BROCADE

BROCADE VALIDATED DESIGN

Brocade Ruckus Campus Fabric

53-1005016-01
17 April 2017

Contents

Introduction

For the past two decades, campus network design has seen little innovation, as market leaders have promoted the status quo with no significant investments beyond incremental "speeds and feeds" enhancements. Years of incremental technology upgrades have turned traditional legacy campus networks into a complex and fragmented patchwork of network devices. Campus networking has reached a point where traditional network architectures are struggling to keep up with users' relentless demands for seamless mobility across the campus and pervasive access to latest-generation applications. To ease all these pains, Brocade came up with the concept of Campus Fabric in the Enterprise. The fabric should have a single point of management, and be self-healing and self-forming. All these features are at the core of the Ruckus Campus Fabric.

Ruckus Campus Fabric technology integrates high-performance, fixed form factor switches to create a single distributed logical switch that is independent of physical location and that allows organizations to add ports whenever and wherever needed across the campus without adding complexity. Switch port extender (abbreviated as SPX in this document) is the engineering name of the Ruckus Campus Fabric solution.

This document describes network designs using the Ruckus Campus Fabric (or SPX) architecture. The configurations and design practices documented here are fully validated and conform to the Ruckus Campus Fabric architecture. The intention of this Brocade validated design document is to provide reference configurations and document best practices for building a scalable Ruckus Campus Fabric using Brocade switches.

It is highly recommended to review the SPX deployment considerations described in the Brocade FastIron Switch Port Extender Deployment Guide for a detailed discussion on various supported and unsupported topologies while designing the Ruckus Campus Fabric.

Brocade Validated Designs

Helping customers consider, select, and deploy network solutions for current and planned needs is our mission. Brocade validated designs offer a fast track to success by accelerating that process.

Validated designs are repeatable reference network architectures that have been engineered and tested to address specific use cases and deployment scenarios. They document systematic steps and best practices that help administrators, architects, and engineers plan, design, and deploy physical and virtual network technologies. Leveraging these validated network architectures accelerates deployment speed, increases reliability and predictability, and reduces risk.

Brocade validated designs incorporate network and security principles and technologies across the ecosystem of service provider, data center, campus, and wireless networks. Each Brocade validated design provides a standardized network architecture for a specific use case, incorporating technologies and feature sets across Brocade products and partner offerings.

All Brocade validated designs follow best-practice recommendations and allow for customer-specific network architecture variations that deliver additional benefits. The variations are documented and supported to provide ongoing value, and all Brocade validated designs are continuously maintained to ensure that every design remains supported as new products and software versions are introduced.

By accelerating time-to-value, reducing risk, and offering the freedom to incorporate creative, supported variations, these validated network architectures provide a tremendous value-add for building and growing a flexible network infrastructure.

Scope

This Brocade validated design provides guidance for designing and implementing Ruckus Campus Fabric technology using Brocade hardware and software. It details the Brocade reference architecture for deploying an IEEE 802.1BR open standard-based solution.

It should be noted that not all features, such as automation practices and monitoring of the Ruckus Campus Fabric, are included in this document.

The design practices documented here follow the best-practice recommendations, but there are variations to the design that are supported as well.

Target Audience

This document is written for Brocade systems engineers, partners, and customers who design, implement, and support the Ruckus Campus Fabric. This document is intended for experienced data communication architects and engineers. It assumes that the reader has a good understanding of basic networking terminologies and design aspects.

About the Author

Amit Mittal is a Staff Engineer on the IP SQA team at Brocade. Amit has extensive experience testing campus fabric and IP routing technologies. He has been involved in validating solution architectures.

The author would like to acknowledge the following Brocadians for their technical guidance in developing this validated design:

- **Abdul Khader**—Technical Director
- **Ashish Sinha**—Senior Software Test Engineer
- **Krish Padmanabhan**—Principal Engineer
- **Sumit Lakhotia**—Senior Manager
- **Vandana Chander**—Software Engineer
- **Vignesh Hariharan**—Senior Product Manager

Document History

Date	Part Number	Description
April, 2017	53-1005016-01	Initial release.

About Brocade

Brocade® (NASDAQ: BRCD) networking solutions help the world's leading organizations transition smoothly to a world where applications and information reside anywhere. This vision is designed to deliver key business benefits such as unmatched simplicity, non-stop networking, application optimization, and investment protection.

Innovative Ethernet and storage networking solutions for data center, campus, and service provider networks help reduce complexity and cost while enabling virtualization and cloud computing to increase business agility.

To help ensure a complete solution, Brocade partners with world-class IT companies and provides comprehensive education, support, and professional services offerings (www.brocade.com).

Terminology

Term	Meaning
ARP	Address Resolution Protocol
BGP	Border Gateway Protocol
BUM	Broadcast, Unknown Unicast, and Multicast
CB	Control Bridge
CSP	Control and Status Protocol
E-CID	E-Channel Identifier
ECMP	Equal-Cost Multi-Path
IP	Internet Protocol
LLDP	Link Level Discovery Protocol
MAC	Media Access Control
ME-CID	Multicast E-Channel Identifier
PE	Port Extender
SPX	Switch Port Extender
TLV	Type-Length-Value
ToR	Top of Rack switch. Also leaf or VxLAN Tunnel Endpoint (VTEP) in an IP fabric context.
UDP	User Datagram Protocol
VLAN	Virtual Local Area Network
VM	Virtual Machine

