  52973  TCP  1.0g     14.5m
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
 13   192.168.20.38       192.168.10.76       59767  52961  TCP  752.1m   10.5m
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
 30   192.168.20.38       192.168.10.76       59767  52978  TCP  755.1m   10.6m
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
 21   192.168.20.38       192.168.10.76       59767  52969  TCP  754.8m   10.6m
-------------------------------------------------------------------------------

Last 5 Connections:
-------------------------------------------------------------------------------
 12   192.168.20.38       192.168.10.76       59767  52960  TCP  1.0g     14.5m
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
 3    192.168.20.38       192.168.10.76       59767  52951  TCP  755.3m   10.6m
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
 1    192.168.20.38       192.168.10.76       59766  52947  TCP  220      168
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
 11   192.168.20.38       192.168.10.76       59767  52959  TCP  755.0m   10.6m
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
 31   192.168.20.38       192.168.10.76       59767  52979  TCP  1.0g     14.5m
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------

Total Connection count:   33
Oldest Entry:             10/17/2017 09:45:16 PDT
Newest Entry:             10/17/2017 09:46:21 PDT
Close RX/TX FIN:          1 / 2
Close RX/TX RST:          30 / 0
Total TX Errors
 Slow Starts:             3
 FastRetrans/RetransTO:   0 / 3
Total RX Errors
 Out of Orders/Dup Ack:   0 / 0

Warning: It is recommended to freeze the table when using detailed stats.
DP1: No historical stats available.

switch:admin>
```

Notice that the information is thawed and a message advises that you freeze the data before displaying detailed statistics.

5. To display all options that you can use to filter the historical statistics, use the `portshow lan-stats --hist-stats -filter -help` command. This command returns a list of all options available for the `portshow` command.

```
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 historicatl 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 Detail

IP pair statistics are useful for understanding traffic flow that occurs across all the TCP connections that are created between a source and destination IP address. See Using IP Extension Flow Monitor on page 194 for information on the default view of IP pair statistics. IP pair detail can provide information about the TCP connections on a specific DP and index value. The following steps show how to display IP pair details.

1. To display detailed information for IP pairs, use the `portshow lan-stats --ip-pair -detail` command.

```
switch:admin> portshow lan-stats --ip-pair -detail

  portshow lan-stats --ip-pair -hist -detail

 DP0: Index:0 192.168.20.38 <-> 192.168.10.76
-------------------------------------------------------------------------------
 No of active connections:  31
 No of closed connections:  33
 Creation time              10/17/2017 09:44:35 PDT
                            TxB       RxB
 Current stats:             4.5t      358
 Mon 10/16 stats:           0         0
 Sun 10/15 stats:           0         0
 Sat 10/14 stats:           0         0
 Fri 10/13 stats:           0         0
 Thu 10/12 stats:           0         0
 Wed 10/11 stats:           0         0
 Tue 10/10 stats:           0         0
 Total stats:               4.5t      358

 DP1: Index:0 1000:17:20::38 <-> 1000:17:10::76
-------------------------------------------------------------------------------
 No of active connections:  2
 No of closed connections:  0
 Creation time              10/17/2017 11:53:33 PDT
                            TxB       RxB
 Current stats:             83        57
 Mon 10/16 stats:           0         0
 Sun 10/15 stats:           0         0
 Sat 10/14 stats:           0         0
 Fri 10/13 stats:           0         0
 Thu 10/12 stats:           0         0
 Wed 10/11 stats:           0         0
 Tue 10/10 stats:           0         0
 Total stats:               83        57
```

The output shows detailed information for a Brocade 7840 Switch. Each IP pair has a unique index value, which you can use to filter the results. In this example, only DP0 has an active IP pair connection.

2. To display IP pair information detail for a specific DP and index value, use the `portshow lan-stats --ip-pair [slot/dp] [index] -detail` command.

```
admin>   portshow lan-stats --ip-pair dp0 0  -detail

 DP0: Index:0 192.168.20.38 <-> 192.168.10.76
-------------------------------------------------------------------------------
 No of active connections:  31
 No of closed connections:  33
 Creation time              10/17/2017 09:44:35 PDT
                            TxB       RxB
 Current stats:             4.6t      358
 Mon 10/16 stats:           0         0
 Sun 10/15 stats:           0         0
 Sat 10/14 stats:           0         0
 Fri 10/13 stats:           0         0
 Thu 10/12 stats:           0         0
 Wed 10/11 stats:           0         0
 Tue 10/10 stats:           0         0
 Total stats:               4.6t      358
```

You can see that the displayed output is limited to DP0 and index 0.

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.

It can take several seconds to accumulate the snapshot data, so the 24-hour time-stamp can vary from day to day. For information on how to display historical flow history, see Displaying Historical Flow Statistics on page 202.

The following steps show how to display IP pair history.

1. To display the IP pair history information, use the `portshow lan-stats --ip-pair -hist` command.

```
switch:admin> portshow lan-stats --ip-pair -hist

 DP       Idx SrcAddr             DstAddr             Active  Closed
             Today   Mon     Sun     Sat     Fri     Thu     Wed     Tue
-------------------------------------------------------------------------------
 DP0      0   192.168.20.38       192.168.10.76       31      33
   TxB       4.8t    -       -       -       -       -       -       -
   RxB       358     -       -       -       -       -       -       -

 DP1      0   1000:17:20::38      1000:17:10::76      21      12
   TxB       432     -       -       -       -       -       -       -
   RxB       381.8g  -       -       -       -       -       -       -

-------------------------------------------------------------------------------
```

The values in the `Today` column are a snapshot as of when the command was entered.

2. To 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 display.

```
switch:admin> portshow lan-stats --ip-pair dp0 0  -hist -detail

 DP0: Index:0 192.168.20.38 <-> 192.168.10.76
--------------------------------------------------------------------------------
 No of active connections:  31
 No of closed connections:  33
 Creation time              10/17/2017 09:44:35 PDT
                            TxB       RxB
 Current stats:             4.2t      358
 Mon 10/16 stats:           0         0
 Sun 10/15 stats:           0         0
 Sat 10/14 stats:           0         0
 Fri 10/13 stats:           0         0
 Thu 10/12 stats:           0         0
 Wed 10/11 stats:           0         0
 Tue 10/10 stats:           0         0
 Total stats:               4.2t      358
```

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. If there is a saved IP pair that has no active connection, 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
- Firmare downgrade
- Platform reboot

NOTE
There is no warning displayed or confirmation required when you reset IP pair statistics.

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 steps show how to reset IP pair statistics:

1. To reset all IP pair statistics, use the `portshow lan-stats --ip-pair reset` command.

```
admin> portshow lan-stats --ip-pair -reset

IP-Pair reset status:Success
```

2. To reset only certain IP pair statistics, 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. 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.

   a) Display 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      0   192.168.20.38       192.168.10.76       31      1
   TxB       22.2t   51.0t   51.0t   51.0t   51.0t   51.0t   51.0t   51.0t
   RxB       0       -       -       -       -       -       -       -

 DP1      0   1000:17:20::38      1000:17:10::76      21      1
   TxB       0       -       -       -       -       -       -       -
   RxB       37.8t   86.7t   86.9t   86.7t   86.7t   86.8t   86.9t   86.7t

-------------------------------------------------------------------------------
```

   b) Reset the IP pair statistics for `DP0` and index `0`.

```
switch:admin> portshow lan-stats --ip-pair dp0 0 -reset

IP-Pair reset status:Success
```

   c) Confirm that the selected IP pair statistics are reset.

```
switch:admin> portshow lan-stats --ip-pair -hist

 DP       Idx SrcAddr             DstAddr             Active  Closed
               Today   Tue     Mon     Sun     Sat     Fri     Thu     Wed
-------------------------------------------------------------------------------
 DP0      0   192.168.20.38       192.168.10.76       0       11265
   TxB       2.6m    -       -       -       -       -       -       -
   RxB       822.3k  -       -       -       -       -       -       -

 -------------------------------------------------------------------------------
```

   Any IP pair that has no active connection is removed from the display table when the statistics are reset.

Troubleshooting Tools

In-band Management

NOTE
In-band management is supported on the Brocade FX8-24 Blade only.

In-band management allows management of an extension switch or blade in conjunction with FCIP traffic through Ethernet ports. This enables a management station located on the WAN side of the Brocade FX8-24 Blade platform to communicate with the control processor (CP) for management tasks, such as SNMP polling, SNMP traps, troubleshooting, and configuration. Through IP forwarding, inband management also allows a management station connected to the management port of one extension switch or blade to manage the blade at the far end of the network through the WAN.

The in-band management path is achieved by receiving the management traffic from the Ethernet port and transmitting the traffic to the CP through the inband interface. The CP then handles the management traffic as it would handle any other management requests from a normal management interface. The in-band management interface is protocol-independent, so any traffic destined for these in-band management interfaces passes through the data processor (DP) to the CP. It is then handled on the CP according to the rules set forth for the normal management interface and follows any security rules that may be in place on the CP.

One in-band management interface can be configured per Ethernet interface to provide redundancy. This allows the management station on the WAN side of the network to have multiple addresses for reaching that switch and provides redundancy if one of the Ethernet ports cannot be reached. Communication is handled through external addresses configured independently for each in-band management interface.

The following functions are not supported by the in-band management interface:

- Downloading firmware
- IPv6 addressing

IP Routing

The in-band management interfaces are separate from the existing IP interfaces currently used for extension tunnel traffic. These interfaces exist on the CP and are added and maintained on the CP routing table to ensure end-to-end connectivity. Because this routing table will be shared among all devices on the CP, including the management interface, precautions must be taken to ensure that proper connectivity is maintained. To ensure proper handling of routes, the in-band management devices should be configured on a different network from the management interface and from every other in-band management interface.

In-band management interface addresses must also be unique and cannot be duplicates of any addresses defined on the Ethernet ports. An in-band management address can exist on the same network as an address defined on one of the GbE ports because the in-band management interfaces use the CP routing table and not the routing table normally used for the GbE ports.

Configuring IP Addresses and Routes

Configure and view IP addresses and routes for in-band management interfaces by using the following Fabric OS commands:

- `portcfg mgmtif [slot/gePort] [create|delete] ipAddress netmask mtu`
- `portcfg mgmtif [slot/gePort] [enable|disable]`
- `portshow mgmtif [slot/gePort]`
- `portcfg mgmtroute [slot/gePort] destination netmask gateway`

Example of a Management Station on the Same Subnet

The following figure illustrates an example of configuring in-band management with the management station attached to the same subnet as managed switches. Note that only the IP address is required for each extension platform.

FIGURE 25 Management station configured on the same subnet

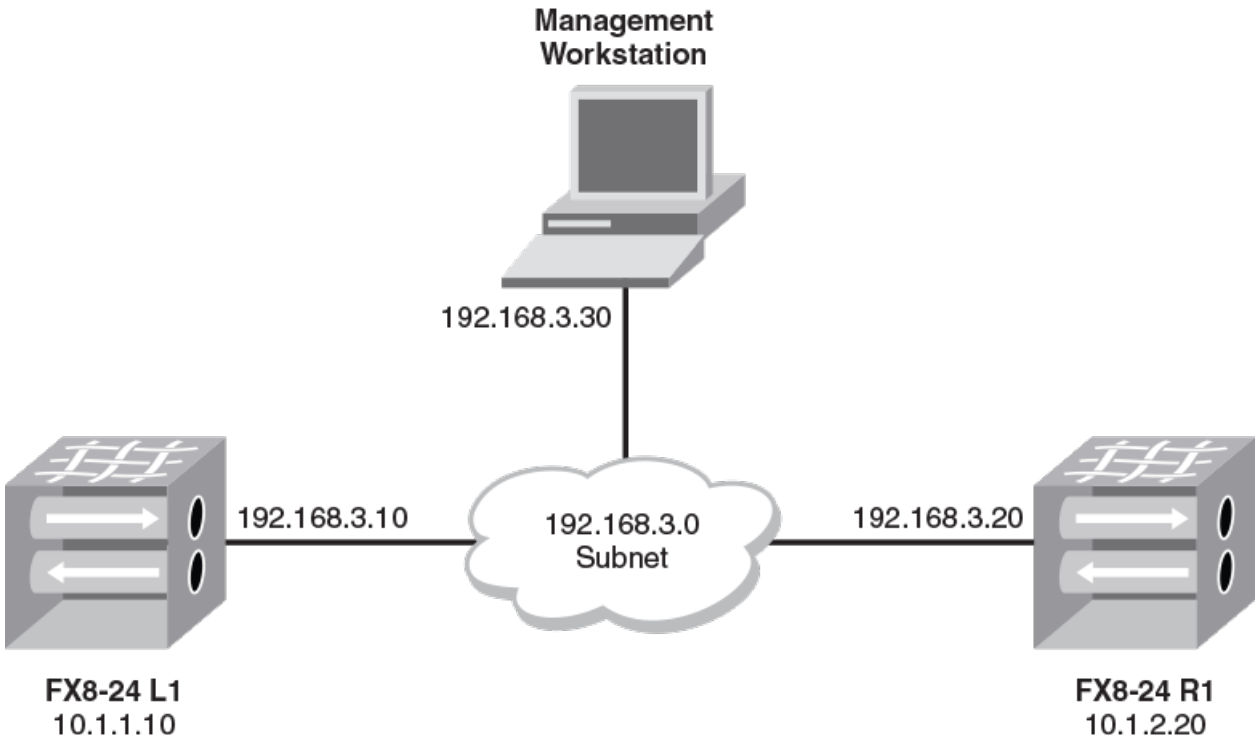

FX8-24 L1

Configure the in-band management interfaces.

