BROCADE
A Broadcom Company

White Paper

Brocade® Extension

Migrate to the Brocade 7850 Extension Platform: Refresh and Modernize Your Long-Distance Disaster Recovery Solutions with Ease

Overview

The lifeblood of any company is its data. Protecting that data—via business continuity and disaster recovery solutions—is paramount to any business's continued existence and success. For decades, Brocade® Extension has been the choice for long-distance business continuity and disaster recovery (BC/DR) solutions. At the heart of these solutions, you will find Brocade Extension Switches and Brocade Extension Blades for Brocade Directors.

The primary technology used to deliver high-performance and reliable data replication has been Fibre Channel over IP (FCIP). Brocade FCIP uses high-efficiency encapsulation to encapsulate Fibre Channel into TCP/IP for transport over a long-distance WAN from anywhere to anywhere—across the state to around the globe. Brocade continues to deliver the industry's most advanced extension products to achieve the most predictable, stable, and high-performance replication fabrics, allowing clients to reliably protect their most critical asset, their data.

The most common data protection scheme enterprise customers use is disk replication or mirroring. There are, and will continue to be, other methods for protecting data; however, array-based mirroring is one of the most commonly used methods in the enterprise market today. The Brocade Extension platforms are designed to excel in this environment, delivering the most outstanding data protection and performance regardless of distance.

With the ever-increasing data and the increase in WAN speeds from 1GbE to 10GbE to 100GbE and beyond, the older Brocade 7840 Extension platform has reached the end of its useful life. The Brocade 7840 was introduced in October 2014 and will enter its End-of-Life (EOL) cycle soon. For those responsible for business continuity and disaster recovery for enterprise data, it is irresponsible to ignore the underlying infrastructure. No organization wants to run replication of data for mission-critical workloads on infrastructure that is end-of-life and not receiving security updates. To avoid that, now is the time to plan a migration off the 7840 platform.

This paper focuses on the methods and strategies for refreshing the Brocade 7840 with the Brocade 7850, and describes the general principles of migrating and refreshing an existing environment to a new, modernized extension infrastructure.

Brocade Extension Platforms

Brocade offers three extension platforms to meet long-distance replication requirements: the Brocade 7810, Brocade 7850, and Brocade SX6 Extension Blade.

To assist our largest enterprise clients with director-class SAN switching platforms, Brocade offers the Brocade SX6 Extension Blade that integrates into the Brocade X7 Director. The Brocade SX6 uses high-efficiency encapsulation to transport Fibre Channel frames across an IP WAN. Data is transported to remote locations where secondary copies are safely kept. The Brocade SX6 interoperates with other Brocade Extension platforms for flexible designs and cost-effective solutions.

For clients who require high WAN throughput and are not inclined to deploy a director blade, Broadcom offers the Brocade 7850 Extension Switch. The Brocade 7850 and Brocade SX6 are considered enterprise-class systems and support all the solutions for open systems and mainframe (FICON) environments.

The Brocade 7850 Extension platform is based on Brocade Gen 7 ASIC technology; its predecessor was the Brocade 7840, based on Gen 5 technology. The Brocade 7850 is a 1U system that incorporates all of the advanced functions of the Brocade SX6. The Brocade 7850 can be integrated into enterprise-class solutions to support high throughput. The Brocade 7850 supports all FCIP and FICON disk replication solutions, IP Extension solutions, and mainframe solutions that leverage emulation techniques.

Figure 1: Brocade 7850 Angled View

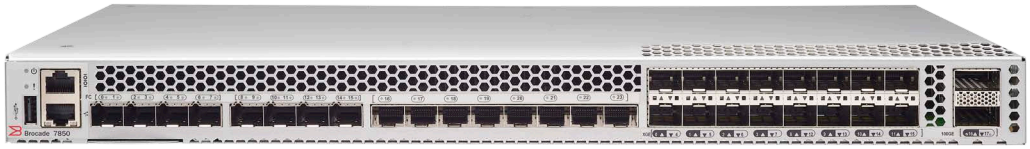

Figure 2: Brocade 7850 Rear View

Brocade Extension Solution Designs

Brocade SAN Extension platforms allow users to create highly effective infrastructure to meet their current and future data replication needs. FCIP technology and its ability to intermix with a Fibre Channel storage area network (SAN) allow for endless possibilities of infrastructure designs. To help narrow the scope of this paper, we will focus on some of the most common network designs for BC/DR, with particular emphasis on migrating from SAN Extension platforms that have gone into their EOL process/cycle.

The most common FCIP deployment contains dual, redundant fabrics or paths that connect the storage array through the IP WAN to the remote site. Having two paths from one site to the other offsets the impact if one of the paths fails. This redundancy and data path protection are inherent in building the most resilient infrastructure.

Users should have two different IP WAN service providers when two or more IP networks are required. This provides additional protection by having separate routes between the two sites. History has shown that users not confirming disparate routes can be impacted by events out of their control, such as a backhoe digging for a new sewer line.

This dual fabric configuration is illustrated in the following diagram.

Figure 3: Example of a Dual Fabric Storage Replication Network

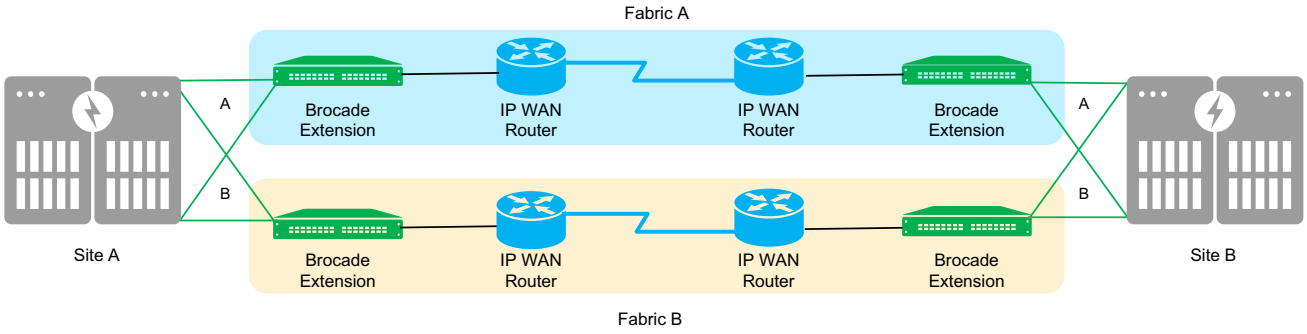

Transition from the Brocade 7840 Extension Switch

The Brocade 7840 Extension Switch is two generations older than the Brocade 7850 Extension Switch, and it is unable to run the same Fabric OS® (FOS) code levels. It's also incapable of delivering the same performance, and cannot support the newer Gen 7 features. Because of this and other incompatibilities between the two platforms, connectivity between Brocade 7840 and 7850 platforms is not supported.

When introducing the Brocade 7850 Extension Switch platform into an environment with existing Brocade 7840 switches, you should plan to deploy a pair of Brocade 7850 units alongside the 7840 platforms in any location that needs to support connectivity to new 7850 platforms. The 7840 core units will continue to support connectivity to any sites with 7840 platforms, whereas the new 7850 core units support connectivity to sites with Brocade 7850 switches. Transition to Brocade 7850 switches that can run at full performance using the entire Gen 7 feature set while gradually moving off the Brocade 7840s.

Figure 4: Gradual Transition from Brocade 7840 to Brocade 7850 Extension Switches

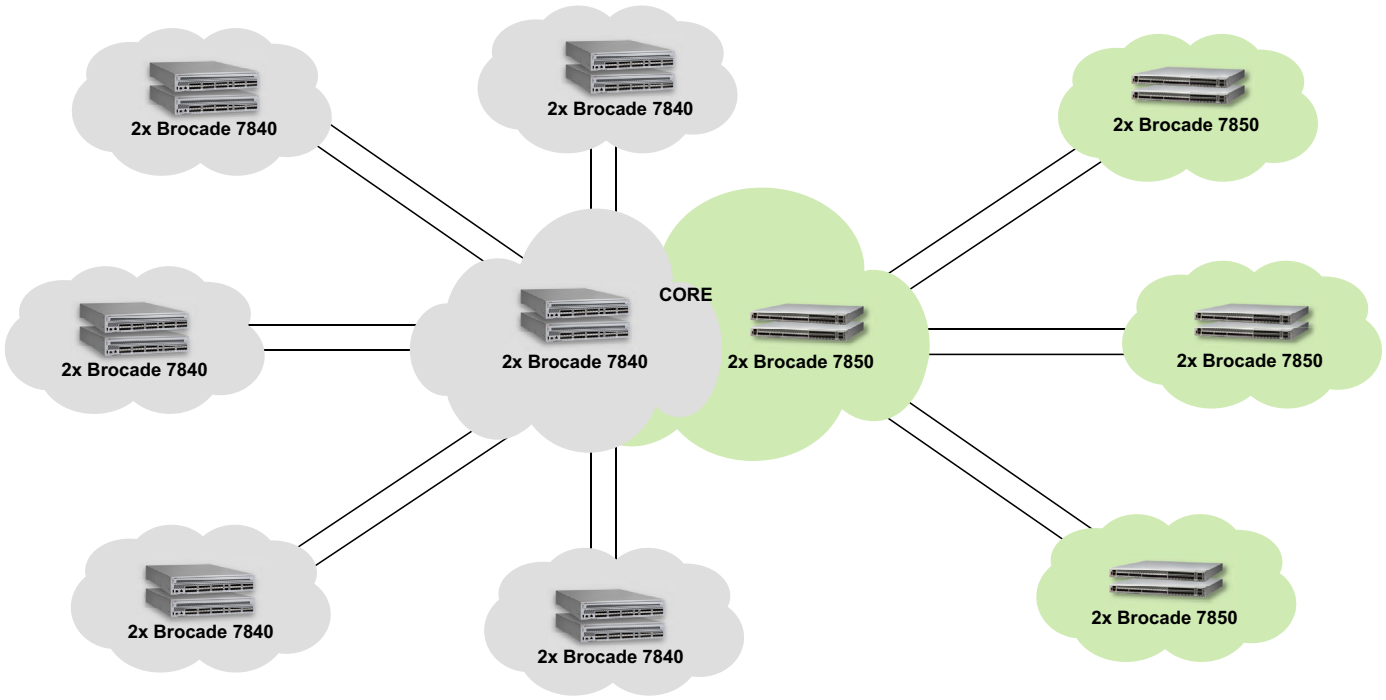

This approach provides flexibility in making the transition. Older Brocade 7840 units can be replaced with the new Brocade 7850 units over time–when it is most suitable to your environment. As the older Brocade 7840 units are replaced, the replication traffic will gradually move to the Brocade 7850 units at the core, and the older Brocade 7840 core units will handle progressively less and less traffic.