Definitions

Control bridge	The control bridge (CB) has ports that link to one or more port extender (PE) units. The CB handles traffic to and from PE ports as if these ports residing on the control bridge itself were local ports. The control bridge can consist of a standalone ICX 7750 or a traditional stack of ICX 7750 stacking units.
Access (or Base) port extender	A port extender (PE) connecting to end hosts. An access PE is located at the end of the CB-PE tree.
Cascade port	An egress CB or PE port connecting to a downstream PE. (The link between the cascade port and the PE port is configured and displayed as an SPX port or SPX LAG.)
CSP	Control and Status Protocol. CSP serves as the SPX communication protocol between the CB master and the PE units in its control plane.
E-channel	The bidirectional path between the external extended port and the corresponding internal extended port. E-channels are identified by E-CID. The ID (E-CID) in the E-tag ranges from 0x0001 to 0x3FFF. E-channels can be point-to-point, point-to-multipoint, or multipoint-to-point.
E-CID	E-channel identifier. An E-CID identifies the PE destination port in CSP.
E-tag	A tag added to SPX packets. The E-tag contains an E-CID.
Extended port	A PE port that serves as an access port to a host. Internal extended ports provide connectivity to the ports of the Customer VLAN ID (C-VLAN) component. Extended ports operate as ports of the extended bridge. Each internal extended port is linked through an E-channel to an external extended port. Additional E-channels provide linkage between an internal extended port and multiple external extended ports in support of multicast frame delivery.
LLDP	Link Level Discovery Protocol. LLDP is the Layer 2 protocol used by the CB to discover and connect to PEs.
Master (Active Controller)	The master of a CB stack. The entire stack is considered one logical device.
ME-CID	Multicast E-channel Identifier. Inserted in the E-tag, it carries a value of 0x1000 to 0x3FFF to indicate a multicast channel.
PE	Port extender. A PE, or PE unit, is a dummy device that contains multiple ports (for example, 24 or 48 ports). A PE forwards all traffic to the connected control bridge (CB). The PE does not perform local switching. The PE can connect to any unit in the CB.
PE mode	A special bootup role that causes a unit to perform as a dummy device. A unit in PE mode does not parse startup configuration flash memory during bootup. It runs protocols such as LLDP only. In PE mode, the unit does not perform local switching. Most commands and configurations are blocked.
Provisional-PE mode	A temporary mode created when a user configures the**spx pe-enable** command on a unit but has not yet reloaded the unit. Because the unit previously booted up in regular mode, it continues to perform as a regular device until the next reload. That is, the unit still acts like a regular switch or router. Most commands and configurations are blocked the same way as in PE mode.
Reserved configuration	An SPX unit or stack unit that is configured and appears in output for the **show running-config** command, but which does not physically exist.
SPX	Switch Port Extender. SPX capability is based on IEEE 802.1BR standards and recommendations.
SPX port	A port on a PE that links to the CB or another PE. SPX port is a general term for PE ports and cascade ports. The **spx-port** command is used in CLI configuration to configure either a cascade port or an SPX port.
SPX LAG	A trunk that contains at least two SPX ports. (Cascade ports may be included in an SPX LAG.)
Standalone	An individual unit that is not part of a stacking system or SPX infrastructure.
TLV	Type-Length-Value information. CSP packets consist of a command TLV and can contain additional TLVs.
Transit PE	A PE that is not an access PE. A transit PE aggregates transmissions from downstream PEs.
Upstream port	A PE port that connects to a transit PE toward the CB or directly to the CB. An upstream port is an "ingress" port from the perspective of the CB.

Technology Overview

A Radically Different Approach to Campus Networking

Campus fabric collapses multiple network layers into a single logical device, combining the power of a "distributed chassis" design with the flexibility and cost-effectiveness of fixed form factor switch building blocks.

For the customers not requiring very high core support, all three layers—core, aggregation, and access—can be collapsed into a single SPX system.

Architecture for Centralized Control

The traditional aggregation and core layers are replaced by a stack of high-performance 10 Gigabit Ethernet (GbE)/40 GbE fixed form factor switches that are connected together through a high-speed "campus ring" that can span up to 10 kilometers (km). These switches are the control bridge (CB) devices. Together they deliver a unified network control plane that acts as the central management and traffic forwarding authority for the entire campus fabric domain. All of the image upgrade, configuration, and management of the access and aggregation switches is handled from a single touchpoint.

Combining Access and Aggregation into One Logical Switch

At the edge of the network, the access layer is replaced by a set of Port Extender (PE) devices connected directly or indirectly to the stack of CB devices. Using the "distributed chassis" metaphor, you can consider these PE devices to be "virtual line cards." They are transparently managed and controlled by the CB, eliminating the need to manually provision and configure individual edge switches. In fact, the entire campus fabric domain is managed as one logical device from a single point of management within the CB. The CB and PE devices use the standard-based 802.1BR protocol to communicate between themselves. An ICX 7750 is used as the CB and the PEs can be ICX 7450 or ICX 7250 switches.