```
portcfg mgmtif 1/ge0 create 192.168.3.10 255.255.255.0
```

FX8-24 R1

Configure the in-band management interfaces.

```
portcfg mgmtif 1/ge0 create 192.168.3.20 255.255.255.0
```

Management Workstation

```
telnet 192.168.3.10
```

Example of a Management Station on a Different Subnet

The following figure illustrates an example configuration consisting of switches and the management station on different networks and attached through a WAN cloud. The routers are assumed to already have route entries to access each other's subnet.

FIGURE 26 Management station configured on different subnets

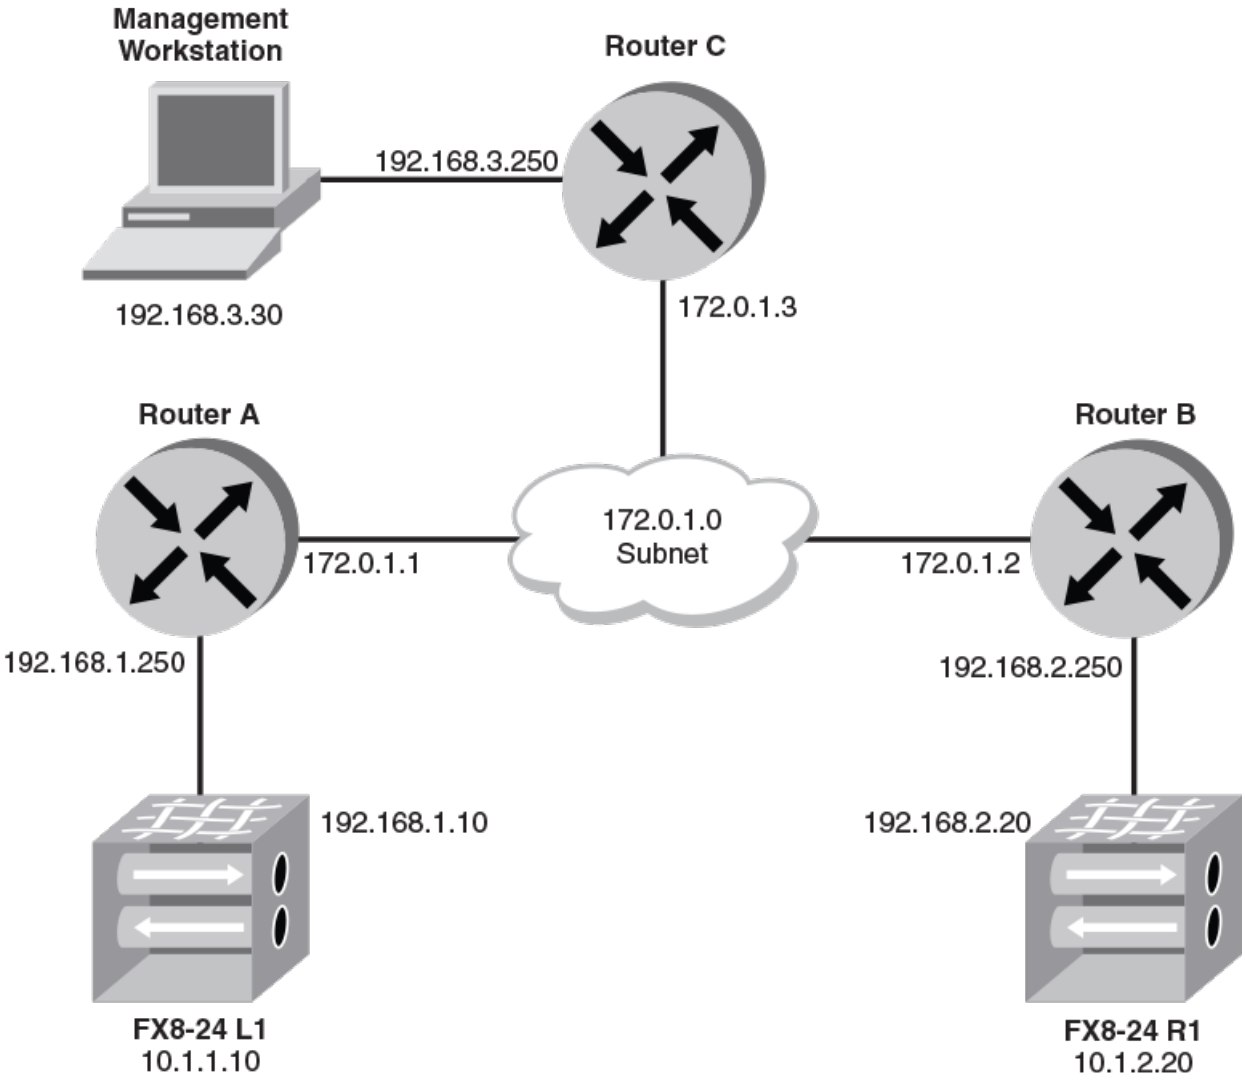

FX8-24 L1

- Configure the in-band management interfaces.

```
portcfg mgmtif 1/ge0 create 192.168.1.10 255.255.255.0
```

- Configure the in-band management route for the management station.

```
portcfg mgmtroute 1/ge0 create 192.168.3.0 255.255.255.0 192.168.1.250
```

FX8-24 R1

- Configure the in-band management interfaces.

```
portcfg mgmtif 1/ge0 create 192.168.2.20 255.255.255.0
```

- Configure the in-band management route for the management station.

```
portcfg mgmtroute 1/ge0 create 192.168.3.0 255.255.255.0 192.168.2.250
```

Management station

- Add route entries to access the Brocade FX8-24 Blade external in-band management interfaces.

```
route add 192.168.1.0 netmask 255.255.255.0 gw 192.168.3.250
route add 192.168.2.0 netmask 255.255.255.0 gw 192.168.3.250
```

- Access the Brocade FX8-24 Blade through the external in-band management interfaces.

```
telnet 192.168.1.10
```

Example of Redundant Connections to the Management Stations

In the following figure, because the in-band management interfaces do not support a multi-homing stack, unique addresses must be used on the management station to communicate with different in-band management interfaces. If both management station interfaces are on the same subnet, then host-specific routes must be added on the Brocade FX8-24 Blade.

FIGURE 27 Redundant connections to management station

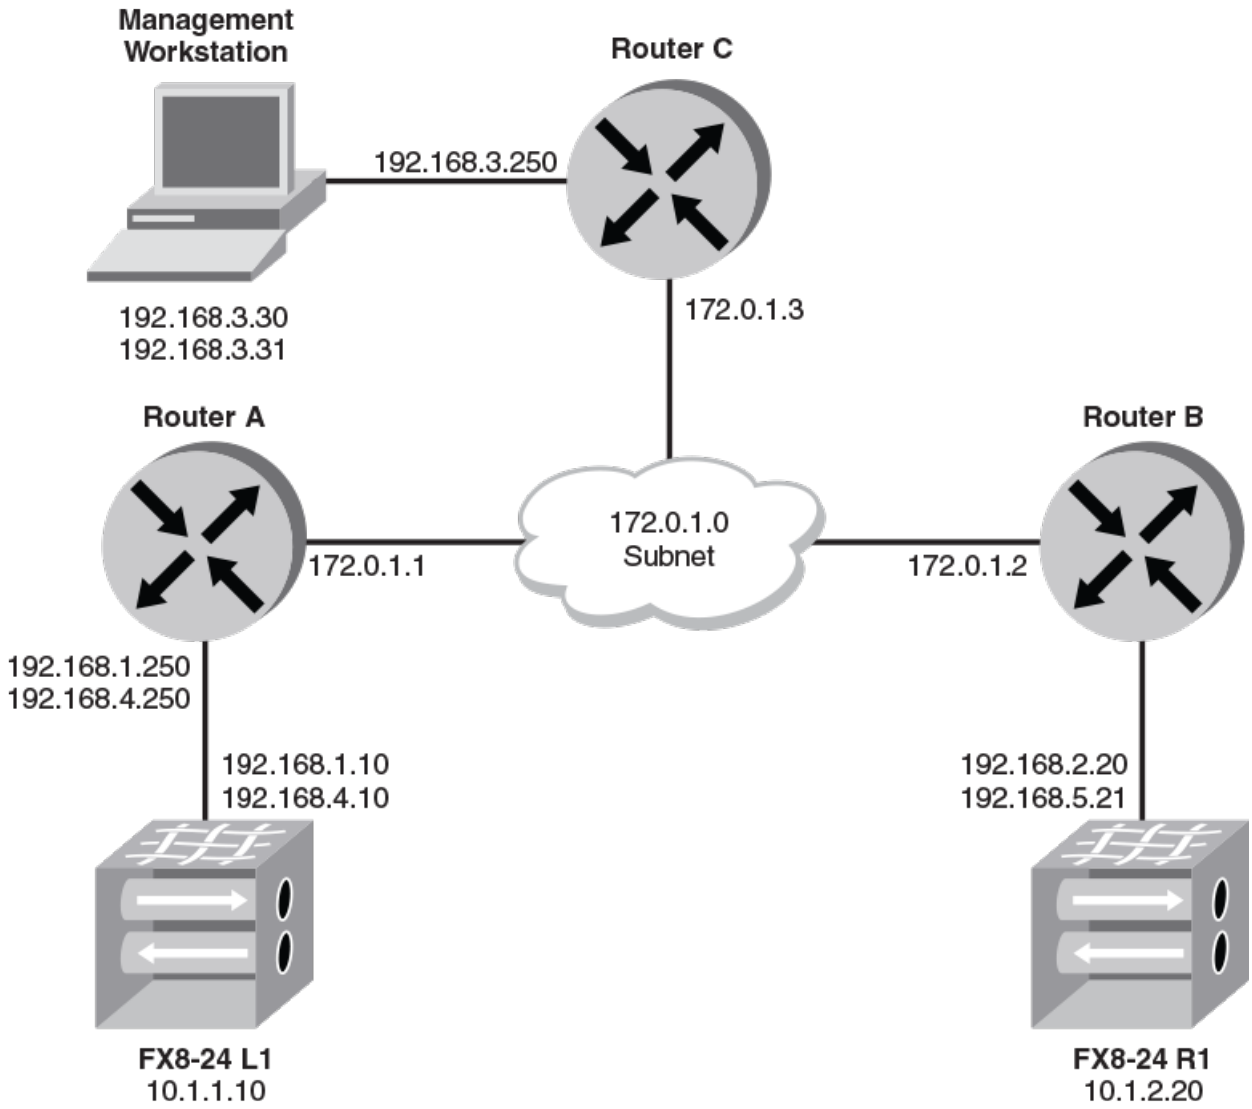

FX8-24 L1

- Configure the in-band management interfaces.