Once the final Brocade 7840 units at remote sites have been replaced, the remaining Brocade 7840 core units can be decommissioned.

Figure 5: Complete Transition to Brocade 7850 Switches

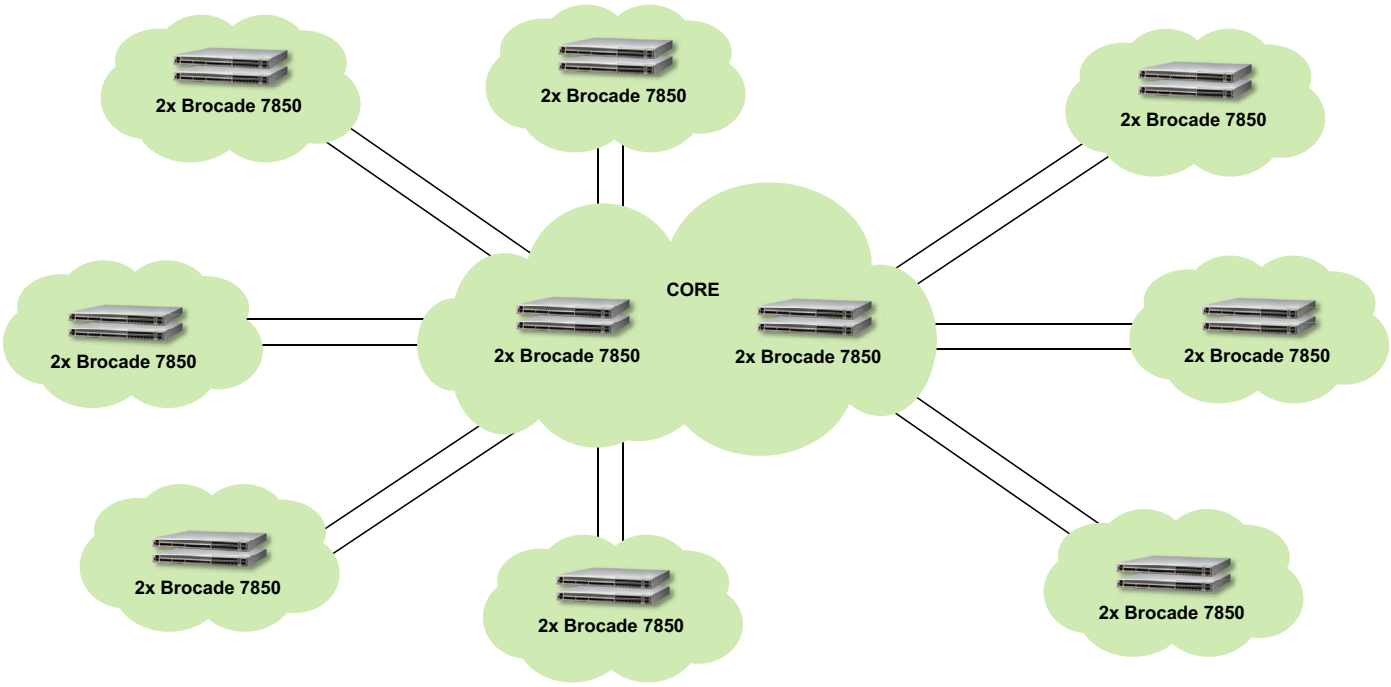

Migration Methodology and Techniques

The Brocade 7850 Extension Switch provides greater performance and improves features on a familiar platform, built on generations of proven Brocade Extension technology.

One of the most common inquiries about migrating to a new Brocade Extension platform is, *Will I need to take an outage, and if so, how long will that last?* The short answer is maybe. Users can avoid an outage and downtime by following proper conversion methods.

To reduce the risk of an outage, stage and configure new platforms to mimic or accept configurations in use today. Suppose you are embarking on a consolidation effort, increasing the number of WAN connections, or upgrading to a higher-speed WAN; additional steps are needed to integrate such changes. For example, new, reused, or changed IP addresses may be required. Additionally, changes to interface connections, IP routes, VE_Port numbers, bandwidth, rate limiting, and Traffic Control List (TCL) must be considered when extension platforms are refreshed.

Refer to the *Brocade Fabric OS Extension User Guide* when performing upgrades and making changes.

A dual-fabric topology provides a seamless transition from one generation of a platform to another. Using one of the two paths, you can keep replication traffic flowing without taking an outage. Use Figure 3 as the example network.

Follow the guidelines below for migrating a one-for-one replacement of a Brocade 7840 to a Brocade 7850 with no changes to Fibre Channel connections or the IP WAN:

1. Preliminary work: Run Brocade SAN Health® to capture a snapshot of your environment and see if you have any current network health problems that need attention before the migration. See Additional Resources for more information about Brocade SAN Health.
2. Gather information from the Brocade 7840s required to build and configure the replication connections on the Brocade 7850s, such as IP interfaces, IP routes, eHCL, IPsec, compression, tunnels, circuits, QoS settings, FastWrite, OSTP, failover metrics, failover groups, and TCLs.
3. Install the Brocade 7850s and configure the management IP address to access the platform.
4. Leave disconnected all Fibre Channel and network connections.
5. Apply power to the Brocade 7850 Extension platforms in Site A and Site B.
6. There are two ways to configure the Brocade 7850s: Brocade SANnav™ or the CLI.
7. With the configuration information gathered from the Brocade 7840s, configure the Brocade 7850s in Site A and Site B.
8. Disable power to the Brocade 7850 Extension platform.
9. Select which fabric you intend to migrate first, and if required by the storage vendor's process, suspend the path for the mirrored interface on the storage array. Review your storage vendor's administration guide for the proper technique to take down one or more mirrored paths.

NOTE: With a dual-fabric architecture, the other fabric continues replication; however, a replication backlog may generate since half the bandwidth is available.

10. To migrate the first path, turn off power to the corresponding Brocade 7840 Extension platforms in Site A and Site B.
11. Remove the Fibre Channel and network connections from each Brocade 7840, and reattach the cables to the appropriate ports on the replacement Brocade 7850s.
12. Apply power to the Brocade 7850 Extension platforms in Site A and Site B.
13. Check for error conditions and messages.
14. While the new path is offline to the mirror, run the Brocade 7850 traffic generator (WAN Test Tool) to confirm that the IP network and WAN reliably support the data flow requirements.
15. Once the path has been verified, re-enable the mirror (if required).
16. Monitor the mirror, the Brocade 7850s, and the network for errors. Confirm that replication is running as expected.
17. Repeat these steps for the other fabric once the fabric is operational and verified.

NOTE: An outage can be eliminated by converting one path at a time.

Configure Extension

This section is an adjunct to the*Brocade Fabric OS Extension User Guide* and the *Brocade Fabric OS Command Reference Manual*. Refer to the user guide and command reference for detailed information.

The example below is specific to FCIP and IP Extension on the Brocade 7850 Extension platform.

Several distinct steps are required to configure extension:

- Plan your extension architecture. Refer to the reference architecture below.
- Gather the necessary information to configure your architecture.
- Configure each Brocade 7850.
- Validate each Brocade 7850.
- Direct end-device storage traffic to the appropriate extension platform.
- Validate storage traffic across extension.

Reference Architecture Example

For the following discussion, refer to the 1x1 architecture and data flow diagrams below. In this example, each site has one Brocade 7850 Extension Switch. There are two sites: local and remote. The 1x1 architecture was used in this example for simplicity; 2x2 is considered the best practice, with two extension platforms at each location.