FIGURE 1 SPX Domain as a Single Logical Switch

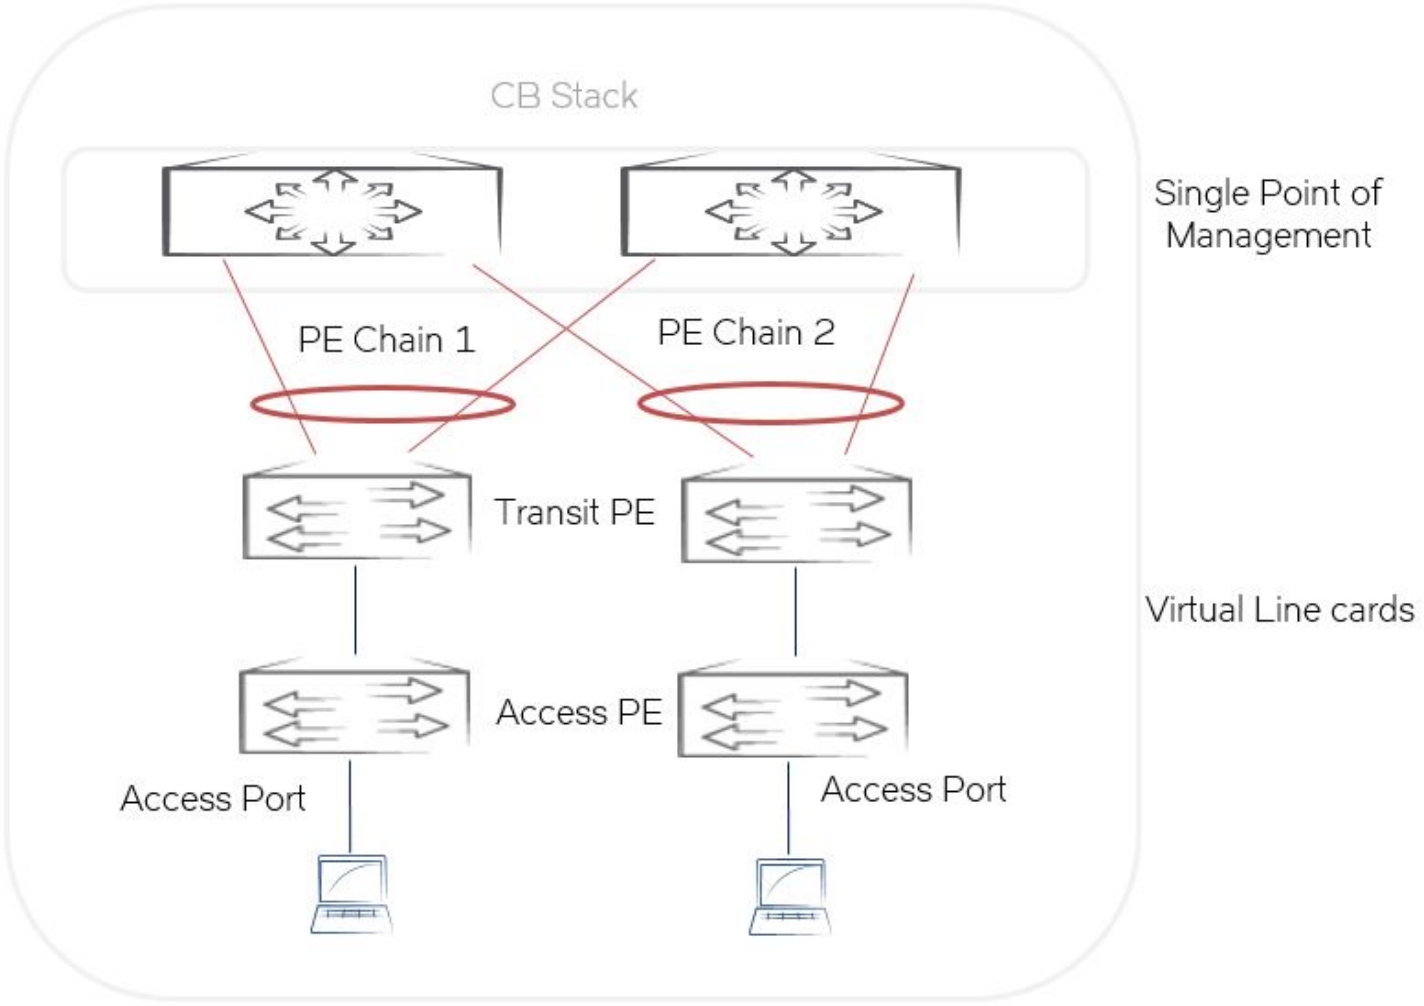

Advantages of Campus Fabric

Elimination of STP Inefficiency

The entire domain runs from a unified control and forwarding plane, eliminating the need to deploy a loop-avoidance protocol, such as Spanning Tree Protocol (STP), within the fabric domain, or complex Layer 3 protocols, such as Open Shortest Path First (OSPF). Multi-pathing is supported by design within a campus fabric domain. All links between switches are active at all times, and traffic is load balanced, optimizing performance while delivering fast failover recovery from link failure with no impact on network service.

Seamless Mobility

The Ruckus Campus Fabric architecture flattens the network, eliminating arbitrary Layer 3 boundaries between physical locations. The Ruckus Campus Fabric provides one large distributed Layer 2 switch which simplifies the deployment of wireless access points and

delivers a better user experience. For example, seamless roaming is enabled between Wi-Fi access points across the campus. Users can enjoy uninterrupted access to their favorite applications, including voice and streaming video services, while they roam across the campus with their portable devices.

Inheriting Features and Services

All devices within a Ruckus Campus Fabric domain offer the same level of network services and software features because they are all part of the same logical switch. All advanced services running at the aggregation layer, such as premium Layer 3 features, are inherited by the network edge ports. Also, a feature exclusive to the control bridge, such as VxLAN, can be inherited by the PE switches, which could not support the feature natively. Additionally, software images running on the various devices are automatically updated and kept in sync, so that there is no risk of version mismatch between the various devices, thus eliminating potential network downtime.

"Pay as You Grow" Design

The Brocade fixed form factor-based design enables cost-effective scale-out networking. It adds PE devices when more ports are needed at the edge and adds CB devices when more ports are needed at the aggregation layer. Unlike traditional chassis-based aggregation switches, no excess idle capacity is required, and no "fork-lift" upgrade is needed to advance to the next capacity level.

Gradual Migration

The Ruckus Campus Fabric interoperates with traditional networks, so there is no need to migrate the whole network at once. Brocade ICX 7750 switches can act simultaneously as a CB on ports that are connected to PE devices and as a regular Layer 2/Layer 3 aggregation switch on ports that are connected to regular access switches, enabling a gradual migration strategy.

Zero Touch Provisioning (ZTP)

In FastIron 08.0.50, Zero Touch Provisioning (ZTP) was added, which requires minimal configuration on the CB. Once fabric ports on the CB are called out, new access switches connected to these ports are image upgraded, rebooted into PE mode, and LAGs are formed on the uplinks automatically. For more complex topologies that involve user interaction, there is an "Interactive Setup" mode that the user can leverage. For the CLI savvy, manual CLI setup is available as well.