```
portcfg mgmtif 1/ge0 create 192.168.1.10 255.255.255.0
portcfg mgmtif 1/ge1 create 192.168.4.10 255.255.255.0
```

- Configure the in-band management route for the management workstation.

```
portcfg mgmtif 1/ge0 create 192.168.1.10 255.255.255.255 192.168.1.250
portcfg mgmtif 1/ge1 create 192.168.4.10 255.255.255.255 192.168.4.250
```

FX8-24 R1

- Configure the in-band management interfaces.

```
portcfg mgmtif 1/ge0 create 192.168.2.20 255.255.255.0
portcfg mgmtif 1/ge1 create 192.168.5.20 255.255.255.0
```

- Configure the in-band management route for the management workstation.

```
portcfg mgmtroute 1/ge0 create 192.168.3.30 255.255.255.255 192.168.2.250
portcfg mgmtroute 1/ge1 create 192.168.3.31 255.255.255.255 192.168.5.250
```

Management Workstation

- Add route entries to get to the Brocade FX8-24 Blade external in-band management interfaces.

```
route add 192.168.1.0 netmask 255.255.255.0 gw 192.168.3.250
route add 192.168.2.0 netmask 255.255.255.0 gw 192.168.3.250
route add 192.168.4.0 netmask 255.255.255.0 gw 192.168.3.250
route add 192.168.5.0 netmask 255.255.255.0 gw 192.168.3.250
```

- Access the Brocade FX8-24 Blade through the external in-band management interfaces.

```
telnet 192.168.1.10
```

VLAN Tagging Support

To add VLAN tag entries to the VLAN tag table for in-band management interfaces, use the `--mgmt` or `-m` options with the `portcfg vlantag` command. Perform the following steps.

1. Configure an IP address and route for in-band management interface using the following command format.

   `portcfg mgmtif [`*`slot/ge_port`*`] [create|delete]` *`ipAddress netmask mtu`*

2. Add the VLAN tag entry for the management interface using the following command format.

   `portcfg vlantag [`*`slot/ge_port`*`] [add|delete]` *`ipAddress`* `L2COS --mgmt`

IP Forwarding Support

IP forwarding is supported over in-band management to allow communication to the remote switch through the WAN connection. This is done by enabling IP forwarding to allow IP packets arriving at the CP interface to be forwarded through the in-band management interface to the remote side. To prevent network routing and actual bridging of the LAN side of the network to the WAN side of the network, the forwarding rules of the `ipfilter` command will default to deny any forwarding traffic. To allow forwarding, new `ipfilter` command rules must be added to specific destinations. This will prevent any unintended network traffic from being forwarded from the LAN side to the WAN side of the network.

The following figure shows an example network where the management station is located on the LAN side of a FX8-24 Blade L1. Using in-band management, the station can also communicate with a Brocade FX8-24 Blade R1.

FIGURE 28 In-band management with IPv4 forwarding

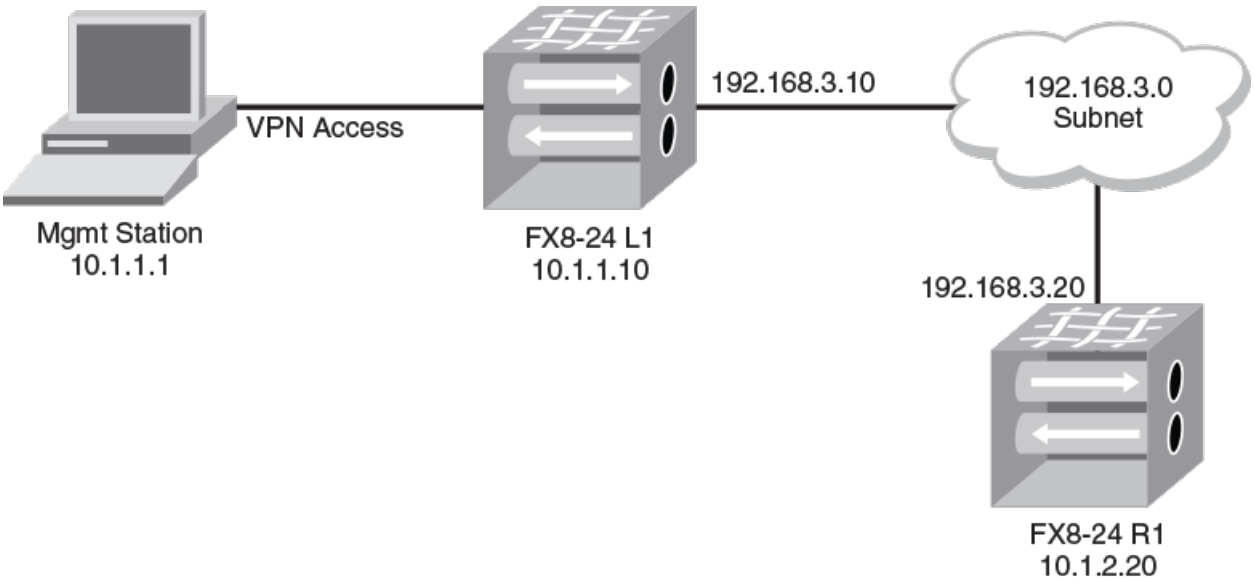

For this example, you must configure the following:

- **Management station**
  - IP address 10.1.1.1/24 (defined)
  - IP route to 192.168.3.20/32 via 10.1.1.10
- **FX8-24 L1**
  - CP management address 10.1.1.10/24
  - In-band management address 192.168.3.10/24
  - IP filter forward rule with destination IP address 192.168.3.20
- **FX8-24 R1**
  - CP management address 10.1.2.20/24
  - In-band management address 192.168.3.20/24
  - In-band management route to 10.1.1.1/32 via 192.168.3.10

Once all of these configurations are complete, proper IP connectivity should occur through the network. In the case where there are routed networks between the Brocade FX8-24 Blades, you will need to add in-band management routes to each Brocade FX8-24 Blade. Using host-specific routes will help eliminate undesired traffic. If network routes are needed, they can be substituted, but you should note that this will allow anything on that network to be forwarded, which could result in undesired disruption of traffic.

NOTE
In all routed network cases, all intermediate hops must have route entries to get to the endpoints.

Using the ipfilter Command

Use the `ipfilter` command to create and manage forwarding rules for use with in-band management. For full details on this command, options, and arguments, refer to the "ipfilter" section of the *Brocade Fabric OS Command Reference*.

To create an IP forwarding rule, you must first create a new policy if one has not yet been created. The easiest way to do this is with the –clone option to create a copy of the default policy.

```
ipfilter --clone inband_ipv4 -from default_ipv4
```

A new rule can be added to allow forwarding traffic.

```
ipfilter --addrule inband_ipv4 -rule rule_number
 -dp dest_port
 -proto protocol
 -act [permit|deny] -type FWD -dip destination_IP
```

Valid *dest_port* values are any TCP or UDP port numbers or a range of port numbers that you want forwarded. Valid *protocol* values are tcp or udp . The *destination_IP* is the IP address of the in-band management interface on the remote side. After a rule is added, save the policy and activate it using the --save and --activate options of the ipfilter command. There can only be a single IPv4 policy active at any time. Each policy can consist of multiple rules.

WAN Analysis Tools

WAN analysis tools 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 on the portCmd command. The following options are available:

- portCmd --tperf. Generates and sends test data over a tunnel to determine the characteristics and reliability of the IP network used by the tunnel at the circuit level. Supported on the Brocade FX8-24 Blade.
- portCmd --ping. Tests connections between a local Ethernet port and a destination IP address.
- portCmd --pmtu. Assists identifying the maximum MTU possible in the WAN network. Supported on the Brocade 7840 Switch and Brocade SX6 Blade.
- portCmd --traceroute. Traces routes from a local Ethernet port to a destination IP address.
- portCmd --wtool. Generates traffic over a pair of IP addresses to test the link for issues such as maximum throughput, congestion, loss percentage, out-of-order delivery, and other network conditions. Supported on the Brocade 7840 Switch and the Brocade SX6 Blade.
- portShow fcipTunnel --perf. Displays performance statistics generated from the WAN analysis.

The tperf Option

The tperf option (portCmd --tperf) is a utility that generates data between a local and remote switch over a tunnel. It reports the data generated and response from the remote switch to determine characteristics and reliability of the IP network used by the tunnel.

The tperf option operates with a pair of Brocade FX8-24 Blades. One blade plays the role of a data sink and the other switch or blade plays the role of the data source. During the data generation process, traffic flows from the source to the sink, then the sink responds to this traffic. The process continues for a duration that you specify with command options or until you terminate (**Ctrl** +**C**).

Normally, you should establish one Telnet or SSH session for the tperf source and one for the tperf sink. Also, open additional Telnet or SSH sessions so that you can periodically display TCP connection statistics using the –tcp or –p options of the portshow fciptunnel [*slot*/*ve_port*] command. These statistics can sometimes help you understand the tunnel bandwidth and IP network infrastructure capability.

To use the `tperf` option, you must first create a tunnel with at least one circuit or modify an existing tunnel using the tperf flag `-T`. As with any tunnel, this must be done on both blades. The following commands create a tperf-enabled tunnel with a committed rate of 10000.

```
portcfg fciptunnel 1/16 create --remote-ip 192.168.10.1 --local-ip 192.168.10.2 10000 -T
portcfg fciptunnel 1/16 create --remote-ip 192.168.10.2 --local-ip 192.168.10.1 10000 -T
```

The `tperf` option will test single and multiple circuit tunnels. The `tperf` option also tests the different priority connections that are provided by a tunnel. When a tperf-enabled tunnel is operative, it is not an active VE_Port. Fabrics will not merge over an operative tperf tunnel. To determine if the tperf tunnel is up, issue the following command:

```
switch:admin>  portshow fciptunnel -c

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
--------------------------------------------------------------------------------
 1/16   -         Up       -------  1h21m43s   0.00    0.00    0      -      -/-
 1/16   0 1/ge3   Up       ---4--s  1h21m43s   0.00    0.00    0  1000/1000  0/-
 1/16   1 1/ge3   Up       ---4--s  1h21m43s   0.00    0.00    0  1000/1000  0/-
--------------------------------------------------------------------------------
 Flags (tunnel): c=compression m=moderate compression a=aggressive compression
                 A=Auto compression f=fastwrite t=Tapepipelining F=FICON
                 T=TPerf i=IPSec l=IPSec Legacy
      (circuit): s=sack v=VLAN Tagged x=crossport 4=IPv4 6=IPv6
                 L=Listener I=Initiator
```

The previous display shows VE_Port 16 as up, but a `switchshow` command for that same VE _Port will show the following:

```
switch:admin> switchshow | grep 16 | grep VE
 128    1   16   028000   --    --         Offline     VE
```

For full details on syntax and using this command, refer to the *Brocade Fabric OS Command Reference*.

The following examples create a tperf data sink and a tperf data source on VE_Port 16.