Figure 6: Layer 3 Deployment Reference Architecture

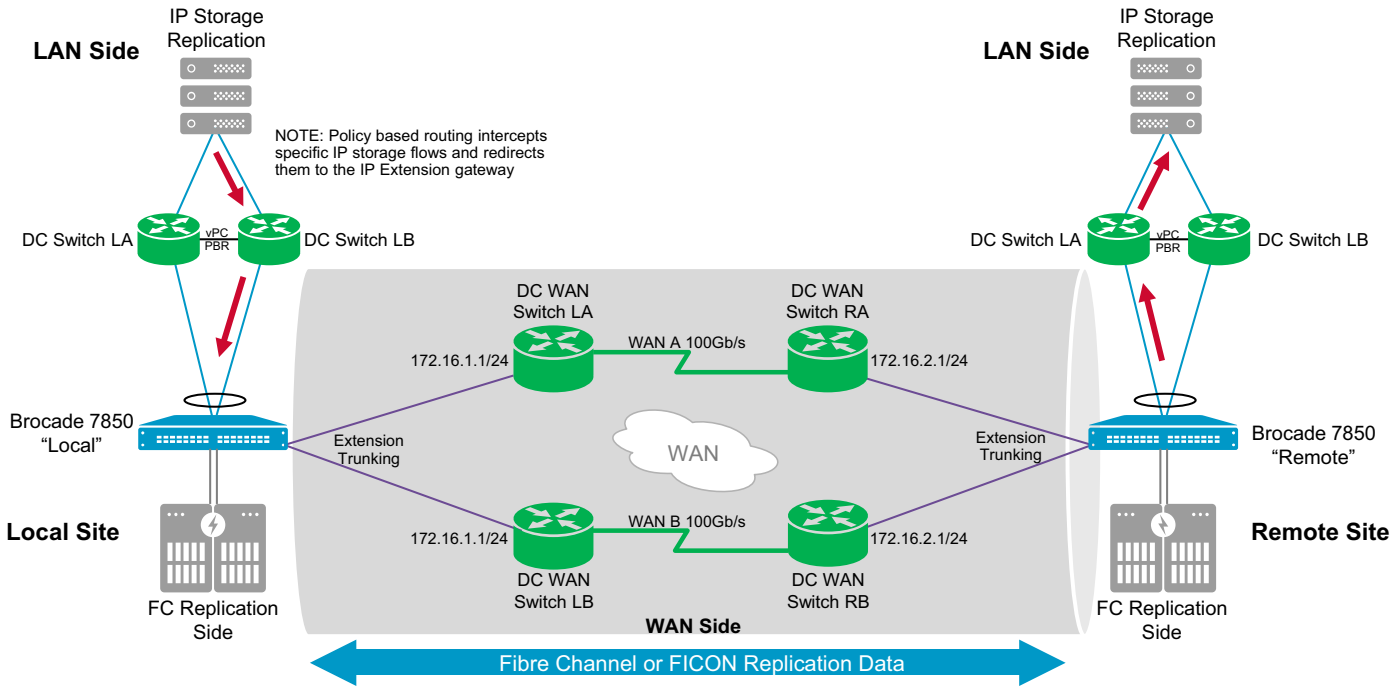

Figure 7: Layer 2 Deployment Data Flow

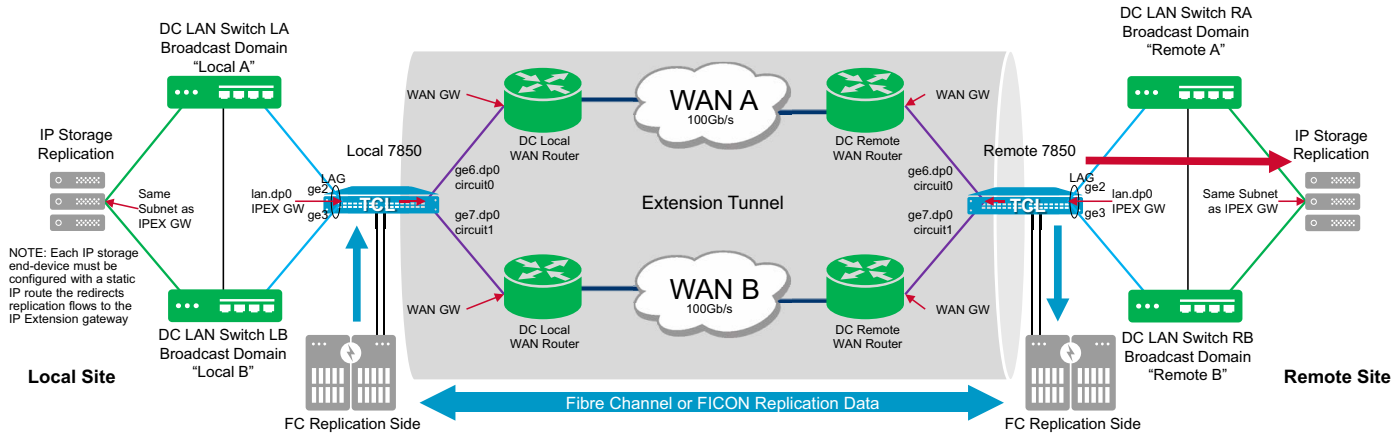

Figure 8: Layer 3 Deployment Data Flow

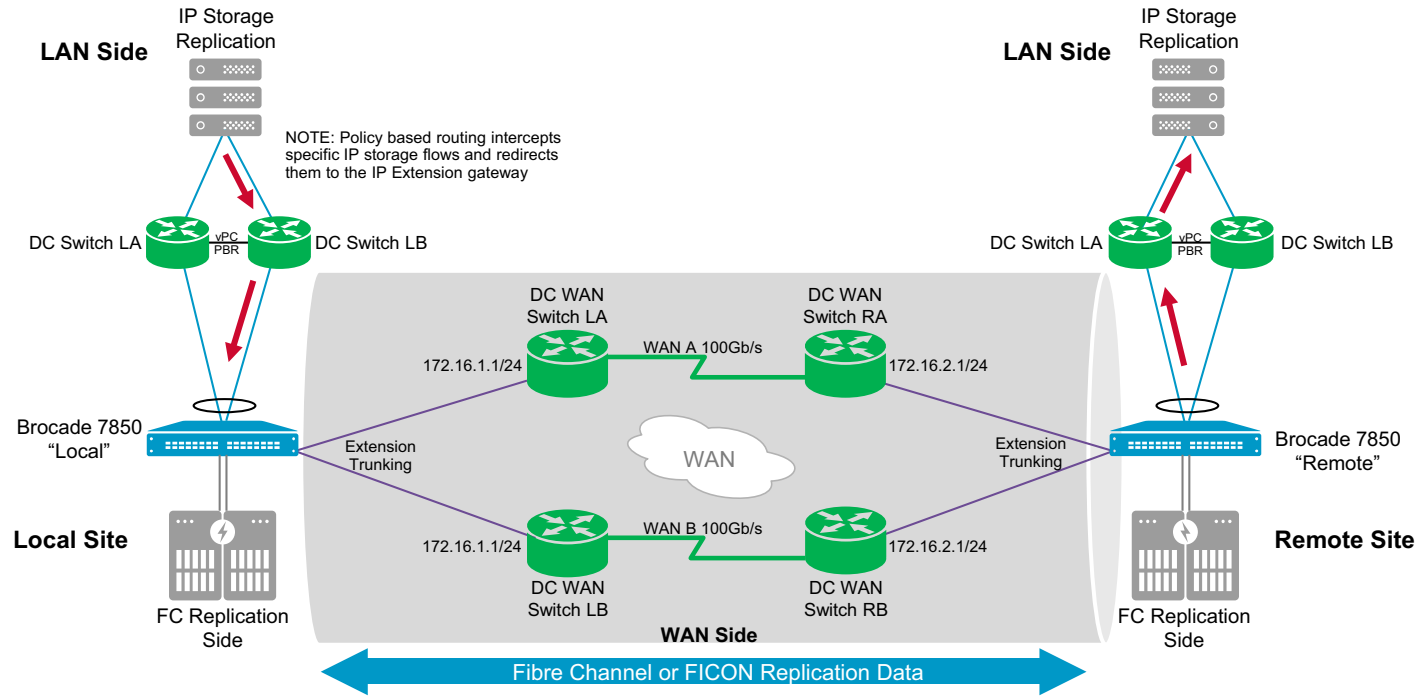

Architecture Considerations

Architecture considerations for this example are as follows:

- Each Brocade 7850 has three reference sides: LAN, WAN, and FC (or FICON). Each side has unique configuration requirements.
- The term local refers to where FCIP and IP Extension data originates, and the term remote refers to where data is replicated.
- The extension tunnel (WAN side) has two circuits. Each WAN link carries one extension circuit. The maximum circuit bandwidth is 25Gb/s.
- For higher availability, a circuit path should be unique through the IP WAN network.

- Brocade Extension Trunking refers to the use of multiple circuits (two circuits in this example) comprising a single tunnel, which is represented by a VE_Port. The circuits aggregate into the overall tunnel bandwidth. Circuits take diverse paths for enhanced availability. This means the tunnel passes through different Ethernet switches and both WAN links. Circuits enable failover/failback while never losing or delivering data out of order. Two or more circuits are best practice and strongly recommended.
- Each site has two LAN switches, which are interconnected to form a single logical switch. This may be referred to in several ways: Virtual PortChannel (vPC), Multi-chassis Link Aggregation Group (MLAG), Virtual Switching System (VSS), Virtual Link Trunking (VLT). Link Aggregation Group (LAG) is a method of using multiple links for redundancy while maintaining a single logical connection. From the Brocade 7850, only a single LAN-side connection is allowed via a single physical Ethernet link or multiple links forming a single logical connection. This reference architecture uses a LAG on each LAN side. Using a LAG is not required but is highly recommended for higher availability.
- Gigibit Ethernet (GE) interfaces on Brocade Extension platforms will not cause network loops and do not require Spanning Tree Protocol (STP). There are two choices for IP traffic that enters a Brocade Extension platform: one, it matches a TCL allow rule and is forwarded across the specified target; two, it does not match a specified TCL allow rule and is dropped. Entering LAN-side traffic cannot be switched back to the LAN from the extension platform.
- If more than one LAN-side Ethernet link is to be used, the links must form a portchannel.
- Each WAN link is 100Gb/s.
- Layer 2 and Layer 3 IP Extension are shown.
  - For Layer 2, the storage end-devices need IP routes that send replication traffic to the IP Extension Gateway. The IP Extension Gateway must be on the same subnet as the replication ports. No LAN-side IP route <ipif lan.dp#> is needed on the extension platform.
  - For Layer 3, the storage end-devices do not need special IP routes; instead, the IP network is configured to intercept specific flows (usually the replication ports' source and destination IP addresses) and redirect those flows to the IP Extension Gateway. A LAN side IP route <ipif lan.dp#> is required on the extension platform.
- FCIP and IP Extension can use the same extension tunnel as this example depicts.
- TCL only applies to IP Extension flows, not FCIP.
- The data processor specified in IPIF and IP Routes is the data processor that owns the tunnel VE_Port; TCL refers to it as the target.
- Transmission Control Protocol (TCP) flows are optimized. A limited number of User Datagram Protocol (UDP) flows are supported and not optimized.
- Broadcast, Unknowns, and Multicast (BUM) traffic is supported in limited numbers and is not optimized.
- The number of TCP sessions is limited to 512 per data processor on the Brocade 7850. In the case of control traffic, which does not require optimization, it is possible to pass the traffic via the TCL without consuming a TCP session. Optimization is not done.

Assumptions and Prerequisites

The assumptions listed below are specific to this deployment guide and example. Change configuration parameters to suit your needs and environment.

- The Brocade 7850 Extension platforms have been installed.
- Cabling for power, management, LAN, and WAN Ethernet has been completed.
- The Brocade 7850 Extension platform fundamental configuration tasks have been completed.
- The extension, LAN, WAN, and IP network are operational end to end.
- The WAN has adequate available bandwidth to add the extension traffic.
- The LAN-side data center switches to form a single logical switch.
- There is one LAN-side Ethernet link to each data center switch.
- The two LAN-side Ethernet links form a portchannel for redundancy.
- The administrator team has access to all devices that need configuration.
- The Brocade 7850 Extension platforms have no preexisting extension configuration.