Advantages of Campus Fabric Over Stacking

- Campus Fabric scales much higher than a stack. A stack is 12 units while Campus Fabric is 36 PEs plus 4 CBs.
- Campus Fabric allows stacking between different switches, such as the ICX 7750, ICX 7450, and ICX 7250 (and other future switches).
- STP is not required between access and aggregation switches when Campus Fabric is deployed.
- A premium license is needed only on the CB stack and all PE switches can use it. Not having to purchase a premium license for every PE switch is a cost savings for the customer.
- Any user-facing ports or uplink ports can be used to form SPX LAGs versus dedicated ports in stacking.
- Zero Touch Provisioning (ZTP) gives automatic image-upgrade capabilities in Campus Fabric. ZTP makes the provisioning more like plug and play.
- A feature supported only on the premium platforms, such as the ICX 7750, can be leveraged by all PEs, even when PE hardware is incapable of performing the feature. For example, a VxLAN, when supported on an ICX 7750, can be leveraged by

all PEs even when the hardware is incapable of supporting VxLAN. Also, the ICX 7150, which may be supported in PE mode in a future release, will be able to leverage features such as GRE, ERSPAN, and VRF available on the ICX 7750.

Functional Components of Ruckus Campus Fabric

The components of Ruckus Campus Fabric are the Brocade ICX 7450 and ICX 7250 switches configured as port extenders (PEs), or PE units, to a set of Brocade ICX 7750 stack units configured as the 802.1BR control bridge. The ICX 7750 control bridge (CB) provides a single point of management for the extended network. Active and standby controller functions are retained in the ICX 7750 (control bridge) stack and continue to provide hitless recovery as well as extended administrative functions. Ruckus Campus Fabric widely increases the number of access devices in the network that can be controlled and managed from a single point. The distributed CB at the center of the SPX architecture manages PE units and hundreds of ports at the network edge.

The following devices from the Brocade FastIron product family support Switch Port Extender (SPX) configurations in FastIron 08.0.50 and later releases:

- Brocade ICX 7750 switches (ICX 7750)—Control Bridge
- Brocade ICX 7450 switches (ICX 7450)—Port Extenders
- Brocade ICX 7250 switches (ICX 7250)—Port Extenders

Campus Fabric Infrastructure Links

All CB units are stacked together to form a functional switch stack. These are interconnected with WAN-edge devices using Layer 3 interfaces. In the validated design:

- 40-GbE links are used as stack ports between the CB units.
- 40-GbE/10-GbE links are used in a LAG configuration to connect the PE chains with the CB stack as a SPX LAG.
- 40-GbE/10-GbE links are used in a LAG configuration (wherever possible) as PE upstream and downstream links.

Server-Facing Links

The server-facing or access links are on the CB devices. In the validated design:

- 10-GbE links are used for server-facing VLANs.
- The links are configured as Layer 3 virtual interfaces with associated VLANs.
- The MTU for these links is set to 1500 bytes (the default).

End Host-Facing Links

The end host-facing or access links are on the CB and PE devices. In the validated design:

- 1-GbE/10-GbE links are used to connect end hosts.
- The links are configured as Layer 2 with associated VLANs.
- The MTU for these links is set to 1500 bytes (the default).
- The links are configured as spanning tree edge ports.

SPX Roles and Communication Protocols

The control bridge (CB) creates, deletes, and manages port extender (PE) ports. It performs switching, routing, and forwarding for PE ports and provides centralized policy management.

The CB uses the Link Level Discovery Protocol (LLDP) to discover PE units. When the CB discovers a PE unit, it connects to the PE unit and creates a control plane using the Control and Status Protocol (CSP) over the uplink and cascade ports.

Each PE port is managed as a virtual port from the CB perspective. The CB sets up each PE unit for traffic forwarding and creates multicast and unicast forwarding tables through CSP.

The PE provides a data path between end hosts and the CB and uses LLDP to advertise 802.1BR capabilities to the CB over the upstream port or LAG. If capabilities match, the PE uses CSP to attach to the CB. The PE reports the number of available ports to the CB, and the CB allocates an E-CID for each PE port. E-tag packets are assigned EtherType 0x893f.

The PE unit creates, deletes, and manages downstream PE units. It also performs hardware-based multicast and broadcast replication. As illustrated in the following figure, the access PE unit assigns an E-tag that is based on the configured ingress port to traffic forwarded upstream. E-tags provide downlink-to-uplink associations. Transit PE units do not assign or change E-tags. PE downstream forwarding is based on the E-CID fields in the E-tags. The E-tag is removed by the access PE unit at the last hop downstream.