```
switch:admin> portcmd --tperf 1/16 -sink -interval 15
switch:admin> portcmd --tperf 1/16 -source -interval 15 -high -medium -low
```

The tperf option generates statistics every 30 seconds by default unless you specify a different value for`-interval`.

TABLE 48 tperf Output

Item	Description
Tunnel ID	Numeric identifier for the Tperf tunnel.
Traffic	Priority High, Medium, or Low.
bytes tx	Number of bytes transmitted.
bytes rx	Number of bytes received.
PDUs tx	Number of protocol data units transmitted.
PDUs rx	Number of protocol data units received.
bad CRC headers rx	Number of bad CRC headers received.
bad CRC payloads rx	Number of bad CRC payloads received.
out of seq PDUs rx	Number of out-of-sequence PDUs received.
flow control count	Flow control count.
packet loss (%)	The percentage of packet loss.
bytes/second	The number of bytes transmitted per second.
last rtt	The time it took for the last round-trip between the Tperf source and the Tperf sink in milliseconds. This is calculated only on the source-side report. It is reported as N/A on the sink-side report.

Using ping to Test a Connection

The `portCmd --ping` command tests the connection between the IP address of a local Ethernet port and a destination IP address. If you want to use this command to test a VLAN connection when you do not have an active tunnel, you must manually add entries to the VLAN tag table on both the local and remote sides of the route, using the `portCfg vlantag` command (applicable to the Brocade FX8-24 Blade only).

The general syntax of the `portCmd --ping` command is as follows:

`portCmd --ping [slot/ge_port] -s source_ip -d destination_ip -n num_request -q diffserv -t ttl -w wait_time -z size -v vlan_id -c L2_Cos`

On the Brocade 7840 Switch, the Brocade 7810 Switch and the Brocade SX6 Blade, because DP complexes share Ethernet ports, identification for the port is*gen.DP n*, for example `ge0.DP0`. This directs the command to a specific DP complex.

Using Traceroute

The `portCmd traceroute` command traces routes from a local Ethernet port to a destination IP address. If you want to use this command to trace a route across a VLAN when you do not have an active tunnel, you must manually add entries to the VLAN tag table on both the local and remote sides of the route using the `portCfg vlantag` command (applicable to the Brocade FX8-24 Blade only).

The general syntax of the `portCmd --traceroute` command is as follows:

`portCmd --traceroute [slot/ge_port] -s source-ip -d destination_ip -h max_hops -f first_ttl -q diffserv -w wait-time -z size -v vlan_id -c L2_Cos`

On the Brocade 7840 Switch, the Brocade 7810 Switch, and the Brocade SX6 Blades, because DP complexes share Ethernet ports, identification for the port is*gen.DP n*, for example`ge0.DP0`. This directs the command to a specific DP complex.

The following example traces the route between IP addresses 192.168.2.22 and 192.168.2.30 over VLAN 12 from a 7840 switch.

```
portcmd --traceroute ge2.dp1 -s 192.168.2.22 -d 192.168.2.30 -v 12
```

The following example traces the route between IP addresses 192.168.10.1 and 192.168.20.1 over VLAN 10 from an FX8-24 blade.

```
portcmd --traceroute 8/ge0 -s 192.168.10.1 -d 192.168.20.1 -v 10
```

NOTE

To trace a route with crossport addresses, see Using traceroute with Crossports on page 180.

For details of command syntax and output examples, refer to the *Brocade Fabric OS Command Reference*.

Using WAN Tool

WAN Tool allows you to generate traffic at a specified rate in Kbps over a pair of IP addresses to test the network link for issues such as maximum throughput, congestion, loss percentage, out of order delivery, and other network conditions. The main purpose of this tool is to determine the health of a link before deploying it for use as a circuit in a tunnel.

WAN Tool is supported exclusively on the Brocade 7840 Switch, the Brocade 7840 Switch, and the Brocade SX6 Blade.

NOTE

For the Brocade 7810 Switch, we have restricted three of the Automated and User WAN Tool sessions: Maximum Automated WAN Tool sessions (AWT), Maximum User WAN Tool sessions, and Maximum Aggregated Bandwidth.

The following table shows the supported features on the three platforms.

Supported Features	Brocade 7840 Switch	Brocade 7810 Switch	Brocade SX6 Blade
Max Automated WAN Tool sessions	20 per DP0 and 20 per DP1	12	20 per DP0 and 20 per DP1
Max User WAN Tool sessions	10	6	10
Max Aggregated BW	20 Gb/s (10VE mode) 10Gb/s (20VE mode)	2.5 Gb/s	20 Gb/s (10VE mode) 10Gb/s (20VE mode)
IPSec	Yes	Yes	Yes
L2cos	Yes	Yes	Yes
Bi-directional traffic	Yes	Yes	Yes
DSCP	Yes	Yes	Yes
L2cos	Yes	Yes	Yes
Connection-type	Yes	Yes	Yes
PMTUd	Yes	Yes	Yes
Jumbo Frames	Yes	Yes	Yes

Following are requirements and considerations for using WAN Tool:

- Each DP complex supports a maximum of ten WAN Tool sessions. This reflects the number of enabled WAN Tool and SLA-configured WAN Tool sessions. One slot can support a maximum of 20 sessions.
- A test session can run over an IP path being used by an existing circuit between two switches; however, you must disable the circuit at each end before configuring the session (applicable to User WAN Tool only).
- You must configure the WAN Tool session on the switch at each end of the circuit.
- Before you can start a test, you must enable the WAN Tool session on the switch at each end of the circuit. When the testing is complete, you can disable the WAN Tool session, which allows you to retain the WAN Tool configuration when you enable the circuit.
- After configuration, you can start a test from one switch to solely examine unidirectional traffic to the opposite switch or you can test 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 configured for all sessions must be equal to or less then the physical speed of the port (40Gb/s, 10Gb/s, or 1Gb/s). For example, on a 10Gb/s port, you could configure four 2.5Gb/s sessions. Alternatively, on a 40Gb/s interface, you could configure four 10Gb/s sessions.
- The MTU size 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` *`ge#.p#t`* `modify` *`ip_addr`* `mtu` *`new-mtu`* command. For details on this command, see *Brocade Fabric OS Command Reference* or Configuring IPIF on page 102.

A tunnel and WAN Tool cannot operate at the same time since they both utilize the 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 insure that the traffic being generated between the IP pair is as balanced as possible. The configured rate is split equally among 500Mb/s connections. For example, if you configure a 10Gb/s rate for the test session, twenty 500Mb/s connections are created. As another example, if you configure a 1Gb/s rate, two 500 Mb/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. See the example output of this command in Configuring a WAN Tool Session and Displaying Results on page 223.

AWT is triggered in the following ways:

- While creating a circuit, if you specify an SLA as one of the configuration options, the AWT will run before bringing the newly-created circuit online.
- After a circuit has come online, an AWT will run provided the circuit bounces for any reason.
- You can manually launch an AWT session by disabling then re-enabling a circuit, disabling/enabling a switch, or bouncing the GE port a specific circuit is configured on.

During a 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 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

Configure a WAN Tool session using the `portcmd --wtool` command. 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 doesn't have to match on each switch, but this is recommended for easier comparison of test results on both ends of the circuit when multiple test sessions are created. Valid IDs are 0 through 15. You can configure 10 sessions for each DP and a maximum of 16 sessions per slot.
- Administrative status (`create --admin-status [enable|disable]`: This 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 will 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 will be the destination address and the destination address will be the source address on the opposite switch.
- Bi-directional or uni-directional (`[`*`--bi-directional|-uni-directional`*`]`: This is an optional parameter, but if used, configure 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 configured on the switch where the test started. If bi-directional is specified, the session runs for the time configured on the switch where the test started, then it runs for the time (if configured) configured on the opposite switch.

  NOTE
  You can create a WAN Tool session without all required parameters. However, all required parameters must be configured before a test session can be enabled. For more details on WAN Tool command and parameters, refer to the *Brocade Fabric OS Command Reference*.

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`.

Following are examples of using the `modify` parameter:

- To modify the rate, use`portcmd --wtool wt-id modify --rate link_rate`.
- To clear test results, use `portcmd --wtool wt-id modify --clear`.

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
User WAN Tool session historical stats are collected when the `start` command is issued. There are no statistics to display until a session runs to completion at least one time. Statistics are also stored when a session is administratively disabled. For automated WAN Tool sessions, historical statistics are stored automatically when the session runs to completion and 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 shoq`, 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`.

NOTE
You can assign a WAN Tool session ID value between 0 and 15, but only 10 sessions can be configured per DP and a maximum of 16 sessions can be configured per slot.

For more details on WAN Tool command and parameters, refer to the *Brocade Fabric OS Command Reference*.

Configuring a WAN Tool Session and Displaying Results

Use 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 in using an account assigned to the admin role.
2. 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 disables circuit 1.

   ```
   Switch1:admin>portCfg fcipcircuit 24 modify 1 --admin-status disable
   ```

3. Verify that the circuit is disabled using the `portshow fciptunnel -c` command. The OpStatus for circuit 1 should be "Down."
4. 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
   ```

5. 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) does not need to match configuration on Switch1, but this is recommended for 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. Note that the source address of Switch1 becomes the destination address for Switch2 and the destination address becomes the source address. See WAN Tool Commands on page 221 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 will be negotiated. The detailed output of the`portcmd --wtool` *`wt-id`* `show -d` command displays the configured, current, and peer rates.

```
Switch1:admin> portcmd --wtool 2 show -d

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%
```

6. 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 -c

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 created between TCP ports 3225 and 3226 to balance the test 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.

7. You can 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% / -
```

8. When using SLA for automated WAN tool sessions, you can show 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%
```

9. If you have created multiple WAN Tool sessions, you can 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.

10. Start traffic on the test connection by entering the `portcmd --wtool` *`wt-id`* `start` command.

```
Switch1:admin> portcmd --wtool 0 start
```

11. 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%)
```

Note that the "State" shows that the test is running and other statistics display as well, such as test "Run Time" and "Time Remaining".

12. 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
```

13. 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
```

14. You can delete the WAN Tool session on both switches, or you can disable the WAN Tool session on both switches. Perform the following steps to delete the WAN Tool session on both switches:
    a) 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 normal traffic.
    b) 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.
15. 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 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 configured the session on both switches in the circuit with appropriate parameters and values. Refer to the list of parameters required for each switch in WAN Tool Commands on page 221.

Common problems in establishing a connection can result from the following WAN Tool configuration problems:

- The test rate doesn't 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. Note that this will generate a warning that the bandwidth has been exceeded and blocks you from creating a session.
- The IPsec policy doesn't match on each switch. Note that the IPsec names need not match, but the names must refer to the same policy.
- Configured source and destination IP addresses are not correct on one or both switches.

Using the portshow Command

Use the `portshow` command to display port operational information on Brocade extension switches and blades. The *Brocade Fabric OS Command Reference* provides complete descriptions of the `portshow` command syntax and options. The following sections identify a few specific outputs that may be useful for maintenance and troubleshooting.

Displaying IP Interfaces

The following example displays IP interface information for a Brocade 7840 Switch.

```
switch:admin> portshow ipif ge0.dp0
```

The following example displays IP interface information for a Brocade FX8-24 Blade.

```
switch:admin> portshow ipif 1/xge0
```

Displaying IP Routes

The following example displays IP route information for a Brocade 7840 Switch.

```
switch:admin> portshow iproute ge5
```

The following example displays IP route information for a Brocade SX6 Blade.

```
switch:admin> portshow iproute 4/ge0
 Port         IP Address               / Pfx  Gateway               Flags
--------------------------------------------------------------------------------
 4/ge0.dp0    192.168.0.0              / 24   192.168.20.1          U G S
 4/ge0.dp0    192.168.20.0             / 24   *                     U C
 4/ge0.dp0    192.168.20.1             / 32   *                     U H L
 4/ge0.dp0    192.168.20.11            / 32   *                     U C
 4/ge0.dp0    192.168.20.12            / 32   *                     U C
 4/ge0.dp0    192.168.20.13            / 32   *                     U C
 4/ge0.dp0    192.168.20.14            / 32   *                     U C
 4/ge0.dp0    192.168.20.15            / 32   *                     U C
 4/ge0.dp0    192.168.20.16            / 32   *                     U C
 4/ge0.dp0    192.168.20.17            / 32   *                     U C
--------------------------------------------------------------------------------
 Flags: U=Usable G=Gateway H=Host C=Created(Interface)
        S=Static L=LinkLayer X=Crossport
```

The following example displays IP route information for a Brocade FX8-24 Blade when the x-port flag is configured.