- The LAN side will use two 10GE interfaces (GE2 and GE3).
- The WAN side will use two 100GE interfaces (GE16 and GE17).
- VE_Port 24 (TCL target 24) is used, which lives on DP0.
- Tunnel 24 (VE24) has two circuits: VE24-Cir0 and VE24-Cir1.
- Each WAN-side Ethernet link is connected to a different WAN IP network switch.
- Circuits for the LAN side and WAN side use the default maximum transmission unit (MTU) of 1500 bytes.
- The IP subnet and address information on the IP storage replication ports is known.
- The IP Extension gateway can be added to the IP storage replication ports' subnet.
- No VLAN tagging is required on the extension platforms to connect to the data center LAN.
- No QoS implementation is required for storage traffic.
- Each circuit will require a minimum bandwidth of 10Gb/s to meet the application requirements.
- Each circuit will require a maximum bandwidth of 25Gb/s to meet the application requirements.
- The following WAN-side IP subnets are used. From these subnets, IP addresses are assigned to circuit IPIFs.

Side	Subnet	Cir0	Cir1	WAN Gateway
WAN (Local)	172.16.1.0/24 (255.255.255.0)	172.16.1.3	172.16.1.4	172.16.1.1
WAN (remote)	172.16.2.0/24 (255.255.255.0)	172.16.2.3	172.16.2.4	172.16.2.1

- The following LAN-side IP subnets are used. From these subnets, an IP Extension gateway is assigned.

Side	Subnet	IP Extension Gateway
LAN (local)	10.10.10.0/24 (255.255.255.0)	10.10.10.1
LAN (remote)	192.168.0.0/24 (255.255.255.0)	192.168.0.1

- The gateways can differ for each circuit. Often, network gateways are virtual and created using Hot Standby Router Protocol or Virtual Router Redundancy Protocol (VRRP). A VRRP virtual gateway might be called a Virtual IP.
- The minimum and maximum values represent the rates for Adaptive Rate Limiting (ARL) for each circuit. If ARL is not used and Committed Information Rate (CIR) is used, set the minimum equal to the maximum. The minimum and maximum values must be set on each circuit. The local and remote settings must be identical.

Limit	ARL Rate per Circuit	per Tunnel
Minimum	10Gb/s	20Gb/s
Maximum	25Gb/s	50Gb/s

- Circuit metrics and groups are not used in this example. There are two circuits; traffic failovers to the remaining circuit when one goes offline.

Configure Brocade 7850 Extension Switches

This section provides a generic configuration sequence offering users the concepts needed to operationalize Brocade Extension. Refer to the *Brocade Fabric OS Extension User Guide* for details on each configuration task and command syntax.

The following is a list of the configuration commands:

```
version
portcfgpersistentdisable < port-num | port-range >
portcfgge < ge# > --set < -lan | -wan >
portcfgge < ge# > --set -speed < 1G | 10G | 25G >
portchannel --create < po_name > -type < static | dynamic > [-key < po-key-num >]
portchannel --add < po_name > -port < port-num | port-range >
portcfg ipif < ge_int.dp# | lan.dp# > <verb> [<args>]
portcfg iproute < ge_int.dp# | lan.dp# > <verb> [<args>]
portcfg ipsec-policy < name > < verb > [<args>]
portcfg fciptunnel  < ge_int.dp# > < verb > [<args>]
portcfg fcipcircuit < ge_int.dp# > < verb > < cir# > [<args>]
portcfg tcl < name > < verb > [<args>]
portcfgpersistentenable < port-num | port-range >
```

The following steps provide extension configuration assistance:

1. Verify each extension platform's Brocade FOS version. The version on both ends must be the same during regular operation. Short periods while upgrading the platforms are acceptable.

   ```
   SW37_7850_A:FID128:admin> version
   Kernel:     5.4.66_rt38
   Fabric OS:  v9.2.0
   Made on:    Mon Apr 17 21:15:45 2023
   Flash:      Tue May 9 06:40:38 2023
   BootProm:   1.0.38-sb
   ```

2. Disable and enable ports. Disabling ports is typically not necessary when configuring extension; however, if the extension platform is a production switch with replication fabric or connected by ISL to a production fabric without replication traffic logical switch isolation, it is best practice not to merge the local and remote fabrics while configuring extension inadvertently. Inadvertently merging fabrics through extension during configuration can result in production fabric instability.

   Turn off the VE_Port during configuration or disable the ISL E_Ports connecting the production fabric to prevent merging. After configuring the extension tunnel, enable the VE_Port and verify proper operation. If ISL E_Ports were disabled, verify proper tunnel operation, then enable the E_Ports. The following command does not query “Are you sure?” and returns no response after being issued.

   The commands to disable or enable VE_Port 24.

   ```
   SW37_7850_A:FID128:admin> portcfgpersistentdisable 24
   SW37_7850_A:FID128:admin> portcfgpersistentenable 24
   ```

   The commands to disable or enable E_Port 0 (example E_Port).

   ```
   SW37_7850_A:FID128:admin> portcfgpersistentdisable 0
   SW37_7850_A:FID128:admin> portcfgpersistentenable 0
   ```

3. Set GE mode. The extension GE interfaces are in either LAN or WAN mode. WAN (shown as FCIP in `switchshow`) is the default setting. The 100GE interfaces cannot be set to LAN mode; they can only be in WAN mode. WAN mode is used for a tunnel's circuits. If a GE interface will be used to connect IP storage via the data center's LAN, the GE interface must be set to LAN mode. Up to eight GE interfaces can be set to LAN mode while eight remain in WAN mode. Changing GE interfaces from WAN to LAN mode or vice-versa does not require a reboot and is only briefly disruptive to the interface being changed.

   The following commands set ge0 to LAN or WAN mode.

```
SW37_7850_A:FID128:admin> portcfgge ge0 --set -lan
Operation Succeeded.

SW37_7850_A:FID128:admin> portcfgge ge0 --set -wan
Operation Succeeded.
```

   The following command shows the status of all GE interfaces.

```
SW37_7850_A:FID128:admin> portcfgge --show
 Port          Speed     Flags     Channel  FEC   Lag Name
-----------------------------------------------------------------------------
 ge0           25G       --L-      N/A      CL108 -
 ge1           10G       ----      N/A      Off   -
 ge2           10G       ----      N/A      Off   -
 ge3           10G       ----      N/A      Off   -
 ge4           10G       ----      N/A      Off   -
 ge5           10G       ----      N/A      Off   -
 ge6           10G       ----      N/A      Off   -
 ge7           10G       ----      N/A      Off   -
 ge8           10G       ----      N/A      Off   -
 ge9           10G       ----      N/A      Off   -
 ge10          10G       ----      N/A      Off   -
 ge11          10G       ----      N/A      Off   -
 ge12          10G       ----      N/A      Off   -
 ge13          10G       ----      N/A      Off   -
 ge14          10G       ----      N/A      Off   -
 ge15          10G       ----      N/A      Off   -
 ge16          100G      ----      N/A      CL91  -
 ge17          100G      ----      N/A      CL91  -
-----------------------------------------------------------------------------
Flags: A:Auto-Negotiation Enabled  C:Copper Media Type
L:LAN Port G: LAG Member
```

4. Set the GE interface to the desired speed: 1, 10, or 25Gb/s.

   The sixteen GE interfaces can be set to any desired speed; however, GE interfaces live in groups, and the speed must be consistent within a group. On WAN side GE interfaces, a slower and faster speed within a GE port group causes the blocking of the faster port. GE blocking does not occur on LAN-side interfaces. Refer to the *Brocade FOS Extension User Guide* for the details. The 100GE interfaces only support 100Gb/s.

   The following command sets the ge0 speed to 25Gb/s or 10Gb/s. The inserted optic must support the selected speed.

```
SW37_7850_A:FID128:admin> portcfgge ge0 --set -speed 25G
Operation Succeeded.

SW37_7850_A:FID128:admin> portcfgge ge0 --set -speed 10G
Operation Succeeded.
```

   The `portcfgge --show` command shows each interface's set speed.

5. If configuring IP Extension, create a portchannel (LAG) to the data center LAN switch. Skip this step if IP Extension will not be used. The following commands show the creation of a dynamic portchannel named MyPO. The portchannel example uses the key number 785 to identify its links to the data center LAN switch. Any unique number between 1 and 1000 can be used; select a number not already used by another portchannel. Two GE ports (two links) are added to the portchannel. Lastly, there is a command that shows the status of portchannels.

   Give the portchannel a human-readable name. Dynamic portchannel is recommended. If connecting the portchannel to more than one LAN switch, the switches must support a portchannel across them, which means the switches are logically a single switch (i.e., vPC, MLAG, VLT).

   ```
   SW37_7850_A:FID128:admin> portchannel --create MyPO -type dynamic -key 785
   ```

   The following commands add or delete a GE interface (ge0) to the MyPO portchannel. Only LAN-side GE interfaces can be added to a portchannel; the WAN side does not support portchannels.

   ```
   SW37_7850_A:FID128:admin> portchannel --add MyPO -port ge0
   SW37_7850_A:FID128:admin> portchannel --delete MyPO -port ge0
   ```

   The following command shows the status of the MyPO portchannel.

   ```
   SW37_7850_A:FID128:admin> portchannel --show
   Name                            Type        Oper-State   Port-Count   Member Ports
   ------------------------------------------------------------------------------
   MyPO                            Dynamic     Offline      2            ge0 ,ge1
   ```

6. If configuring IP Extension, an IP Extension gateway must be created. The IP Extension gateway is an IPIF designated as lan.dp# with an assigned IP address and mask, and optionally the MTU and VLAN tag. The Brocade 7850 has two data processors (DP0 and DP1); the DP number must be specified when configuring the IP Extension gateway. The DP number is the data processor that owns the VE_Port used for the tunnel to the remote site.

   One or more IP Extension gateways may exist depending on the environment. For a Layer 2 deployment, one IP Extension gateway is needed per subnet that IP storage end devices are on. A maximum of eight IP Extension gateways can be configured per data processor. For a Layer 3 deployment, only one IP Extension gateway is needed for communicating with the local data center router. The local data center router forwards the IP storage traffic to the end device's subnet.

   The following command creates an IP Extension gateway on DP0. The gateway IP address is 192.168.0.4 and has a mask of 255.255.255.0. The command shows if the operation was successful or not.