FIGURE 2 E-Tags Added Upstream and Removed at the Last Hop Downstream

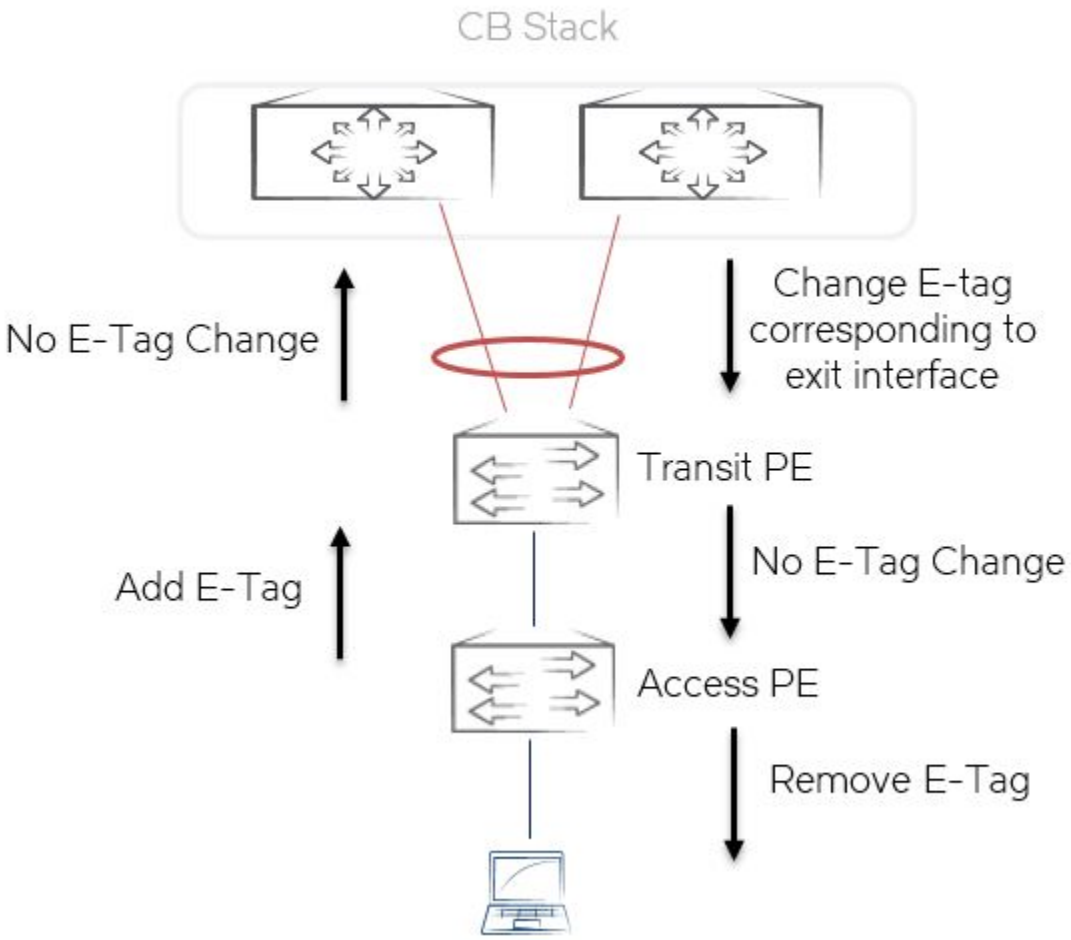

The bidirectional path between the external extended port and the corresponding internal extended port is referred to as an E-channel.

E-channels are identified by the E-channel ID (E-CID) in the E-tag (channel range 0x00 0001 to 0x3F FFFE). E-channels can be point-to-point, point-to-multipoint, or multipoint-to-point.

Control and Status Protocol

Control and Status Protocol (CSP) runs between the CB and attached PE units and is used to bring the PE units up or down.

PE CSP executes as an upper layer protocol over the Edge Control Protocol (ECP). ECP provides acknowledgment and retransmission of packets. PE CSP assumes that once a protocol data unit (PDU) is delivered to ECP, the PDU is reliably delivered to its peer.

FIGURE 3 SPX Communication Protocol Hierarchy

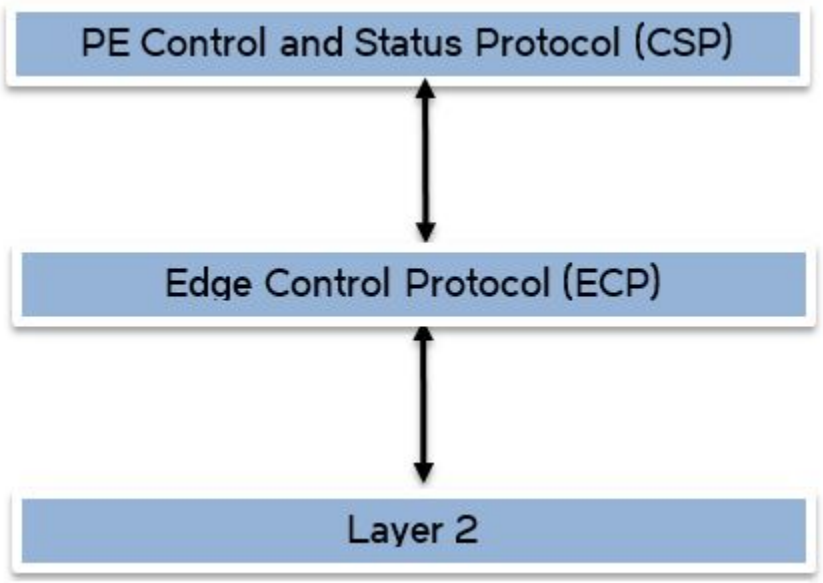

PE CSP is implemented as a simple command-and-response protocol with protocol information packaged in Type-Length-Value (TLV) triplets. Each PE CSP PDU consists of a command TLV and zero or more additional TLVs. A separate control E-channel is created between the CB and PE unit to carry PDUs, and a single instance of CSP runs between the PE unit and the CB on each link.

SPX Data Path

To understand packet flow in an SPX system, consider what happens when one host pings another through the SPX network. In the following figure, both hosts are connected to the same PE unit. All the MAC addresses are learned on the CB.