```
switch:admin> portshow iproute
Port         IP Address               / Pfx  Gateway               Flags
-------------------------------------------------------------------------------
9/xge0       1.250.250.0              / 24   *                     U C
9/xge0       1.250.250.1              / 32   *                     U H L
9/xge0       1.250.250.205            / 32   *                     U C
9/xge0       1.250.250.209            / 32   *                     U C
9/xge0       1.250.250.213            / 32   *                     U C
9/xge0       1.250.250.217            / 32   *                     U C
9/xge0       1.250.250.221            / 32   *                     U C
9/xge0       1.250.250.225            / 32   *                     U C
9/xge0       1.250.250.229            / 32   *                     U C
9/xge0       1.250.250.233            / 32   *                     U C
9/xge0       1.250.250.237            / 32   *                     U C
9/xge0       1.250.251.0              / 24   1.250.250.1           U G S
9/xge0       1.250.250.0              / 24   *                     U C X
9/xge0       1.250.250.1              / 32   *                     U H L X
9/xge0       1.250.250.207            / 32   *                     U C X
9/xge0       1.250.250.211            / 32   *                     U C X
9/xge0       1.250.250.215            / 32   *                     U C X
9/xge0       1.250.250.219            / 32   *                     U C X
9/xge0       1.250.250.223            / 32   *                     U C X
9/xge0       1.250.250.227            / 32   *                     U C X
9/xge0       1.250.250.231            / 32   *                     U C X
9/xge0       1.250.250.235            / 32   *                     U C X
9/xge0       1.250.250.239            / 32   *                     U C X
9/xge0       1.250.251.0              / 24   1.250.250.1           U G S X
9/xge1       1.250.250.0              / 24   *                     U C
9/xge1       1.250.250.1              / 32   *                     U H L
9/xge1       1.250.250.208            / 32   *                     U C
9/xge1       1.250.250.212            / 32   *                     U C
9/xge1       1.250.250.216            / 32   *                     U C
9/xge1       1.250.250.220            / 32   *                     U C
9/xge1       1.250.250.224            / 32   *                     U C
9/xge1       1.250.250.228            / 32   *                     U C
9/xge1       1.250.250.232            / 32   *                     U C
9/xge1       1.250.250.236            / 32   *                     U C
9/xge1       1.250.250.240            / 32   *                     U C
9/xge1       1.250.251.0              / 24   1.250.250.1           U G S
9/xge1       1.250.250.0              / 24   *                     U C X
9/xge1       1.250.250.1              / 32   *                     U H L X
9/xge1       1.250.250.206            / 32   *                     U C X
9/xge1       1.250.250.210            / 32   *                     U C X
9/xge1       1.250.250.214            / 32   *                     U C X
9/xge1       1.250.250.218            / 32   *                     U C X
9/xge1       1.250.250.222            / 32   *                     U C X
9/xge1       1.250.250.226            / 32   *                     U C X
9/xge1       1.250.250.230            / 32   *                     U C X
9/xge1       1.250.250.234            / 32   *                     U C X
9/xge1       1.250.250.238            / 32   *                     U C X
9/xge1       1.250.251.0              / 24   1.250.250.1           U G S X
-------------------------------------------------------------------------------

Flags: U=Usable G=Gateway H=Host C=Created(Interface)

       S=Static L=LinkLayer X=Crossport
```

Displaying Switch Mode Information with the extncfg Command

The Brocade 7810 switch, the Brocade 7840 switch, and Brocade SX6 Blade operate in *Hybrid mode* to support the IP Extension features. (Be aware that the Brocade 7810 switch operates only in Hybrid mode.)

The following example displays the operating mode for the Brocade 7840 Switch or the Brocade SX6 Blade.

```
switch:admin> extncfg --show
slot 4:
```

```
APP Mode is FCIP
VE-Mode: configured for 10VE mode.
slot 8:
APP Mode is FCIP
VE-Mode: configured for 10VE mode.
```

The following example displays the operating mode for the Brocade 7810 Switch.

NOTE

With release 8.2.1 and the introduction of the Brocade 7810 Switch, we introduced a new CLI under the `extncfg` command to change the GE port mode between copper and optical.

```
switch:admin> extncfg –show
APP Mode is HYBRID (FCIP with IPEXT)
VE-Mode: Not Applicable.
GE-Mode: Copper
```

Displaying GbE Port Information with the portcfgge Command

The following example displays GbE port configuration for a Brocade 7840 Switch, a Brocade 7810 Switch, or a Brocade SX6 Blade.

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

Listing the MAC Addresses of LAN and GE Ports

Use the `portcfgge` *`ge#/lan`* `--show -lmac`command to display the MAC addresses of the LAN and GE ports on the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade. See the following sample outputs for more information.

Examples

The following examples show the output with and without supplying the ge_port.

```
switch:admin> portcfgge 12/xge0 --show -lmac
Port       dpid    MAC Address
--------------------------------------------------------------------------------
12/xge0    -       00:05:33:42:87:6a
--------------------------------------------------------------------------------

switch: admin> portcfgge --show -lmac
Port       dpid    MAC Address
--------------------------------------------------------------------------------
3/ge0      -       00:05:33:3a:1d:04
3/ge1      -       00:05:33:3a:1d:05
3/ge2      -       00:05:33:3a:1d:06
3/ge3      -       00:05:33:3a:1d:07
3/ge4      -       00:05:33:3a:1d:08
3/ge5      -       00:05:33:3a:1d:09
3/ge6      -       00:05:33:3a:1d:0a
3/ge7      -       00:05:33:3a:1d:0b
3/ge8      -       00:05:33:3a:1d:0c
3/ge9      -       00:05:33:3a:1d:0d
3/xge0     -       00:05:33:3a:1d:0e
3/xge1     -       00:05:33:3a:1d:0f
12/ge0     -       00:05:33:42:87:60
12/ge1     -       00:05:33:42:87:61
12/ge2     -       00:05:33:42:87:62
12/ge3     -       00:05:33:42:87:63
12/ge4     -       00:05:33:42:87:64
12/ge5     -       00:05:33:42:87:65
12/ge6     -       00:05:33:42:87:66
12/ge7     -       00:05:33:42:87:67
12/ge8     -       00:05:33:42:87:68
12/ge9     -       00:05:33:42:87:69
12/xge0    -       00:05:33:42:87:6a
12/xge1    -       00:05:33:42:87:6b
--------------------------------------------------------------------------------
```

The following example lists the MAC addresses of all the LAN and GE ports in the switch:

```
switch:admin> portcfgge --show -lmac
 Port       dpid    MAC Address
-------------------------------------------------------
 4/ge0      dp0     00:27:f8:f0:aa:d0
 4/ge0      dp1     00:27:f8:f0:aa:d0
 4/ge1      dp0     00:27:f8:f0:aa:d1
 4/ge1      dp1     00:27:f8:f0:aa:d1
 4/ge2      dp0     00:27:f8:f0:aa:d2
 4/ge2      dp1     00:27:f8:f0:aa:d2
 4/ge3      dp0     00:27:f8:f0:aa:d3
 4/ge3      dp1     00:27:f8:f0:aa:d3
 4/ge4      dp0     00:27:f8:f0:aa:d4
 4/ge4      dp1     00:27:f8:f0:aa:d4
 4/lan      dp0     00:00:00:00:00:00
 4/lan      dp1     00:00:00:00:00:00
 4/lan1     dp0     00:00:00:00:00:00
 4/lan1     dp1     00:00:00:00:00:00
 4/ha0      dp0     00:00:00:00:00:01
 4/ha0      dp1     00:00:00:00:00:01
-------------------------------------------------------
```

The following example lists the MAC addresses of the LAN ports on slot 8 in the switch:

```
switch:admin>  portcfgge 8/lan --show -lmac

 Port       dpid    MAC Address
-------------------------------------------------------
 8/lan      dp0     00:00:00:00:00:00
```

```
 8/lan      dp1     00:00:00:00:00:00
----------------------------------------------------
```

The following example lists the MAC addresses of the GE4 port on slot 4 in the switch:

```
switch:admin> portcfgge 4/ge4 --show -lmac

 Port       dpid    MAC Address
----------------------------------------------------
 4/ge4      dp0     00:27:f8:f0:aa:d4
 4/ge4      dp1     00:27:f8:f0:aa:d4
----------------------------------------------------
```

Displaying LAG Information

You can display link aggregation group (LAG) information for the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade only when it is operating in Hybrid mode. (Recall that the Brocade 7810 Switch operates only in Hybrid mode.)

> **NOTE**
> Brocade 7840 Switch and Brocade SX6 Blade support a maximum of 8 LAGs (combination of static and dynamic) whereas Brocade 7810 Switch supports a maximum of 2 LAGs.

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
--------------------------------------------------------------------------------
dlag101         Dynamic      Online        1            ge6*
slag101         Static       Online        3            ge15 ,ge16 ,ge17
slag102            Static          Offline          0
```

You can display more detailed information by using the detail option.

```
switch:admin> portchannel --show -detail
 Name :dlag101
 Type :Dynamic
 Key  :555
 Speed  :1G
 Autoneg  :Off
 Admin-state: Enable
 Oper-state : Online
  Admin Key: 0555 - Oper Key 0555
  LACP System ID: 0x8000,00-05-33-65-7b-c2
  PART System ID: 0x0001,00-24-38-9c-00-00
  Portchannel Member count = 1
  Port          Oper state      Sync    Timeout
  ---------------------------------------------
  *ge6          Online          1       Long

 Name :slag101
 Type :Static
 Key  :1
 Speed  :10G
 Autoneg  :Off
 Admin-state: Enable
 Oper-state : Online
  Portchannel Member count = 3
  Port                  Oper state
  ----------------------------
   ge15                 Online
   ge16                 Online
   ge17                 Online

 Name :slag102
 Type :Static
```

```
 Key  :123
 Speed  :10G
 Autoneg  :Off
 Admin-state: Disable
 Oper-state : Offline
  Portchannel Member count = 0
```

Displaying Tunnel HCL Information

The following example displays the tunnel HCL status for a Brocade 7840 Switch, a Brocade 7810 Switch, or a Brocade SX6 Blade.

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

Checking FCIP Tunnel HA Status.

Current Status         : Ready
CP Version             : v8.2.1_bld30
DP0 Status:
   State               : Online - Inactive
   Version             : v8.2.1_bld30
   Current FC HA Stage : IDLE
   Current IP HA Stage : IDLE
   IP SVI Swapped      : NO
   DP COMM Status      : UP
DP1 Status:
   State               : Online - Inactive
   Version             : v8.2.1_bld30
   Current FC HA Stage : IDLE
   Current IP HA Stage : IDLE
   IP SVI Swapped      : NO
   DP COMM Status      : UP

Tunnel 24 (FID:40) FC:HA Ready IP:Disabled - FC traffic will be disrupted.
Tunnel 25 (FID:65) FC:HA Online IP:Disabled - Traffic will not be disrupted.
Tunnel 26 (FID:20) FC:HA Offline IP:HA Offline - FC and IP traffic will be disrupted.
Tunnel 27 (FID:65) FC:HA Ready IP:Disabled - FC traffic will be disrupted.
Tunnel 34 (FID:40) FC:HA Online IP:HA Online - Traffic will not be disrupted.
Tunnel 35 (FID:65) FC:HA Ready IP:Disabled - FC traffic will be disrupted.
Tunnel 36 (FID:65) FC:HA Ready IP:Disabled - FC traffic will be disrupted.
Tunnel 37 (FID:20) FC:HA Online IP:Disabled - Traffic will not be disrupted.
Tunnel 38 (FID:20) FC:HA Ready IP:HA Ready - FC and IP traffic will be disrupted.
```

Displaying TCL Information

You can display traffic control list (TCL) configuration information for the Brocade 7840 Switch, the Brocade 7810 Switch, or the Brocade SX6 Blade only when it is operating in Hybrid mode. (Recall that the Brocade 7810 Switch operates only in Hybrid mode.)

> **NOTE**
> Brocade 7840 Switch and Brocade SX6 Blade support a maximum of 1024 defined and 128 active TCL sessions. In contrast, the Brocade 7810 Switch supports a maximum of 256 defined and 32 active TCL sessions.

The following example displays the TCL configuration information. 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
-------------------------------------------------------------------------------
  0     test1                DI---  -         ANY     ANY  ANY  ANY   ANY  0
                             ANY                      ANY

.....................[TUNCATED OUTPUT].....................................
```

```
*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.


Active TCL Limits:   Cur / Max
---------------------------------------------
  4/DP0               1 / 128
  4/DP1               2 / 128
  8/DP0               - / -
  8/DP1               - / -
---------------------------------------------
Configured Total:     5 / 1024
```

You can display more detailed information by using the `detail` option.

```
switch:admin> portshow tcl --detail
```

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

You can display IP Extension LAN statistics for a Brocade 7840 Switch, a Brocade 7810 Switch, or a Brocade SX6 Blade only when it is operating in Hybrid mode. (Recall that the Brocade 7810 Switch operates only in Hybrid mode.)