   ```
   SW37_7850_A:FID128:admin> portcfg ipif lan.dp0 create 192.168.0.4/24
   Operation Succeeded.
   ```

   The following command shows a list of IPIFs that have been created. It shows the MTU; the default is 1500 bytes. It also indicates if a VLAN ID has been configured; 0 means no VLAN is configured and is just an access port.

   ```
   SW37_7850_A:FID128:admin> portshow ipif
    Port         IP Address                  / Pfx  MTU   VLAN  Flags
   ------------------------------------------------------------------------------
    lan.dp0      192.168.0.4                 / 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
   ```

7. If configuring IP Extension for Layer 3, a LAN side IP route must be created. If configuring a Layer 2 deployment, skip this step. In a Layer 3 deployment, on the LAN side, IP Extension sends the IP storage data to a local data center router to deliver the data to the end device. IP Extension does not need to be configured with a gateway for each subnet; the end devices do not need to be configured with static routes to the IP Extension gateway. However, the network must be configured to intercept the IP storage flows and forward them to the IP Extension gateway.

   The following command creates a LAN-side IP route that forwards all traffic heading for subnet 10.10.10.0/24 to router gateway 10.0.0.1.

```
SW37_7850_A:FID128:admin> portcfg iproute lan.dp0 create 10.10.10.0/24 10.0.0.1
Operation Succeeded.
```

8. Create a WAN-side IPIF for the tunnel's circuits and HA. Each circuit passes through a WAN-side GE interface and terminates on its data processor's VE_Port. When creating an IPIF, the GE interface, data processor, IP address, and subnet mask are specified. An HA IPIF is on the opposite data processor as the circuit it protects, although the GE interface can be the same. When the IP address is added to a tunnel's circuit, the IPIF indicates which GE interface will be used for the circuit. Each circuit below has its own GE interface for redundancy, and the eHCL (HA) circuits are configured to use the same GE interfaces.

   The following commands show the IPIF creation on ge16 and ge17 for DP0. The IP addresses are 172.16.1.3 and 172.16.1.4 and have a mask of 255.255.255.0 (/24).

```
SW37_7850_A:FID128:admin> portcfg ipif ge16.dp0 create 172.16.1.3/24
Operation Succeeded.

SW37_7850_A:FID128:admin> portcfg ipif ge17.dp0 create 172.16.1.4/24
Operation Succeeded.
```

   Configure eHCL (HA) IPIFs on DP1 for the above circuits. The same IP addresses above cannot be used for eHCL (HA); they must be unique.

```
SW37_7850_A:FID128:admin> portcfg ipif ge16.dp1 create 172.16.1.5/24
Operation Succeeded.

SW37_7850_A:FID128:admin> portcfg ipif ge17.dp1 create 172.16.1.6/24
Operation Succeeded.
```

9. Create WAN-side IP routes for circuits and eHCL (HA). The IP routes for HA are the same if the identical subnet and GE interface are used. As shown below, the difference is that the IP route needs to be added to the opposite DP.

   If the WAN IP network is Layer 2, the same subnet end-to-end, the WAN-side IP routes do not need to be configured. In the example below, the subnets are different at each end.

   LAN-side IP routes <lan.dp#> do not forward traffic to the WAN side; TCL is used. WAN-side IP routes <ge#.dp#> do not forward traffic toward the LAN side; LAN-side IP routes are used in a Layer 3 deployment.

   The following commands create IP routes from the 172.16.1.0/24 site to the 172.16.2.0/24 site. The IP network's gateway to the remote side is 172.16.1.1. An IP route must be created for each unique GE interface/DP pair configured.

```
SW37_7850_A:FID128:admin> portcfg iproute ge16.dp0 create 172.16.2.0/24 172.16.1.1
Operation Succeeded.

SW37_7850_A:FID128:admin> portcfg iproute ge17.dp0 create 172.16.2.0/24 172.16.1.1
Operation Succeeded.
```

The following commands create IP routes for eHCL (HA). HA IP routes are configured on the opposite DP. Without these IP routes, eHCL will not operate.

```
SW37_7850_A:FID128:admin> portcfg iproute ge16.dp1 create 172.16.2.0/24 172.16.1.1
Operation Succeeded.

SW37_7850_A:FID128:admin> portcfg iproute ge17.dp1 create 172.16.2.0/24 172.16.1.1
Operation Succeeded.
```

10. Create an IPsec policy, which is required when implementing data encryption in flight. Encryption is optional but strongly recommended; otherwise, data is sent in the clear. Enabling encryption causes no negative performance impact. Preshared key (PSK) and public-key infrastructure (PKI) authentication methods are supported. A PSK can be 16 to 64 alphanumeric characters in length. Assign the IPsec policy a human-readable name.

    The following command creates an IPsec policy named MyIPsec. The example below uses PSK. The user can randomly generate a key. Both ends must use the same key, but remembering or recording the key is not essential. If the tunnel needs to be altered, create a new PSK and update the policy at both sites. Losing a key will not result in data loss, as the key is only used for data in flight. IPsec is applied to each circuit of the tunnel.

    ```
    portcfg ipsec-policy MyIPsec create --preshared-key
    w7ffuffx9zvgt6wrr3pka2mzd9o6bz30vwb0y3u67p3taw5wjxmwp9t7brwau98v
    ```

11. Create the extension tunnel. In the following command, tunnel 24 is created. IP Extension was enabled, and compression was set to deflate. IP Extension is not enabled on tunnels by default. FCIP compression was set to fast-deflate. IPsec was enabled with the policy MyIPsec.

    The configuration staging method is used. No circuit configuration is done in the `fciptunnel` command. All circuit configuration is done later in the `fcipcircuit` commands.

    ```
    SW37_7850_A:FID128:admin> portcfg fciptunnel 24 create --ipext enable --ipsec MyIPsec --ip-
    compression deflate --fc-compression fast-deflate
    Operation Succeeded.
    ```

12. Create two circuits for tunnel 24: circuit 0 and circuit 1. Up to ten circuits can be created per tunnel. At least one circuit is required before a tunnel can transmit data. Creating multiple circuits for a tunnel transforms the tunnel into a Brocade Extension Trunk. Each circuit requires identical settings on each end, but each circuit can be configured uniquely. Use multiple circuits to establish different network paths, which enhances availability.

    ```
    SW37_7850_A:FID128:admin> portcfg fcipcircuit 24 create 0
    Operation Succeeded.

    SW37_7850_A:FID128:admin> portcfg fcipcircuit 24 create 1
    Operation Succeeded.
    ```

13. Add local and remote IP addresses to the circuits. These IP addresses are required. The extension platforms will verify the local IPIF (IP address); the remote IPIF is not verified during configuration. The IPIF must exist before configuring the circuit. If required, the IP route must exist before configuring the circuit.

    Circuit 0

    ```
    SW37_7850_A:FID128:admin> portcfg fcipcircuit 24 modify 0 --local-ip 172.16.1.3 --remote-ip
    172.16.2.3

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

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

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

Circuit 1

```
SW37_7850_A:FID128:admin> portcfg fcipcircuit 24 modify 1 --local-ip 172.16.1.4 --remote-ip
172.16.2.4

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

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

14. Add local and remote eHCL (HA) IP addresses to the circuits.

    Circuit 0 eHCL (HA)

```
SW37_7850_A:FID128:admin> portcfg fcipcircuit 24 modify 0 --local-ha-ip 172.16.1.5 --remote-ha-ip
172.16.2.5

!!!! WARNING !!!!

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

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

    Circuit 1 eHCL (HA)

```
SW37_7850_A:FID128:admin> portcfg fcipcircuit 24 modify 1 --local-ha-ip 172.16.1.6 --remote-ha-ip
172.16.2.6

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

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

15. Add ARL minimum and maximum bandwidth values. Rates can be added in Kb/s, Mb/s, or Gb/s. The default is Kb/s. Each circuit requires setting its minimum and maximum values before coming online. ARL is activated when the maximum value is greater than the minimum value. CIR is activated when the maximum and minimum values are equal.

    Brocade FOS supports M and G CLI syntax when entering circuit minimum and maximum bandwidth rates. For example, 10G indicates 10Gb/s and 100M indicates 100Mb/s. If M or G are excluded, the default is Kb/s. For example, 5000 indicates 5Mb/s.

    The following commands set the minimum and maximum ARL values for circuits 0 and 1. Both circuits have a minimum set of 10Gb/s and a maximum of 25Gb/s. The tunnel's minimum bandwidth is 20Gb/s, and its maximum is 50Gb/s. A data processor's maximum WAN-side rate is 50Gb/s, and the minimum WAN-side rate is 50Mb/s.

Circuit 0

```
SW37_7850_A:FID128:admin> portcfg fcipcircuit 24 modify 0 --min 10G --max 25G
Operation Succeeded.
```

Circuit 1

```
SW37_7850_A:FID128:admin> portcfg fcipcircuit 24 modify 1 --min 10G --max 25G
Operation Succeeded.
```

16. If configuring IP Extension, create a Traffic Control List (TCL). If IP Extension is not being configured, skip this step.

    A TCL is required; otherwise, all LAN-side ingress IP Extension traffic is dropped. Give each TCL rule a meaningful, human-readable name. Priority rules are evaluated in the order from smallest to largest number. Leave space between each rule; counting by 100 makes it easier to change and add TCL rules, which is helpful for troubleshooting. A TCL rule must be administratively enabled; by default, rules are disabled. Allow is the default rule action and requires a target. A target is the tunnel's VE_Port number.

    In this example, the source and destination IP subnets are used to identify the flows to be sent across tunnel 24. All IP storage devices on these subnets that send data to the IP Extension Gateway will communicate across tunnel 24. Put IP storage replication ports on the same subnet to simplify configuration and troubleshooting.

    The final TCL rule is a deny all; this final rule cannot be modified or deleted. Any traffic flow failing all previous allow rules falls to the bottom and is dropped.

```
SW37_7850_A:FID128:admin> portcfg tcl MyRule1 create --priority 100 --admin-status enable --target
24 --src-addr 10.10.10.0/24 --dst-addr 192.168.0.0/24
Operation Succeeded.
```

    The following command lists the specifics of the TCL rules. TCL rule 100 is enabled and named MyRule1. It matches a source IP address from subnet 10.10.10.0/24 and a destination IP address from subnet 192.168.0.0/24. A match sends the IP storage traffic into tunnel VE24. All other matching criteria are set to ANY. At the time of the capture, the rule had been queried 1829 times. A rule is only queried once when a TCP session's three-way handshake first arrives. If allowed, the session is permitted going forward without having to query the TCL. If denied, the session is dropped going forward.