FIGURE 4 SPX Data Path Example

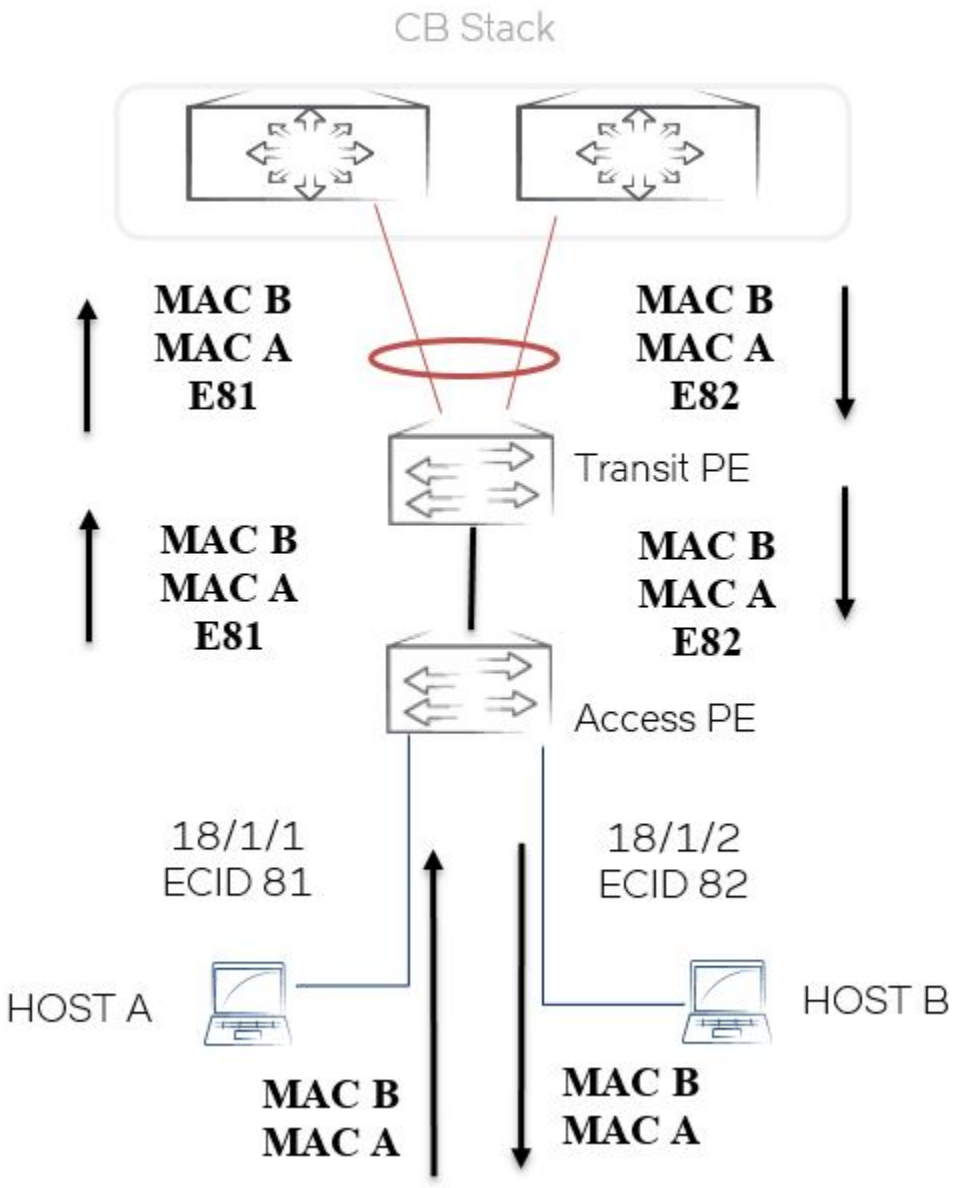

The following steps describe the series of events required for Host A to communicate with Host B.

1. Host A intends to ping Host B, both of which reside on the same LAN.
2. Because the MAC address of Host B is unknown, an ARP request is sent by Host A.
3. The packet sent by Host A travels over port 18/1/1 to access PE unit 18, where an E-tag with the E-CID 81 is added. (The CB assigned this E-CID to the PE port when it was initialized.)
4. PE unit 18 sends the packet to the CB over the uplink SPX LAG.
5. The CB installs the MAC A in its MAC table.

6. The CB floods this ARP request. It adds E-CID 82 (which is the E-CID for port 18/1/2) to the packet and sends it back out to PE unit 18 over the SPX LAG.
7. PE unit 18 looks up E-CID 82 and matches it with port 18/1/2. The PE unit removes the E-tag and sends the ARP request to Host B over port 18/1/2.
8. Host B replies to the ARP request. The packet sent by Host B travels over port 18/1/2 to access PE unit 18, where an E-tag with the E-CID 82 is added.
9. PE unit 18 sends the packet to the CB over the uplink SPX LAG.
10. The CB installs the MAC B in its MAC table.
11. The CB looks up the destination address (MAC A) and associates it in forwarding tables with port 18/1/1. The CB adds E-CID 81 (which is the E-CID for port 18/1/2) to the packet and sends it back out to PE unit 18 over the SPX LAG.
12. PE unit 18 looks up E-CID 81 and matches it with port 18/1/1. The PE unit removes the E-tag and sends the ARP reply to Host A over port 18/1/1.
13. Host A builds the required ping packet and sends it over to port 18/1/1 of PE unit 18.
14. The ping request and reply transactions are completed in the same way as the ARP request and reply transactions get completed.

FastIron Devices and SPX Communication

An ICX 7750 stack or an ICX 7750 standalone serves as the IEEE 802.1BR control bridge (CB) for attached ICX 7450 units configured as port extender (PE) units. The PE units provide connectivity to PCs, laptops, IP phones, and other access devices.

The CB communicates with the attached PE units downstream using protocols defined in the IEEE 802.1BR standards. PE units in the extended topology tag packets from attached user devices and send the packets upstream to the CB for switching and network management. The simplified network topology with centralized control and lower-cost PE devices expands capacity by hundreds of additional access ports.