> **NOTE**
> The Brocade 7840 Switch and Brocade SX6 Blade support a maximum of 8 LAN ports whereas the Brocade 7810 Switch supports a maximum of 4 LAN ports.

The following example displays the global LAN statistics for a Brocade 7840 Switch.

```
switch:admin> portshow lan-stats --global
```

For additional information about displaying IP Extension LAN statistics, see Using IP Extension Flow Monitor on page 194.

Displaying Performance Statistics

Display a summary of performance statistics for all tunnels and circuits using the circuit, perf, and summary options as in the following example.

```
switch:admin> portshow fciptunnel --all --circuit --perf --summary
```

Display a summary of performance statistics for current FV tunnels and circuits using the circuit, perf, and summary options as in the following example.

```
switch:admin> portshow fciptunnel all --circuit --perf --summary
```

Displaying QoS Statistics

Display QoS statistics for all tunnels using the `--qos` and `--summary` options, as in the following example.

```
switch:admin> portshow fciptunnel --all --qos --summary
```

Display QoS statistics for current FV tunnels using the `--qos` and `--summary` options, as in the following example.

```
switch:admin> portshow fciptunnel all --qos --summary
```

Displaying Configuration Details

You can display all configuration details using the `--all --detail` option as in the following example.

```
switch:admin> portshow fciptunnel --all --detail
```

You can display current FV configuration details using the `all --detail` option as in the following example.

```
switch:admin> portshow fciptunnel all --detail
```

Filtering portshow Display Output

You can filter the display output of the `portshow filter` command 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.

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 \
              -tcp-port 336 -and
```

The following example creates a filter set called `tcpErrors` that looks for any retransmits greater than 100 and a byte count greater than 1000000 bytes. This 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 additional details on using the `portshow filter` command, refer to the *Brocade Fabric OS Command Reference*.

Displaying Tunnel Status

The following example of the `portshow fciptunnel` command is used most often to determine tunnel status.

```
switch:admin> portshow fciptunnel --all -c

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
-------------------------------------------------------------------------------
 25     -        Up       --------I   14d15h    0.00    0.00   3      -       -
 25     0 ge11   Up       ----ah--4   14d15h    0.00    0.00   3  5000/5000  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

The following example displays general tunnel information related to port 24.

```
switch:admin> portshow fciptunnel 24

 Tunnel: VE-Port:24 (idx:0, DP0)
 ================================================
  Oper State           : Online Warning
  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 / 5d1h52m
 Uptime               : 5d1h52m
 Stats Duration       : 5d1h52m
 Receiver Stats       : 7697068 bytes / 51760 pkts /   16.00 Bps Avg
 Sender Stats         : 6297068 bytes / 51757 pkts /   16.00 Bps Avg
 TCP Bytes In/Out     : 4460323284 / 4458865888
 ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
 RTT (min/avg/max)    : 1 / 1 / 6 ms
 Wan Util             : 0.0%
 TxQ Util             : 0.0%
```

Displaying a Tunnel with Circuit Information

The following example adds circuit information to the `portshow fciptunnel` command output using the `-c` option.

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


 Tunnel: VE-Port:24 (idx:0, DP0)
 ================================================
 Oper State           : Online Warning
 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 / 5d1h51m
 Uptime               : 5d1h51m
 Stats Duration       : 5d1h51m
 Receiver Stats       : 7695768 bytes / 51751 pkts /    8.00 Bps Avg
 Sender Stats         : 6295972 bytes / 51748 pkts /    8.00 Bps Avg
 TCP Bytes In/Out     : 4459623112 / 4458165680
 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 ge2 <-> 192.168.1.2
 HA IP Addr (L/R)     : 192.168.5.12 ge2 <-> 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 / 5d1h51m
 Uptime               : 5d1h51m
 Stats Duration       : 5d1h51m
 Receiver Stats       : 7695768 bytes / 51751 pkts /    8.00 Bps Avg
 Sender Stats         : 6295972 bytes / 51748 pkts /    8.00 Bps Avg
 TCP Bytes In/Out     : 4459625608 / 4458168176
 ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
 RTT (min/avg/max)    : 1 / 1 / 6 ms
 Wan Util             : 0.0%
```

Displaying Tunnel Performance

The following example displays performance statistics for a tunnel associated with port 24.

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

 Tunnel: VE-Port:24 (idx:0, DP0)
 ================================================
  Oper State           : Online Warning
  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 / 5d1h50m
  Uptime               : 5d1h50m
  Stats Duration       : 5d1h50m
  Receiver Stats       : 7694964 bytes / 51746 pkts /   16.00 Bps Avg
  Sender Stats         : 6295372 bytes / 51743 pkts /   16.00 Bps Avg
  TCP Bytes In/Out     : 4459054876 / 4457597464
  ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
  RTT (min/avg/max)    : 1 / 1 / 6 ms
  Wan Util             : 0.0%
  TxQ Util             : 0.0%
```

Displaying Tunnel TCP Statistics

The following example displays TCP connection statistics for a tunnel associated with port 24.

```
switch:admin>portshow fciptunnel 24 -c --tcp
```

You can reset statistics counters to zero to display only new statistics with the `--tcp` option from the time you issue the reset using the following command.

```
switch:admin> portshow fciptunnel 24 -c --tcp --reset
```

You can display the entire lifetime of statistics for the tunnel using the following command. The time basis for the statistics will display in the output.

```
switch:admin> portshow fciptunnel 24 -c --tcp --lifetime
```

Displaying Circuits

The following example displays circuit information for all VFs.

```
switch:admin> portshow fcipcircuit --all
```

Displaying a Single Circuit

The following example displays information for circuit 1 on tunnel 24.

```
switch:admin> portshow fcipcircuit 24 1
```

Displaying TCP Statistics for Circuits

The following example displays TCP statistics for circuits associated with VE_Port 12 of a Brocade FX8-24 Blade.

```
switch:admin> portshow fcipcircuit 3/12 --tcp
```

You can reset statistics counters to zero to display only new statistics with the `--tcp` option from the time you issue the reset using the following command.

```
switch:admin> portshow fcipcircuit 3/12 --tcp --reset
```

You can display the entire lifetime of statistics for the circuit using the following command. The time basis for the statistics will display in the output.

```
switch:admin> portshow fcipcircuit 3/12 --tcp --lifetime
```

Displaying Circuit Performance

The following example will display circuit performance information for circuit 1 on tunnel 24.

```
switch:admin> portshow fcipcircuit 24 1 --perf
```

Displaying GbE Port Performance

The following example displays GbE throughput port performance. The display updates until you press the **Enter** key. Use the `--help` option to show available command options.

```
switch:adming> geportperfshow
Throughput  of GE port
slot 4:
  ge 0    ge 1    ge 2    ge 3    ge 4    ge 5    ge 6    ge 7    ge 8
================================================================================
    0       0       0      40       0      40       0       0       0

  ge 9    ge10    ge11    ge12    ge13    ge14    ge15    ge16    ge17   Total
================================================================================
   50       0       0       0       0       0       0       0       0     130

slot 8:
  ge 0    ge 1    ge 2    ge 3    ge 4    ge 5    ge 6    ge 7    ge 8
```

```
=========================================================================
    0       0      50       0      51       0       0       0       0

  ge 9    ge10    ge11    ge12    ge13    ge14    ge15    ge16    ge17   Total
=================================================================================
    0       0       0       0       0       0       0       0       0     101
```

Displaying QoS Prioritization for a Circuit

The following example displays QoS prioritization for circuit 0 on tunnel 24. The QoS performance statistics are shown for all QoS tunnels. The example shows a FCIP-only tunnel (no IP traffic) and the Control, High, Medium, and Low statistics.

```
switch:admin> portshow fcipcircuit 24 0 --perf --qos

Circuit 24.0 (DP0)
 ====================================================
  Admin/Oper State     : Enabled / Online
  Priority             : Control
  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: 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
  SLA                  : (none)
  IPSEC                : Disabled (0)
  Failover Group       : 0
  VLAN-ID              : NONE
  L2Cos (Ctrl)         : 0
   DSCP (Ctrl)         : 32
  cfgmask              : 0x40000000 0x00213def
  Flow Status          : 0
  ConCount/Duration    : 1 / 5d12m
  Uptime               : 5d12m
  Stats Duration       : 5d12m
  Receiver Stats       : 7592900 bytes / 51059 pkts /   16.00 Bps Avg
  Sender Stats         : 6211872 bytes / 51056 pkts /   16.00 Bps Avg
  TCP Bytes In/Out     : 1148929804 / 1147502276
  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
  Priority             : FC-High
  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: 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 24.0 (DP0)
==================================================
 Admin/Oper State     : Enabled / Online Warning
 Priority             : FC-Medium
 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: 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 ge2 <-> 192.168.1.2
 HA IP Addr (L/R)     : 192.168.5.12 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%
```

Displaying Tunnel Information (Brocade FX8-24 Blade)

You can use the `portShow fcipTunnel` command to view the performance statistics and monitor the behavior of an online tunnel.

The following example shows using the `portShow fcipTunnel` command with the `-c` option to display the circuits of tunnel 8/12.

```
switch:admin> portshow fciptunnel 8/12 -c
```

Tunnel Issues

The following are common tunnel issues and recommended actions for you to follow to fix them .

Tunnel Does Not Come Online

If a tunnel fails to come online, troubleshoot this issue using the following steps.

1. Confirm that the Ethernet port is online.

```
switch:admin> portshow ge2
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 commands.

```
switch:admin> portshow ipif ge2

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge2.dp0      192.168.5.2                   / 24   1500  0     U R M I
 ge2.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 ge2 <-> 192.168.1.2
  HA IP Addr (L/R)     : 192.168.5.12 ge2 <-> 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 possible problem and the information identified by Configuration Warnings tells where the problem might exist. Both ends of the circuit should be configured with the same parameters.

3. Enter the `portCmd --ping` command to the remote tunnel endpoint from both endpoints.

   The `ge1.dp0` value identifies a port on a Brocade 7840 switch or Brocade SX6 blade. The `-s` value is the source IP address; the `-d` value is the destination IP address.

   ```
   portcmd --ping ge1.dp0 -s 11.1.1.1 -d 11.1.1.2
   ```

   If the command is successful, then you have IP connectivity and your tunnel should come up. If not, continue to the next step.

   When using VLANs, VLAN tagging ensures that test traffic traverses the same path as real traffic. A VLAN tag entry for both the local and remote sides of the route must exist prior to issuing the `portCmd --ping` or `portCmd --traceroute` commands. See Configuring VLANs on page 106 for details.

4. Enter the `portCmd --traceroute` command to the remote tunnel endpoint from both endpoints.

   ```
   portcmd --traceroute ge1.dp0 -s 11.1.1.1 -d 11.1.1.2
   ```

5. If there are routed IP connections that provide for the tunnel, confirm that both ends of the tunnel have defined IP routes, and the route gateways are correct. The tunnel or route lookup may fail to come online because of a missing or incorrect IP route.

   See Configuring IP Route on page 104 to review the set up of the IP route.

6. Confirm that the tunnel is configured correctly using the following command.

   ```
   portshow fciptunnel tunnel ID (VE-port number)
   ```

   Confirm that the compression, FastWrite, and OSTP and IPsec settings match at each endpoint or the tunnel might not come up. Confirm that the local and destination IP address and WWN 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 port 3225, and for the Brocade 7840 Switch, the Brocade 7810 Switch, and the Brocade SX6 Blade, TCP port 3226) and IPsec traffic, if IPsec is used (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.

Tunnel Goes Online and Offline

A tunnel that goes offline and then online (a bouncing tunnel) is a common problem. This bouncing usually occurs because of an overcommitment of available bandwidth, resulting in the following behaviors:

- Too much data tries to go over the link.
- Management data gets lost, or is queued too long, and timeouts expire.
- Data times out multiple times.

Perform the following steps to gather information.

1. Verify what link bandwidth is available.
2. Confirm that the IP path is being used exclusively for 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. For the Brocade FX8-24 Blade, run the `portcmd --tperf` command to gather WAN performance data. For the Brocade 7840 Switch, the Brocade 7810 Switch, and the Brocade SX6 Blade, use WAN Tool.