```
SW37_7850_A:FID128:admin> portshow tcl

 Pri    Name                 Flgs   Target     L2COS  VLAN DSCP Proto Port Hit
                             Src-Addr                 Dst-Addr
-------------------------------------------------------------------------------
 *100   MyRule1              AI---  24-Med     ANY    ANY  ANY  ANY   ANY  1829
                             10.10.10.0/24            192.168.0.0/24
 *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
-----------------------------------------
  DP0                  2 / 128
  DP1                  1 / 128
-----------------------------------------
Configured Total:      2 / 1024
```

17. FCIP typically does not require additional configuration other than establishing WAN-side connectivity and zoning the replication ports. Creating Virtual Fabric Logical Switches is beyond the scope of this document.

Validating Brocade 7850 Extension Configuration

This section provides the validation steps required to ensure proper operation. Refer to the Brocade Fabric OS Command Reference Manual, 9.2.0, for syntax details and parameters for each command.

The following is a list of applicable commands:

```
switchshow
sfpshow
portcfgge --show
portshow < option > [<SlotNum>/]<portNum> [<Args>]
lldp --show
lldp --show -nbr [<portNum | port-range>] [-detail]
portchannel --show [ -detail | -static | -dynamic | -all | <poName> ]
portcmd --ping
portcmd --traceroute
portcmd --wtool
```

1. Validate that the required ports are online and in the correct mode.

   The output below is from the `switchshow` command. If creating an independent replication SAN with autonomous A and B fabrics, the A fabric should have only one Principal, and the B fabric should have only one Principal. The other connected switches in the replication SAN should be Subordinate. switchDomain must be a unique number within the replication SAN. VE24 is configured and connected; it shows online with the switch it is connected to. Ge0 and ge1 are configured as LAN-side interfaces; the other GE interfaces are WAN side (WAN side is shown as FCIP). Ge16 and ge17 are the 100GE interfaces and show as online; the circuits are not shown with this command.

   The best practice is to persistently disable unused ports.

NOTE: Ensure switchState is online.

```
SW37_7850_A:FID128:admin> switchshow
switchName:     SW37_7850_A
switchType:     190.0
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:d8:1f:cc:fb:41:20
zoning:         OFF
switchBeacon:   OFF
FC Router:      OFF
HIF Mode:       OFF
Allow XISL Use: OFF
LS Attributes:  [FID: 128, Base Switch: No, Default Switch: Yes, Ficon Switch: No, Address Mode 0]
Index Port Address  Media Speed   State       Proto
==================================================
   0   0   010000   id    N64     No_Sync     FC  Disabled (Persistent) (None)
   1   1   010100   id    N64     Mod_Uns     FC  "SFP in a DD-SFP port"
   2   2   010200   id    N64     No_Light    FC
   3   3   010300   id    N64     Mod_Uns     FC  "SFP in a DD-SFP port"
   4   4   010400   id    N64     No_Light    FC
   5   5   010500   id    N64     Mod_Uns     FC  "SFP in a DD-SFP port"
   6   6   010600   id    N64     No_Light    FC
   7   7   010700   id    N64     Mod_Uns     FC  "SFP in a DD-SFP port"
   8   8   010800   id    N64     No_Light    FC
```

```
     9   9   010900   id    N64     Mod_Uns     FC  "SFP in a DD-SFP port"
    10  10   010a00   id    N64     No_Light    FC
    11  11   010b00   id    N64     Mod_Uns     FC  "SFP in a DD-SFP port"
    12  12   010c00   id    N64     No_Light    FC
    13  13   010d00   id    N64     Mod_Uns     FC  "SFP in a DD-SFP port"
    14  14   010e00   id    N64     No_Light    FC
    15  15   010f00   id    N64     Mod_Uns     FC  "SFP in a DD-SFP port"
    16  16   011000   --    N64     No_Module   FC
    17  17   011100   --    N64     No_Module   FC
    18  18   011200   --    N64     No_Module   FC
    19  19   011300   --    N64     No_Module   FC
    20  20   011400   --    N64     No_Module   FC
    21  21   011500   --    N64     No_Module   FC
    22  22   011600   --    N64     No_Module   FC
    23  23   011700   --    N64     No_Module   FC
    24  24   011800   --    --       Online      VE  VE-Port  10:00:d8:1f:cc:fb:91:e0 "SW38_7850_B"
(downstream)
    25  25   011900   --    --       Offline     VE
    26  26   011a00   --    --       Offline     VE
    27  27   011b00   --    --       Offline     VE  Disabled (6VE Mode)
    28  28   011c00   --    --       Offline     VE  Disabled (6VE Mode)
    29  29   011d00   --    --       Offline     VE  Disabled (6VE Mode)
    30  30   011e00   --    --       Offline     VE  Disabled (6VE Mode)
    31  31   011f00   --    --       Offline     VE  Disabled (6VE Mode)
    32  32   012000   --    --       Offline     VE  Disabled (6VE Mode)
    33  33   012100   --    --       Offline     VE
    34  34   012200   --    --       Offline     VE
    35  35   012300   --    --       Offline     VE
    36  36   012400   --    --       Offline     VE  Disabled (6VE Mode)
    37  37   012500   --    --       Offline     VE  Disabled (6VE Mode)
    38  38   012600   --    --       Offline     VE  Disabled (6VE Mode)
    39  39   012700   --    --       Offline     VE  Disabled (6VE Mode)
    40  40   012800   --    --       Offline     VE  Disabled (6VE Mode)
    41  41   012900   --    --       Offline     VE  Disabled (6VE Mode)
        ge0          id    10G     No_Light    LAN
        ge1          id    10G     No_Light    LAN
        ge2          id    10G     No_Light    FCIP
        ge3          id    10G     No_Light    FCIP
        ge4          id    10G     No_Light    FCIP
        ge5          id    10G     No_Light    FCIP
        ge6          id    10G     No_Light    FCIP
        ge7          id    10G     No_Light    FCIP
        ge8          --    10G     No_Module   FCIP
        ge9          --    10G     No_Module   FCIP
        ge10         --    10G     No_Module   FCIP
        ge11         --    10G     No_Module   FCIP
        ge12         --    10G     No_Module   FCIP
        ge13         --    10G     No_Module   FCIP
        ge14         --    10G     No_Module   FCIP
        ge15         --    10G     No_Module   FCIP
        ge16         id    100G    Online      FCIP
        ge17         id    100G    Online      FCIP
```

2. Validate that the SFP optics are recognized and support the desired speed. There are a variety of SFP error messages that can be shown in this output.

   The `sfpshow` command is used to gather information about inserted and recognized optics. It shows the optic's state, vendor, serial number, and supported speeds. FC optics must be Brocade-branded, and the best practice is to use Brocade-branded GE optics as well. If a regular SFP is inserted into an SFP Double Density (SFP-DD) bay, the second FC port displays Data is not available because no optic is connected.

```
SW37_7850_A:FID128:admin> sfpshow
Port  0: id (sw) Vendor: BROCADE          Serial No: MAA12229C185255S Speed: 16,32,64_Gbps
Port  1: Data is not available
Port  2: id (sw) Vendor: BROCADE          Serial No: MAA12229C185305S Speed: 16,32,64_Gbps
Port  3: Data is not available
Port  4: id (sw) Vendor: BROCADE          Serial No: MAA12229C185365S Speed: 16,32,64_Gbps
Port  5: Data is not available
Port  6: id (sw) Vendor: BROCADE          Serial No: MAA12229C108245S Speed: 16,32,64_Gbps
Port  7: Data is not available
Port  8: id (sw) Vendor: BROCADE          Serial No: MAA12229C185375S Speed: 16,32,64_Gbps
Port  9: Data is not available
Port 10: id (sw) Vendor: BROCADE          Serial No: MAA12229C185335S Speed: 16,32,64_Gbps
Port 11: Data is not available
Port 12: id (sw) Vendor: BROCADE          Serial No: MAA12229C185565S Speed: 16,32,64_Gbps
Port 13: Data is not available
Port 14: id (sw) Vendor: BROCADE          Serial No: MAA12229C185685S Speed: 16,32,64_Gbps
Port 15: Data is not available
Port 16: Media not installed
Port 17: Media not installed
Port 18: Media not installed
Port 19: Media not installed
Port 20: Media not installed
Port 21: Media not installed
Port 22: Media not installed
Port 23: Media not installed
 GE: Port  0: id (id) Vendor: BROCADE          Serial No: DAA122445010384  Speed: 10_Gbps
 GE: Port  1: id (id) Vendor: BROCADE          Serial No: DAA122445010174  Speed: 10_Gbps
 GE: Port  2: id (id) Vendor: BROCADE          Serial No: DAA122445010504  Speed: 10_Gbps
 GE: Port  3: id (id) Vendor: BROCADE          Serial No: DAA122445010474  Speed: 10_Gbps
 GE: Port  4: id (sw) Vendor: BROCADE          Serial No: CAA423021001011  Speed: 25_Gbps
 GE: Port  5: id (sw) Vendor: BROCADE          Serial No: CAA423021000541  Speed: 25_Gbps
 GE: Port  6: id (sw) Vendor: BROCADE          Serial No: CAA423021001041  Speed: 25_Gbps
 GE: Port  7: id (sw) Vendor: BROCADE          Serial No: CAA423021000861  Speed: 25_Gbps
 GE: Port  8: --
 GE: Port  9: --
 GE: Port 10: --
 GE: Port 11: --
 GE: Port 12: --
 GE: Port 13: --
 GE: Port 14: --
 GE: Port 15: --
 GE: Port 16: id (sw) Vendor: BROCADE          Serial No: YTA42308PH00071  Speed: 100_Gbps
 GE: Port 17: id (sw) Vendor: BROCADE          Serial No: YTA42308PH00063  Speed: 100_Gbps
```

3. Validate the GE settings, including speed, auto-negotiation, media type, LAN or WAN side, and if the GE interface is a portchannel (LAG) member.

   The `portcfgge --show` command output shows the following: GE interface number, interface speed setting, interface LAN or WAN mode (WAN is default and has no flag), FEC setting, and the LAG Name.