FIGURE 5 IEEE 802.1BR CB-to-PE Communication

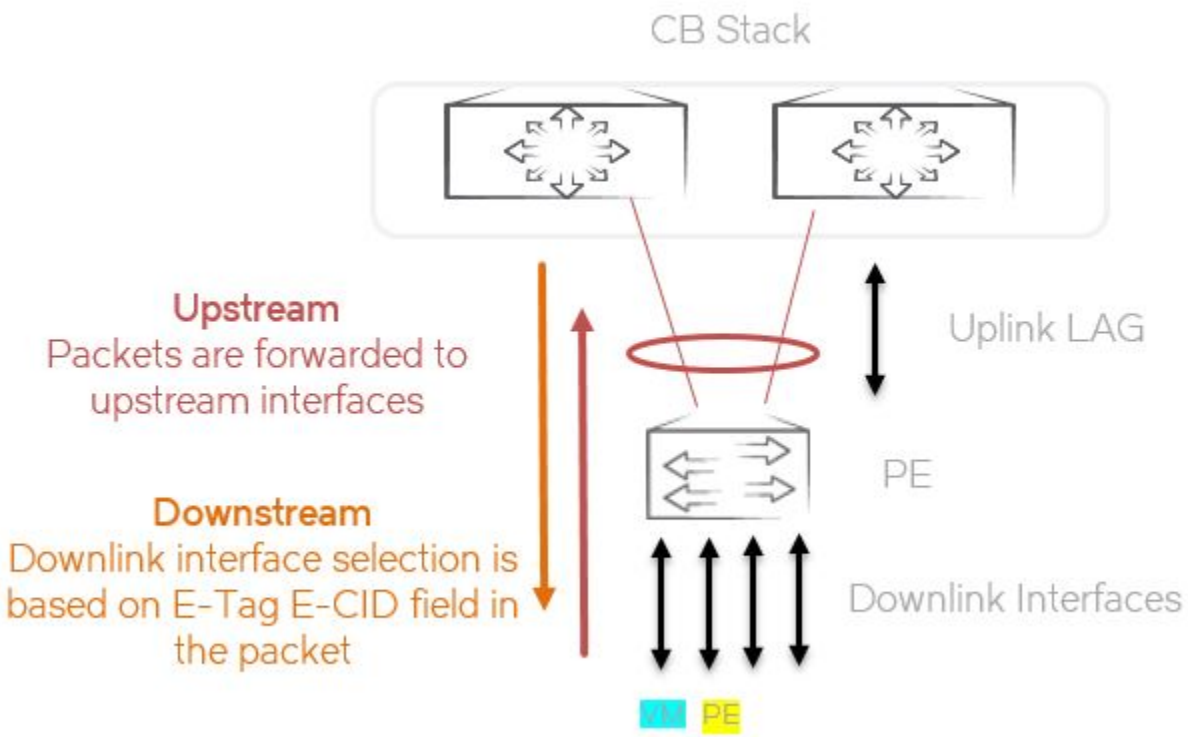

Campus Fabric Validated Design

Ruckus Campus Fabric Validation Topology

The validation design topology is defined using a two-tier architecture in which the distribution and access layers are merged together using the SPX architecture to provide network services in two different campus buildings consisting of multiple floors. In effect, this architecture creates a single distributed logical switch per building.

The buildings can communicate with each other over the core layer having the Layer 3 peering with the SPX infrastructure in each building.

The port density in each building is increased by provisioning the port extender switches (PEs) on each floor connected to the control bridge. This results in the centralized management and easy configuration of the access devices.

In this topology, link and device redundancy is built in, so as to avoid a single point of failure:

- The control bridge (CB) is built in a stack formation using ring stack links.
- Connectivity between the CB and core devices is over two different LAGs, providing the benefits of equal-cost multi-path (ECMP) routing.
- Connectivity between the CB devices and PE devices is over the LAG links.
- The SPX ring (recommended) is provisioned, providing protection against PE device failure.