Troubleshooting Extension Links

The following list 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 only verifies physical connectivity. This command does not verify that you have configured the ports correctly for tunnels.
- Ports at both ends of the tunnel must be configured correctly for a tunnel to work correctly. These ports can be either VE_Ports or VEX_Ports. A VEX_Port must be connected to a VE_Port.
- When configuring routing over an Extension link for a fabric, the edge fabric will use VE_Ports and the backbone fabric will use VEX_Ports for a single tunnel.
- If a tunnel fails and a "Disabled (Fabric ID Oversubscribed)" message displays, the solution is to reconfigure the VEX_Port to the same fabric ID as all of the other ports connecting to the edge fabric.
- Because of an IPsec RASlog limitation, you might not be able to determine an incorrect configuration that causes an IPsec tunnel to not become active. 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 Additional Information

The following commands should be executed and their data collected before the `supportsave` command is run. Using the `supportsave` command can take ten minutes or more to run, and some of the information is time critical.

- `tracedump -n`
- `porttrace --show all`
- `porttrace --status`

For issues specific to tunnel ports, run and collect the data from the following commands:

- `slotshow`
- `portshow` *`slot/ge_port`* /

For a Brocade 7840 or Brocade 7810 Switch, run and collect the data from the following commands:

- `extncfg --show`
- `portcfgge --show`

For a Brocade 7810 or a Brocade 7840 switch that is running in Hybrid mode (for IP Extension features), run and collect the data from the following commands:

- `portshowlag`
- `portshowlag --detail`
- `portshowtcl`
- `portshowtcl --detail`

- `portshowlan-stats --global`
- `portshowlan-stats --per-flow`
- `portshowlan-stats --hist-stats`

If possible, run and collect the data from the following commands:

- `portshow ipif --all` *`slot/ge_port`*
- `portshow arp --all` *`slot/ge_port`*
- `portshow iproute --all` *`slot/ge_port`*
- `portshow fciptunnel --all` *`slot/ge_portall|tunnel`*
- `portshow fciptunnel --all --perf`
- `portshow fciptunnel --all -c`
- `portshow fciptunnel --all --circuit --perf --summary`
- `portshow fciptunnel --all --circuit --perf --tcp --qos`
- `portCmd --ping --traceroute --perf`
- `portCmd -ping`
- `portCmd traceroute`

Finally, gather the data from the `supportsave` command.

Refer to the *Brocade Fabric OS Administration Guide* or *Brocade Fabric OS Command Reference* for complete details on these commands.

Using FTRACE

FTRACE is a support tool used primarily by your switch support provider. FTRACE can be used in a manner similar to that of a channel protocol analyzer. You can use FTRACE to troubleshoot problems through a Telnet session rather than using an analyzer or sending technical support personnel to the installation site.

CAUTION
FTRACE is meant to be used solely as a support tool and should be used only by Brocade support personnel or at the request of Brocade support personnel. The FTRACE command is restricted to the root switch user.

FTRACE is always enabled on extension switches and blades, and the trace data is automatically captured.

FTRACE Configuration

A default configuration for FTRACE is provided for each of the two DP complexes on the Brocade 7840 Switch, Brocade SX6 Blade and the Brocade FX8-24 Blade, and on the DP complex for the Brocade 7810 Switch. This allows tracing of events related to the DP complexes.

The `portcfg ftrace` *`slot/ve_port`* `cfg` command is interactive.

Use this command under the direction of an authorized support representative. FTRACE configuration settings are described in Changing FTRACE Configuration Settings on a Brocade 7840 Switch on page 250.

Brocade FX8-24 Blade

The default configuration creates four FTRACE buffers of 100,000 trace events that will be used until a trigger event (programmed trigger point in the logic) occurs. Trigger events include unexpected events or events that include FC abort sequences or other errors when emulation features are enabled on the tunnel.

The default configuration does not allow reuse of a trace buffer that includes one or more trigger events. The FTRACE configuration item that controls this function is called Auto Checkout (ACO). The default configuration of FTRACE provides for capturing, at a minimum, the first four error time periods in the four FTRACE buffers. That is because the default setting has enabled FTRACE ACO processing. When a buffer is checked out, it will not be reused until it is manually checked in or cleared through the supportsave process.

If the FTRACE configuration is changed so that ACO is disabled, then instead of post-filling and then checking out, the buffer is marked as triggered. If multiple trigger events subsequently occur so that all buffers are marked triggered, FTRACE will find the oldest triggered buffer and make it the current buffer. In this configuration, FTRACE will be set up to capture the last three error time periods.

FTRACE data contents are included in a switch supportsave capture. After the supportsave has been captured, the FTRACE buffers will be reset and all buffers that were previously either "checked out" or "triggered" return to an "unused" state.

Change the FTRACE ACO configuration using the following root command:

```
portcfg ftrace [slot/]vePort cfg
```

See Changing Configuration Settings on page 249 for more information.

Brocade 7810 Switch, Brocade 7840 Switch, and the Brocade SX6 Blade

FTRACE has been enhanced on the Brocade 7840 Switch, Brocade 7810 Switch, and the Brocade SX6 Blade to allow more trace saving options than for the Brocade FX8-24 Blade. The default FTRACE configuration has been changed on this platform as a result of those enhancements. For a display of the default configuration for the Brocade 7840 Switch or the Brocade SX6 Blade using the `portshow trace ve_port stats` command, see Brocade 7840 Switch Example on page 251.

The Brocade 7840 Switch and Brocade SX6 Blade include two data processing (DP) complexes. Each DP complex has an FTRACE instance. The default configuration for FTRACE on the switch or blade defines eight FTRACE buffers for trace events on each DP complex. The default configuration defines 300,000 trace entries (trace records) per trace buffer. The default FTRACE configuration enables auto checkout (ACO) for the first four buffers and disables ACO for the last four. The Brocade 7840 switch and Brocade SX6 Blade have a solid state disk (SSD) file system in each DP complex. This can be used to save copies of triggered FTRACE buffers. Use of the SSD to save FTRACE buffers is enabled by default and by the "Save to Flash" `portcfg ftrace ve_port cfg` command.

On the Brocade 7840 Switch or Brocade SX6 Blade, you can enable ACO for each defined FTRACE buffer. FTRACE processing varies when the FTRACE buffer is defined with ACO enabled or 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 will post fill that buffer to the end (or add the post fill percentage of more trace entries). When the post filling process is occurring the FTRACE buffer state will be reported as "post fill". When the post filling process has completed, the buffer state will be reported as "checked out," and the next sequential available buffer number will be assigned to the current buffer (state "current"). If all FTRACE buffers are marked as "checked out," FTRACE will no longer be recording trace entries. The default configuration therefore will capture at least the first four error traces, permanently check out those buffers, and then move them to the ACO-off buffers. FTRACE buffers that have been checked out will be saved in a supportsave capture. When the supportsave is complete, the buffers will return to an "unused" state and will be available for new traces. You can use the `portshow ftrace ve_port cmd` command to check in a checked out buffer.

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 will complete the same post filling process as described for ACO enabled. When completed, if the "Save to Flash" configuration option was enabled, the buffer will move to a "saving" state, and the next available buffer will be made as the current trace buffer. The Brocade 7840 Switch or the Brocade SX6 Blade will save as many as eight FTRACE buffers in the DP SSD file

system. If there are already eight saved FTRACE buffers in the file system, the oldest trace buffer will be replaced by the current buffer being saved. When the save-to-flash processing completes, the buffer will be marked as "triggered". If the "Save to Flash" option is not enabled, the buffer will be immediately marked as "triggered" and the next sequentially available FTRACE buffer will be marked as the "Current" buffer.

In the default configuration, FTRACE will therefore capture at least the first four error events (in buffers 0, 1, 2, and 3). It will capture the last three error events in triggered buffers (4-7) and will always have a current buffer. Buffers 4-7 will also potentially have as many as 10 saved prior trigger events 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 will be reset and all buffers that were previously either "Checked Out" or "Triggered" return to an "unused" state.

Change the FTRACE ACO configuration using the roo`portcfg ftrace` *`slot/ve_port`* `cfg` command. See Changing Configuration Settings on page 249 for more information.

Changing Configuration Settings

Use the root`portcfg ftrace` *`slot/ve_port`* `cgf` command to change FTRACE configuration settings. The configuration for FTRACE is defined using the first VE_Port on the switch or blade DP complex as follows:

- Brocade FX8-24 blade: VE_Port 22 on DP0 and VE_Port 12 on DP1
- Brocade 7840 switch: VE_Port 24 on DP0 and VE_Port 34 on DP1
- Brocade 7810 switch: VE_Port 12 on DP0
- Brocade SX6 blade: VE_Port 16 on DP0 and VE_Port 26 on DP1

To change FTRACE configuration settings on the first DP complex (DP0) on a Brocade 7840 Switch, if applicable, set the context where VE_ Port 24 is defined, and then issue the following command as the root user only:

```
portcfg ftrace 24 cfg
```

To change FTRACE configuration settings on the first DP complex (DP0) on a Brocade FX8-24 Blade, if applicable, set the context where the VE_Port 22 is defined, and then issue the following command as the root user only:

```
portcfg ftrace slot_number/22 cfg
```

To change FTRACE configuration settings on the second DP complex on a Brocade FX8-24 Blade (DP1), if applicable, set the context to where VE_port 12 is defined, and then issue the following command as the root user only:

```
portcfg ftrace slot_number/12 cfg
```

To change FTRACE configuration settings on the first DP complex (DP0) on a Brocade 7840 Switch, if applicable, set the context where the VE_Port 24 is defined, and then issue the following command as the root user only:

```
portcfg ftrace 24 cfg
```

To change FTRACE configuration settings on the second DP complex (DP1) on a Brocade 7840 Switch, if applicable, set the context to where VE_port 34 is defined, and then issue the following command as the root user only:

```
portcfg ftrace 34 cfg
```

Note that `portcfg` is an interactive command sequence and will prompt you for configuration items.

Changing FTRACE Configuration Settings on a Brocade 7840 Switch

Following is an example of the interactive command sequence that illustrates where you are prompted to change FTRACE configuration settings on a Brocade 7840 Switch. To change the settings, set the context where VE_Port 34 is defined, and then issue the `portcfg ftrace 34 cfg` command as root user only.

> **NOTE**
>
> User input lines in following example of this interactive command have been annotated to help you select configuration options. Those notes in italic font, such as ** Enables FTRACE (default is y) **, indicate options that you can modify. Those in between double asterisk characters indicate options that you should not modify without direction from a support representative.

```
switch:admin> portcfg ftrace 34 cfg

  ***  FTRACE INTERACTIVE CONFIGURATION  ***

  ***  Note: A reboot is necessary to    ***
  ***  activate a change in the number   ***
  ***  of buffers or records.            ***
Enable FTRACE?                    (Y,y,N,n): [y] y          *Enables FTRACE -default y*
Buffers                           (0-16): [8]               *Sets number of trace buffers -default 8*
Records (decimal, no commas)      (0-262,144): [300,000]    *Sets  number of trace records per buffer -
default 200,000*
Auto Checkout?                    (Y,y,N,n): [y]            *Enables ACO (default y)*
Auto Checkout is on, config at least 1 buffer accordingly.
    Auto Checkout buffer 0        (Y,y,N,n): [y]            *Enables ACO for buffer 0 -default y*
    Auto Checkout buffer 1        (Y,y,N,n): [y]            *Enables ACO for buffer 1 -default y*
    Auto Checkout buffer 2        (Y,y,N,n): [y]            *Enables ACO for buffer 2 -default y*
    Auto Checkout buffer 3        (Y,y,N,n): [y]            *Enables ACO for buffer 3 -default y*
    Auto Checkout buffer 4        (Y,y,N,n): [n]            *Disables ACO for buffer 4 -default n*
    Auto Checkout buffer 5        (Y,y,N,n): [n]            *Disables ACO for buffer 5 -default n*
    Auto Checkout buffer 6        (Y,y,N,n): [n]            *Disables ACO for buffer 6 -default n*
    Auto Checkout buffer 7        (Y,y,N,n): [n]            *Disables ACO for buffer 7 -default n*
Save to Flash?                    (Y,y,N,n): [y]            *Enables saving non-ACO to flash -default y*
Post Percentage (decimal)         (0-100): [5]              *Sets the post fill percentage -default 5*
Trace Mask (*)                    (0-ffffffff): [8000dffb]  **Sets the trace mask -default 8000dffb**
Trigger Mask (T)                  (0-ffffffff): [1]         **Sets the trigger mask -default 1**
Display Mask (-)                  (0-ffffffff): [ffffffff]  **Sets the trace display mask -default ffffffff**
Enable VE Traces?                 (Y,y,N,n): [y]            **Enables VE event traces -default y**
Enable FCIP Traces?               (Y,y,N,n): [y]            **Enables FCIP event traces -default y**
Enable TCPIP Traces?              (Y,y,N,n): [y]            **Enables TCP/IP event traces -default y**
Enable TCPIP Conn Traces?         (Y,y,N,n): [y]            **Enables TCP/IP Connection event traces -default
y**
Enable IP Traces?                 (Y,y,N,n): [y]            **Enables IP Event traces -default y**
Enable ARL Traces?                (Y,y,N,n): [y]            **Enables ARL Event traces -default y**
Enable Ethernet Traces?           (Y,y,N,n): [n]            **Disables Ethernet traces -default n**
Enable IP API Traces?             (Y,y,N,n): [y]            **Enables IP/API even traces -default y**
Enable FCIP MSG Traces?           (Y,y,N,n): [y]            **Enables FCIP Msg traces -default y**
Enable VDM Traces?                (Y,y,N,n): [n]            **Disables VDM traces -default n**
Configuration complete.
Operation Succeeded
switch:admin>
```