```
SW37_7850_A:FID128:admin> portcfgge --show
 Port          Speed     Flags     Channel  FEC   Lag Name
-------------------------------------------------------------------------
 ge0           10G       --LG      N/A      Off   MyPO
 ge1           10G       --LG      N/A      Off   MyPO
 ge2           10G       ----      N/A      Off   -
 ge3           10G       ----      N/A      Off   -
 ge4           10G       ----      N/A      Off   -
 ge5           10G       ----      N/A      Off   -
 ge6           10G       ----      N/A      Off   -
 ge7           10G       ----      N/A      Off   -
 ge8           10G       ----      N/A      Off   -
 ge9           10G       ----      N/A      Off   -
 ge10          10G       ----      N/A      Off   -
 ge11          10G       ----      N/A      Off   -
 ge12          10G       ----      N/A      Off   -
 ge13          10G       ----      N/A      Off   -
 ge14          10G       ----      N/A      Off   -
 ge15          10G       ----      N/A      Off   -
 ge16          100G      ----      N/A      CL91  -
 ge17          100G      ----      N/A      CL91  -
-------------------------------------------------------------------------
Flags: A:Auto-Negotiation Enabled  C:Copper Media Type
L:LAN Port G: LAG Member
```

4. Validate the GE interface status: state, speed, and MAC address.

   GE0, which is set to the LAN side (Offline)

```
SW37_7850_A:FID128:admin> portshow ge0
Eth Mac Address: d8.1f.cc.fb.41.21
Port State: 2   Offline
Port Phys:  4   No_Light
Port Flags: 0x1  PRESENT
Port Speed: 10G
```

   GE16, which is set to the WAN side (Online)

```
SW37_7850_A:FID128:admin> portshow ge16
Eth Mac Address: d8.1f.cc.fb.41.31
Port State: 1   Online
Port Phys:  6   In_Sync
Port Flags: 0x4003       PRESENT ACTIVE LED
Port Speed: 100G
```

5. Validate that Link Layer Discovery Protocol (LLDP) is operational. If LLDP is not working, it will not prevent extension functionality. LLDP is primarily used as a verification and troubleshooting tool; the connected network switches must support LLDP and have it enabled.

   The LLDP output below shows how LLDP has been configured. LLDP is enabled by default; it must be enabled on both ends of the Ethernet link (extension platform to data center LAN switch).

```
LLDP Global Information
  system-name: SW37_7850_A
  system-description: Brocade_7850_Fabric_OS_Version_9_2_0
  State:                    Enabled
  Mode:                     Receive/Transmit
  Advertise Transmitted:    30 seconds
  Hold time for advertise: 120 seconds
  Tx Delay Timer:           1 seconds
  Transmit TLVs:            Chassis ID              Port ID
                            TTL                     Port Description
                            System Name             System Description
                            System Capabilities     Management Address
```

The LLDP output below helps validate and troubleshoot connected GE interfaces. If a GE interface appears in the LLDP list, it is communicating with the connected data center switch. LLDP only communicates at the Ethernet link level, not through routers. There are timers, and the entry will be removed upon expiration after a link goes offline. The connected system name with the local and remote interfaces is shown. Verify the cabling was connected to the proper data center LAN switch and interface. LLDP works on WAN, LAN, and Mgmt ports.

```
SW37_7850_A:FID128:admin> lldp --show -nbr
Local Intf  Dead Interval  Remaining Life  Remote Intf     Chassis ID      Tx    Rx      System Name
ge16        120            95              ge16            0acd.1fd8.0000 2507  95077   SW38_7850_B
ge17        120            103             ge17            0acd.1fd8.0000 2507  92157   SW38_7850_B
```

6. Validate the LAN-side portchannel (LAG). If IP Extension is being configured, using a portchannel is the best practice. If IP Extension is not being deployed, skip this step. The portchannel name is shown, and the type is dynamic. A dynamic portchannel uses Link Aggregation Control Protocol (LACP) to form the LAG. An operational portchannel will have a state of Online. The portchannel has two links; there are two member GE interfaces.

```
SW37_7850_A:FID128:admin> portchannel --show
Name                            Type        Oper-State   Port-Count   Member Ports
--------------------------------------------------------------------------------
MyPO                            Dynamic     Offline      2            ge0 ,ge1
```

Use the following command to gather portchannel details:

```
SW37_7850_A:FID128:admin> portchannel --show -detail
 Name: MyPO
 Type: Dynamic
 Key: 785
 Speed: 10G
 Admin-state: Enable
 Oper-state: Offline
  LACP System Priority: 32768
  LACP System MAC: d8:1f:cc:fb:41:3a
  LACP PARTNER System Priority: 65535
  LACP PARTNER System MAC: 00:00:00:00:00:00
  Portchannel Member count: 2
  Port          Oper state      Sync    Timeout         Auto-Negotiation
  ------------------------------------------------------------------------
   ge0          Offline         0       Long            Disabled
   ge1          Offline         0       Long            Disabled
```

7. Validate that the IP interfaces (IPIF) have been created correctly.

   The output of the `portshow ipif` command shows the DP the GE interface is associated with, the WAN-side (circuit endpoint) IPIFs, and the LAN-side (IP Extension gateway) IPIFs. Additionally, the command shows the MTU and VLAN ID. A VLAN ID of zero indicates that no VLAN was set, and VLAN tagging for this IPIF is not in use. The output also shows if IPIFs are Up, Running, and InUse.

```
SW37_7850_A:FID128:admin> portshow ipif

 Port         IP Address                    / Pfx  MTU   VLAN  Flags
-------------------------------------------------------------------------------
 ge16.dp0     172.16.1.3                    / 24   1500  0     U R M I
 ge16.dp1     172.16.1.5                    / 24   1500  0     U R M I
 ge17.dp0     172.16.1.4                    / 24   1500  0     U R M I
 ge17.dp1     172.16.1.6                    / 24   1500  0     U R M I
 lan.dp0      192.168.0.4                   / 24   1500  0     U R M
 lan.dp0      10.0.0.4                      / 29   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
```

8. Validate the WAN-side IP Routes.

   The `portcfgshow iproute` command shows the DP, and the GE interface the route was assigned. Routes are specific to the GE interface/DP pairs. A route that indicates a GE interface is a WAN-side route. A route indicating LAN is a LAN-side route. The destination subnet, mask, and local gateway are shown.

```
SW37_7850_A:FID128:admin> portcfgshow iproute

 Port         IP Address                / Pfx  Gateway                Flags
-------------------------------------------------------------------------------
 ge16.dp0     172.16.2.0                / 24   172.16.1.1
 ge16.dp1     172.16.2.0                / 24   172.16.1.1
 ge17.dp0     172.16.2.0                / 24   172.16.1.1
 ge17.dp1     172.16.2.0                / 24   172.16.1.1
 lan.dp0      10.10.10.0                / 24   10.0.0.1
-------------------------------------------------------------------------------
 Flags: S=Static X=Crossport
```

9. Validate that the IPsec-Policy was created.

   The `portshow ipsec-policy` command shows IPsec policies configured on the platform. The IPsec policy name and flags are displayed. PKI and PSK are supported. Use the -p argument to show the PSK; otherwise, leave it off. IPsec encryption uses AES 256 and SHA512.

```
SW37_7850_A:FID128:admin> portshow ipsec-policy -p

 IPSec Policy    Flg Authentication data
-------------------------------------------------------------------------------
 MyIPsec         S-- w7ffuffx9zvgt6wrr3pka2mzd9o6bz30vwb0y3u67p3taw5wjxmwp9t7brwau98v
-------------------------------------------------------------------------------
Flags: *=Name Truncated. Use "portshow ipsec-policy -d for details."
       P=PKI Profile  S=Shared-Key Profile
       X=Expired Cert  M=Hash Mismatch
```

10. Validate the state of the tunnel and circuits and that they were created correctly.

    The output from the `portshow fciptunnel -cs` command shows the VE_Port number (tunnel), the circuits and the local IPIF GE interface in use, operational status, flags, uptime, TX and Rx in Mbps, connection count, min and max rate settings, circuit metric, and failover group. Note: tunnel and circuit flags are different; reference the proper section of the legend below the output.

```
SW37_7850_A:FID128:admin> portshow fciptunnel -c --summary

 Tunnel Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps ConnCnt CommRt Met/G
--------------------------------------------------------------------------------
 24    -         Up       --i----PI    21h4m    0.00    0.00   1      -       -
 24    0 ge16    Up       ----ah-i4    21h4m    0.00    0.00   1 10000/25000 0/0
 24    1 ge17    Up       ----ah-i4    21h4m    0.00    0.00   1 10000/25000 0/0
--------------------------------------------------------------------------------
 Flags (tunnel): l=Legacy QOS Mode
                 i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
                 a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                 I=IP-Ext
      (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4 6=IPv6
                 ARL a=Auto r=Reset s=StepDown t=TimedStepDown  S=SLA
```

    For tunnel details, use the following command:

```
SW37_7850_A:FID128:admin> portshow fciptunnel 24

 Tunnel: VE-Port:24 (idx:0, DP0)
 ================================================
  Oper State           : Online
  TID                  : 24
  Flags                : 0x00000000
  IP-Extension         : Enabled
  Compression          : None
  FC-Compression       : Fast Deflate (Override)
  IP-Compression       : Deflate (Override)
  QoS Distribution     : Protocol (FC:50% / IP:50%)
  FC QoS BW Ratio      : 50% / 30% / 20%
  IP QoS BW Ratio      : 50% / 30% / 20%
  Fastwrite            : Disabled
  Tape Pipelining      : Disabled
  IPSec                : Enabled
  IPSec-Policy         : MyIPsec
  Legacy QOS Mode      : Disabled
  Load-Level (Cfg/Peer): Failover (Failover / Failover)
  Local WWN            : 10:00:d8:1f:cc:fb:41:20
  Peer WWN             : 10:00:d8:1f:cc:fb:91:e0
  RemWWN (config)      : 00:00:00:00:00:00:00:00
  Peer Platform        : 7850
  cfgmask              : 0x4001024c 0x00c0001f
  Uncomp/Comp Bytes    : 0 / 0 / 1.00 : 1
  Uncomp/Comp Byte(30s): 0 / 0 / 1.00 : 1
  Flow Status          : 0
  ConCount/Duration    : 1 / 1d22h41m
  Uptime               : 21h12m
  Stats Duration       : 21h12m
  Receiver Stats       : 5260620 bytes / 19727 pkts /   58.00 Bps Avg
  Sender Stats         : 5502212 bytes / 19730 pkts /   59.00 Bps Avg
```

```
  TCP Bytes In/Out     : 6261382048 / 5437178684
  ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
  RTT (min/avg/max)    : 1 / 1 / 27 ms
  Wan Util             : 0.0%
  TxQ Util             : 0.0%
```

11. Validate the tunnel and circuit performance and compression ratio.

    Using the `--perf` argument, the output shows the compression ratio, Round Trip Time (RTT), and number of retransmits (ReTx). Only transmit compression is shown, not received data.

```
SW37_7850_A:FID128:admin> portshow fciptunnel -c --perf

Tunnel Circuit St Flg TxMBps RxMBps CmpRtio RTTms  ReTx  TxWAN%  TxQ%/BW  Met/G
-------------------------------------------------------------------------------
 24    -       Up -I-    0.0    0.0   1.0:1   -       0     0        0       -
 24    0 ge16  Up ---    0.0    0.0       -   1       0     0   10000/25000 0/-
 24    1 ge17  Up ---    0.0    0.0       -   1       0     0   10000/25000 0/-
-------------------------------------------------------------------------------
   Flg (tunnel): 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)
```

12. If configuring IP Extension, validate the TCL; otherwise, skip this step.

    Observe the following output: An enabled rule shows an `*` before the priority number. The last TCL rule is 65535 and cannot be deleted or modified; it is a deny-all rule that drops traffic that did not match a previous allow rule.

    Verify that any IP addresses, subnets, mask lengths, VLAN, QoS, and protocol ports a rule uses are correct. When the application sends data to the IP Extension gateway, does the anticipated rule's hit count? If not, either the traffic is not getting to the IP Extension gateway or the rule is not matching the traffic.

```
SW37_7850_A:FID128:admin> portshow tcl

 Pri    Name                 Flgs   Target     L2COS  VLAN DSCP Proto Port Hit
                             Src-Addr                 Dst-Addr
-------------------------------------------------------------------------------
 *100   MyRule1              AI---  24-Med     ANY    ANY  ANY  ANY   ANY  0
                             10.10.10.0/24            192.168.0.0/24
 *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
-----------------------------------------
  DP0                  2 / 128
  DP1                  1 / 128
-----------------------------------------
Configured Total:      2 / 1024
```

13. Validate basic connectivity on the WAN side.

    You can ping into the WAN side from a WAN-side GE interface. Use the `portcmd --ping <ge#.dp#>` command to ping. See the command output below.

    The destination IP address does not need to be another extension GE interface. It is not possible to ping IP addresses within the same extension platform. Note: this command does not apply to the management port.

```
SW37_7850_A:FID128:admin> portcmd --ping ge16.dp0 -s 172.16.1.3 -d 172.16.1.11

PING 172.16.1.11 (172.16.1.3) with 64 bytes of data.
64 bytes from 172.16.1.11: icmp_seq=1 ttl=20 time=1 ms
64 bytes from 172.16.1.11: icmp_seq=2 ttl=20 time=1 ms
64 bytes from 172.16.1.11: icmp_seq=3 ttl=20 time=1 ms
64 bytes from 172.16.1.11: icmp_seq=4 ttl=20 time=1 ms

--- 172.16.1.11 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 21 ms
rtt min/avg/max = 1/1/1 ms
```

14. If you are configuring IP Extension, validate basic connectivity on the LAN side. Validating connectivity on the LAN side is useful when you are troubleshooting why a TCL hit count is not increasing. Start by pinging from the IP Extension gateway to the local router gateway, as shown in the following command. They must both be on the same subnet.

    You can send pings between IP storage end devices through IP Extension to see if the hit count increases.

NOTE: A TCL rule is needed to match an ICMP echo, and the ping header needs the proper source and destination IP addresses for the TCL rule to match.

```
SW37_7850_A:FID128:admin> portcmd --ping lan.dp0 -s 10.0.0.4 -d 10.0.0.1

PING 10.0.0.1 (10.0.0.4) with 64 bytes of data.
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1 ms
64 bytes from 10.0.0.1: icmp_seq=3 ttl=64 time=1 ms
64 bytes from 10.0.0.1: icmp_seq=4 ttl=64 time=1 ms

--- 10.0.0.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 11 ms
rtt min/avg/max = 1/1/1 ms
```

15. If you are implementing IP Extension, validate the flow stats.

    There are various ways to view flow information from IP Extension. The output below shows each TCP flow that IP Extension manages and optimizes. It shows the source and destination IP addresses and ports, the protocol, transmission, and receive rates in bytes per second.

    At the bottom of the output, active TCP sessions for each DP are shown. Also, the number of attempted TCP sessions that exceeded the DP's maximum is shown.

```
SW37_7850_A:FID128:admin> portshow lan-stats --per-flow

 *** Displaying Top 17 connections by throughput ***

 DP     Idx  Src-Address      Dst-Address      Sport Dport Pro   Tx(B/s) Rx(B/s)
-------------------------------------------------------------------------------
 DP0    682  10.150.20.14     10.150.25.15     54994 64482 TCP   1.3m    1.3m
```

```
 DP0     684  10.150.20.14     10.150.25.15     54996 64482 TCP   1.3m    1.3m
 DP0     683  10.150.20.14     10.150.25.15     54995 64482 TCP   1.3m    1.3m
 DP0     686  10.150.20.14     10.150.25.15     54998 64482 TCP   1.3m    1.3m
 DP0     685  10.150.20.14     10.150.25.15     54997 64482 TCP   1.3m    1.3m
 DP0     697  10.150.20.14     10.150.25.15     55009 64482 TCP   1.3m    1.3m
 DP0     696  10.150.20.14     10.150.25.15     55008 64482 TCP   1.3m    1.3m
 DP0     690  10.150.20.14     10.150.25.15     55002 64482 TCP   1.0m    1.0m
 DP0     693  10.150.20.14     10.150.25.15     55005 64482 TCP   1.0m    1.0m
 DP0     687  10.150.20.14     10.150.25.15     54999 64482 TCP   1.0m    1.0m
 DP0     691  10.150.20.14     10.150.25.15     55003 64482 TCP   1.0m    1.0m
 DP0     692  10.150.20.14     10.150.25.15     55004 64482 TCP   1.0m    1.0m
 DP0     694  10.150.20.14     10.150.25.15     55006 64482 TCP   1.0m    1.0m
 DP0     688  10.150.20.14     10.150.25.15     55000 64482 TCP   1.0m    1.0m
 DP0     689  10.150.20.14     10.150.25.15     55001 64482 TCP   1.0m    1.0m
 DP0     695  10.150.20.14     10.150.25.15     55007 64482 TCP   1.0m    1.0m
 DP0     681  10.150.25.15     10.150.20.14     64481 54991 TCP   0       0
--------------------------------------------------------------------------------
 Sport=Source-Port Dport=Destination-Port Pro=Protocol

 DP       ActTCP  ExdTCP  TCLDeny TCLFail
----------------------------------------------------------
 DP0      17      0       0       0
 DP1      0       0       0       0
----------------------------------------------------------
 ActTCP=Active TCP Conns  ExdTCP=Exceeded TCP Conn Cnt
```

Suppose multiple IP storage end device ports communicate across IP Extension, and visibility into a specific pair of IP addresses is desired. In that case, the following command shows the proper output for the transmit and receive rates for active sessions.

```
SW37_7850_A:FID128:admin> portshow lan-stats --ip-pair

 DP       Idx SrcAddr             DstAddr             Active  TxB       RxB
--------------------------------------------------------------------------------
 DP0      0   10.150.25.15        10.34.196.159       0       0         0
 DP0      1   10.75.16.12         10.150.20.14        0       0         0
 DP0      2   192.19.189.10       10.150.20.14        0       0         0
 DP0      3   10.34.112.19        10.150.20.14        0       0         0
 DP0      4   10.150.20.1         10.150.20.14        0       0         0
 DP0      5   20.228.85.55        10.150.20.14        0       0         0
 DP0      6   52.137.108.250      10.150.20.14        0       0         0
 DP0      7   104.91.122.87       10.150.20.14        0       0         0
 DP0      8   10.34.176.127       10.150.20.14        0       0         0
 DP0      9   40.119.249.228      10.150.20.14        0       0         0
 DP0      10  52.191.219.104      10.150.20.14        0       0         0
 DP0      11  52.248.96.54        10.150.20.14        0       0         0
 DP0      12  10.34.112.21        10.150.20.14        0       0         0
 DP0      13  10.150.25.15        10.150.20.14        17      10.0g     10.0g
 DP0      14  10.34.176.126       10.150.20.14        0       0         0
--------------------------------------------------------------------------------
```

Summary

As the digital world explodes with new workloads and the risk of losing your most critical data becomes more intense, securing your data with the most current and performant infrastructure becomes necessary to establish business sustainability. Brocade Extension platforms are relevant to all-flash arrays and new technologies like NVMe over Fabrics (Fibre Channel). Lockstep your replication network with server and storage advances to ensure investment in those endpoints returns the expected performance.

In addition, it is crucial to understand the need for product life-cycle management and the eventual migration to new and supported technology. Remote data replication is the gold standard for protecting workloads. You do not want to explain to your CIO why the replication infrastructure is without support, incapable of functionality, or not interoperable.

Migration to a new platform can be stressful, but with the methodology outlined in this paper, you can move with little to no impact on day-to-day operations.

Additional Resources

To access the Brocade Fabric OS Extension User Guide, go to: techdocs.broadcom.com/us/en/fibre-channel-networking/fabric-os/fabric-os-extension/9-2-x.html

To access product details about the Brocade 7850 Extension Switch, go to: www.broadcom.com/products/fibre-channel-networking/extension/7850-extension-switch

For more information on Brocade SAN health, go to: www.broadcom.com/products/fibre-channel-networking/software/sanhealth