To correctly and completely delete an FTRACE configuration and reset to defaults, perform the following command sequences.

```
switch:admin> portcfg ftrace 34 del

  *** Note: This command will clear out  ***
  *** the current config and FTRACE will ***
  *** be reset to default values.        ***

Do you wish to continue?           (Y,y,N,n): [n] y

Operation Successful

switch:admin> reboot
/* After switch completes reboot sequence */
switch:admin> portcfg ftrace 34 cfg
/* repeat the configuration or leave as default */
```

Displaying FTRACE Status on a DP Complex

To display the current FTRACE status on an DP complex, issue the following command as the root user:

```
portshow ftrace [slot/]vePort stats
```

The *vePort* is in the current logical switch context.

Brocade 7840 Switch Example

Following is an example of displaying FTRACE status using the `portshow ftrace` *`slot/ve_port`* `stats` command. The example shows the default configuration for the Brocade 7840 Switch. A similar command can be used for the Brocade SX6 Blade.

```
switch:admin> portshow ftrace 34 stats

VE traces:          On-all        Trace Mask:        0x8000dffb (*)
FCIP Tunnel traces: On-all        Trigger Mask:      0x00000001 (T)
TCPIP traces:       On-all        Display Mask:      0xffffffff (-)
TCPIP Conn. traces: On-all        Tunnel Mask:       Inactive
IP traces:          On-all        Post trigger:      5% - 10000 events
ARL traces:         On-all        Record Size:       128
ETHERNET traces:    Off           Save to Flash:     Enabled
IP API traces:      On-all        FTRACE is:         Enabled
FCIP MSG traces:    On-all        Debug level:       4-Normal (low)
VDM traces:         Off           CLIB / HAL:        Off / Off

*-Bit 31 [0x80000000]: Software Structure
  -Bit 19 [0x00080000]: EtRX - Ethernet Received Frame
  -Bit 18 [0x00040000]: EtSX - Ethernet Send Frame to Peer
  -Bit 17 [0x00020000]: TnTX - Tunnel Received Peer Frame
  -Bit 16 [0x00010000]: TnSX - Tunnel Send Frame to Peer
*-Bit 15 [0x00008000]: FcT - FC FWD Frame From Peer
*-Bit 14 [0x00004000]: FcR - FC FWD Received Frame
  -Bit 13 [0x00002000]: Dsc - Discarded Frame
*-Bit 12 [0x00001000]: Data - Frame Data
*-Bit 11 [0x00000800]: State Change
*-Bit 10 [0x00000400]: CpRX - Frame Received From CP
*-Bit  9 [0x00000200]: CpSX - Frame Sent To CP
*-Bit  8 [0x00000100]: ToP - Sent To Peer
*-Bit  7 [0x00000080]: Tfx - Emulation FC Frame From Peer
*-Bit  6 [0x00000040]: Rfx - Emulation FC Received Frame
*-Bit  5 [0x00000020]: Sfx - Send Frame
*-Bit  4 [0x00000010]: Gfx - Generated Frame
*-Bit  3 [0x00000008]: FC SOFi1/2/3 or Class F Frames
  -Bit  2 [0x00000004]: FC SOFn1/2/3 Frames
*-Bit  1 [0x00000002]: Msg - Information
T*-Bit  0 [0x00000001]: Err - Error

+-----+----------+-----+--------+------------+-----+------+------+-------+--------+
|     |          |     |        |Trace Header|Wrap | In   | Out  |Switch | Switch |
|  Id | State    | ACO |  Size  |  Address   |Count| OXID | OXID |  Date |  Time  |
+-----+----------+-----+--------+------------+-----+------+------+-------+--------+
|   0 |  Current |  on | 200000 | 0x0b0f7480 |   0 | FFFF | FFFF |       |        |
|   1 |   unused |  on | 200000 | 0x0b0f7780 |   0 | FFFF | FFFF |       |        |
|   2 |   unused |  on | 200000 | 0x0b0f7a80 |   0 | FFFF | FFFF |       |        |
|   3 |   unused |  on | 200000 | 0x0b0f7d80 |   0 | FFFF | FFFF |       |        |
|   4 |   unused | off | 200000 | 0x0b0f8080 |   0 | FFFF | FFFF |       |        |
|   5 |   unused | off | 200000 | 0x0b0f8380 |   0 | FFFF | FFFF |       |        |
|   6 |   unused | off | 200000 | 0x0b0f8680 |   0 | FFFF | FFFF |       |        |
|   7 |   unused | off | 200000 | 0x0b0f8980 |   0 | FFFF | FFFF |       |        |
+-----+----------+-----+--------+------------+-----+------+------+-------+--------+
```

The table at the bottom of the output example has the following columns:

- Id: The FTRACE trace buffer identifier or buffer number.

- State: The FTRACE buffer state for that buffer number. The state can be one of the following:
  - Current: The buffer is the current active buffer in use for events.
  - Triggered: The buffer has been used to record an error event from the DP complex. This state is used only when the Auto Checkout option was disabled.
  - Checked Out: The buffer has been used to record an error event from the DP complex, and the buffer will not be overwritten.
  - Post Fill: A trigger event has been encountered, and the FTRACE buffer is currently being post-filled with a number of post-error events. Once the post-filling has been completed, the buffer will transition to either a "Checked Out" or "Triggered" state.
  - Unused: The buffer has not been used to capture any events. The buffer will be used when the prior buffer in the list transitions to either a "Checked Out" or "Triggered" state.
- ACO: Auto Checkout enabled (on) or disabled (off) status.
- Size: The number of trace records that are in the buffer.
- Trace Header Address: A memory address used internally for controlling access to the trace buffer.
- Wrap Count: The number of times that a trace buffer has been wrapped. The trace is a circular buffer that wraps after the size number of trace events has been exceeded.
- In OXID and Out OXID: Not used until the buffer is being analyzed.
- Switch Date: Indicates the system date when the buffer transitioned to either a "Checked Out" or "Triggered" state.

Brocade FX8-24 Blade Example

Following is an example of displaying FTRACE status using the `portshow ftrace` *`slot/ve_port`* `stats` command.

```
Slot 0:
VE traces (0-31): (0xffffffff)  On     Trace Mask:    0x8000fefb (*)
FCIP Tunnel traces (32-64):  On     Trigger Mask:  0x00000001 (T)
TCPIP traces (65):           On     Display Mask:  0x8000fefb (-)
TCPIP Conn. traces (66):     Off    Tunnel Mask:   Inactive
IP traces (67-83):           Off    Post trigger:  3% - 3600 events
ARL traces (84):             Off    Record Size:   128
ETHERNET traces (85-103):    Off    Auto Checkout: Enabled
IP API traces (104):         Off    FTRACE is:     Enabled
FCIP MSG traces (105):       Off    Debug level:   4-Normal (low)
VDM traces (106):            Off
 *-Bit 31 [0x80000000]: Software Structure
   Bit 19 [0x00080000]: EtRX - Ethernet Received Frame
   Bit 18 [0x00040000]: EtSX - Ethernet Send Frame to Peer
   Bit 17 [0x00020000]: TnTX - Tunnel Received Peer Frame
   Bit 16 [0x00010000]: TnSX - Tunnel Send Frame to Peer
 *-Bit 15 [0x00008000]: FcT - FC FWD Frame From Peer
 *-Bit 14 [0x00004000]: FcR - FC FWD Received Frame
 *-Bit 13 [0x00002000]: Dsc - Discarded Frame
 *-Bit 12 [0x00001000]: Data - Frame Data
 *-Bit 11 [0x00000800]: State Change
 *-Bit 10 [0x00000400]: CpRX - Frame Received From CP
 *-Bit  9 [0x00000200]: CpSX - Frame Sent To CP
   Bit  8 [0x00000100]: ToP - Sent To Peer
 *-Bit  7 [0x00000080]: Tfx - Emulation FC Frame From Peer
 *-Bit  6 [0x00000040]: Rfx - Emulation FC Received Frame
 *-Bit  5 [0x00000020]: Sfx - Send Frame
 *-Bit  4 [0x00000010]: Gfx - Generated Frame
 *-Bit  3 [0x00000008]: FC SOFi1/2/3 or Class F Frames
   Bit  2 [0x00000004]: FC SOFn1/2/3 Frames
 *-Bit  1 [0x00000002]: Msg - Information
T*-Bit  0 [0x00000001]: Err - Error
+-----+----------+--------+------------+-------+------+------+--------+--------+
|     |          |        |Trace Header| Wrap  | In   | Out  | Switch | Switch |
|  Id | State    |  Size  |  Address   | Count | OXID | OXID |  Date  |  Time  |
+-----+----------+--------+------------+-------+------+------+--------+--------+
```

```
|   1 | Current  | 100000 | 0x001f2f00 | 12344 | FFFF | FFFF |        |        |
|   1 |   unused | 100000 | 0x001f3180 |     0 | FFFF | FFFF |        |        |
|   2 |   unused | 100000 | 0x001f3400 |     0 | FFFF | FFFF |        |        |
|   3 |   unused | 100000 | 0x001f3680 |     0 | FFFF | FFFF |        |        |
+-----+----------+--------+------------+-------+------+------+--------+--------+
```

The table at the bottom of the output example has the following information:

- Id: The FTRACE trace buffer identifier or buffer number.
- State: The FTRACE buffer state for that buffer number. The state can be one of the following:
  - Current: The buffer is the current active buffer in use for events
  - Triggered: The buffer has been used to record an error event from the DP complex. This state is used only when the Auto Checkout option was disabled.
  - Checked Out: The buffer has been used to record an error event from the DP complex, and the buffer will not be overwritten.
  - Post Fill: A trigger event has been encountered, and the FTRACE buffer is currently being post-filled with a number of post-error events. Once the post-filling has been completed, the buffer will transition to either a “Checked Out” or “Triggered” state.
  - Unused: The buffer has not been used to capture any events. The buffer will be used when the prior buffer in the list transitions to either a “Checked Out” or “Triggered” state.
- Size: The number of trace records that are in the buffer.
- Trace Header Address: A memory address used internally for controlling access to the trace buffer.
- Wrap Count: The number of times that a trace buffer has been wrapped. The trace buffer is a circular buffer that wraps after the size number of trace events has been exceeded.
- In OXID and Out OXID: Not used until the buffer is being analyzed.
- Switch Date: Indicates the system date when the buffer transitioned to either a “Checked Out” or “Triggered” state.
- Switch Time: Indicates the system time when the buffer transitioned to either a “Checked Out” or “Triggered” state.