FIGURE 6 Example SPX Validated Design Topology

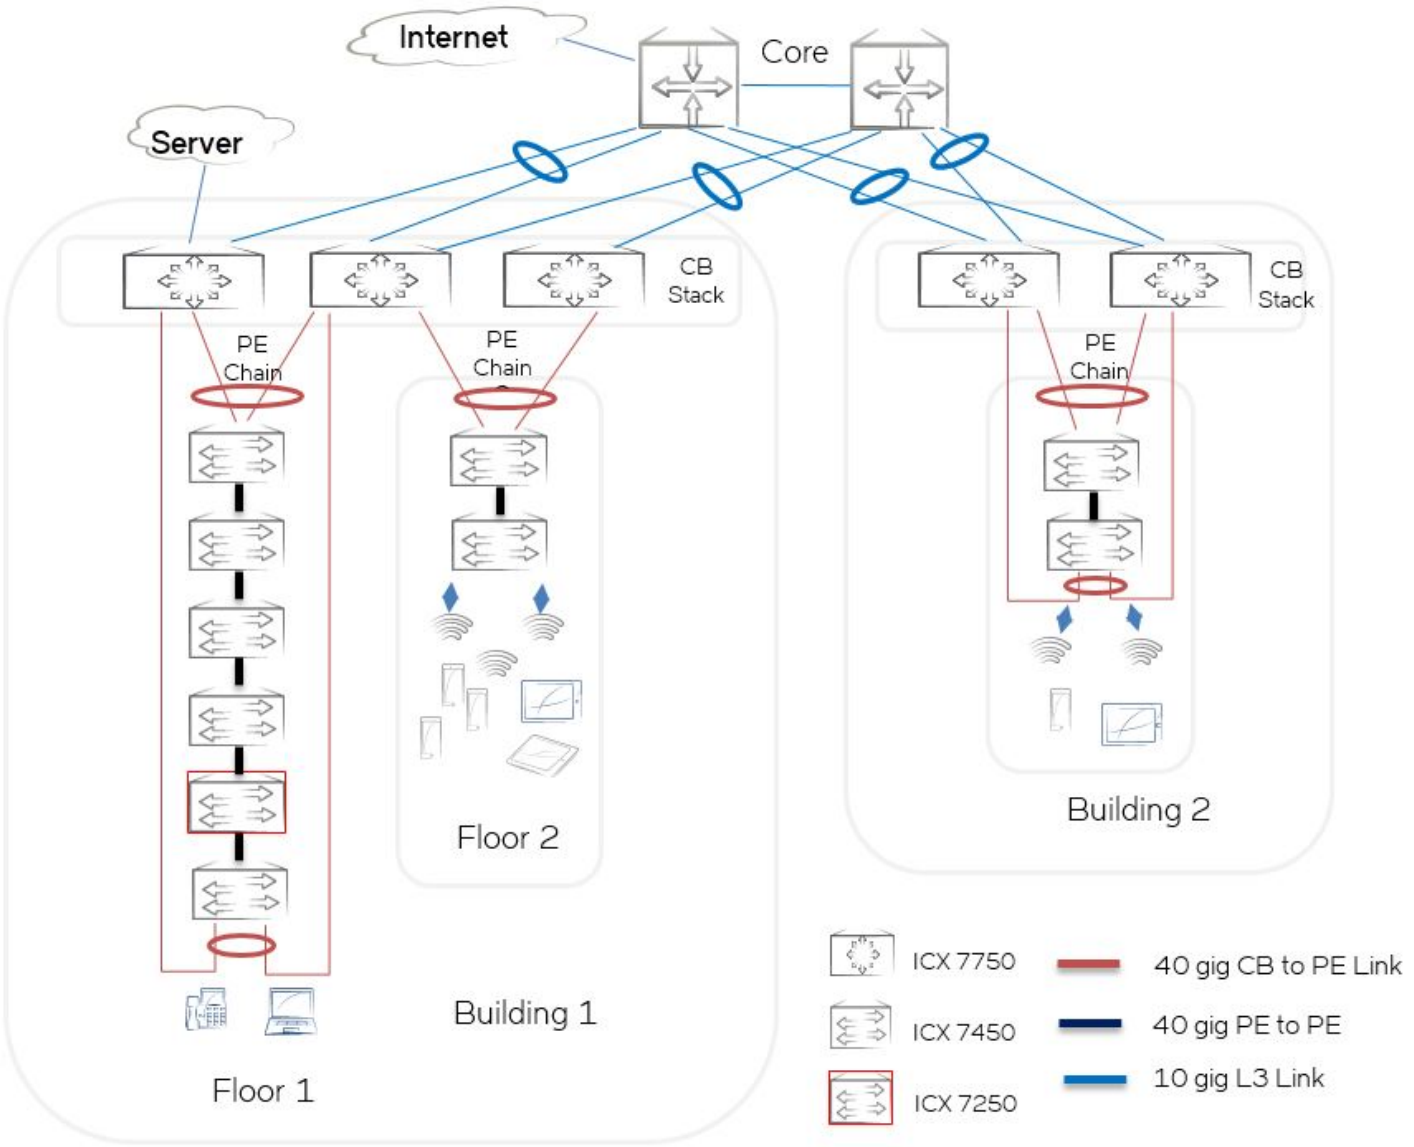

Hardware Matrix

TABLE 1 Hardware and Software Matrix

Places in the Network	Brocade Platform	Minimum Software Version
PE units	ICX 7450 ICX 7250	FastIron 08.0.50
CB stack	ICX 7750	FastIron 08.0.50
Core devices	MLXe	

Feature Matrix

TABLE 2 Supported Features Matrix

Feature Group	Feature
Platform Features	Stacking
	SPX manual bringup
	SPX ZTP
	SPX interactive setup
	LAG
Layer 2 Technologies	VLANs
	STP edge port
Layer 3 Technologies	VE Interface
	OSPF
Layer 3 Multicast	PIM-SM
	Static RP
Layer 2 Multicast	IGMP Snooping
	MLD Snooping
Security Technologies	AAA-RADIUS
	Secure-Shell
	802.1x
	Max-auth
	ACLs
Management Features	PoE
	NTP
	DHCPv4
	DHCPv6
	SNMP
Traffic Optimization	sFlow
	QoS

Design Components

Ruckus Campus Fabric Architecture

Deploying the Ruckus Campus Fabric means collapsing multiple network layers into a single logical device, combining the power of a "distributed chassis" design with the flexibility and cost-effectiveness of fixed form factor switch building blocks.

To deploy the Ruckus Campus Fabric, perform the following recommended tasks:

1. Bring up the control bridge stack.
2. Enable control bridge functionality on the control bridge stack.

3. Add the PEs in the SPX infrastructure by using one of the following options:
    a. Manual provisioning of PE: The PE chains provisioned in the "Building 1" block of the validated design topology (refer to Figure 6 on page 24) are validated using manual provisioning.
    b. SPX Zero Touch Provisioning: The PE chain provisioned in the "Building 2" block of the validated design topology (refer to Figure 6 on page 24) is validated using SPX ZTP.
    c. SPX interactive setup.

To successfully deploy the SPX infrastructure, one should take care of following design considerations.

CB Stacking

Stacking provides the ability to operate multiple devices logically as a single device. This helps in better network management and low administrative overhead, because you can use the active controller to manage the entire stack with a single configured management IP address. This provides redundancy and the ability to grow as needed (more units can be added as required). Devices in a stack can assume different roles, such as the active controller to handle stack management and configuration of all units and interface-level features, the standby controller to take over if the current active controller fails, or a stack member, which is a functional unit in the stack other than the active controller or standby controller. In a 2-unit stack, one device is elected as the active controller, and the other device is the standby controller. In a 3- to 12-unit stack, one device is the active controller, another is the standby controller, and the remaining units are stack members.

In the SPX infrastructure, the fundamental building block of a control bridge is based on the Brocade stacking technology. **The Brocade SQA team has validated a control bridge stack with a maximum of four units using the ring and linear topologies.** Brocade stacking supports both linear and ring stack topologies. Brocade highly recommends using the stack in a ring topology for the best redundancy and the most resiliency. Unicast switching follows the shortest path in a ring topology. When the ring is broken, the stack recalculates the forwarding path and resumes the flow of traffic within a few seconds. In a ring topology, all stack members must have two stacking ports; however, in a linear topology, both end units use only one stacking port, leaving the other port available as a data port.

Consider the following factors when designing the stack bandwidth:

- Failure scenarios
- Number of uplinks and their capacity
- Per-unit edge capacity
- Oversubscription ratio
- Traffic direction assumptions (north-south or east-west)

FIGURE 7 Ring Stack

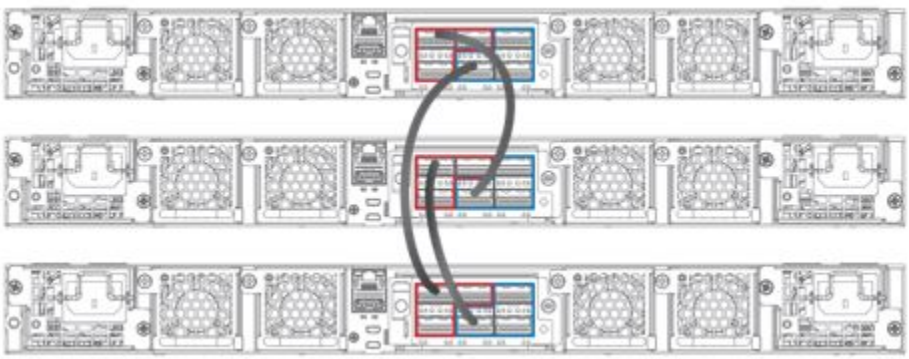

Link Aggregation Groups

Link aggregation allows a network administrator to combine multiple Ethernet links into a larger logical trunk known as a Link Aggregation Group (LAG). This results in traffic load sharing by way of redundant, alternate paths for traffic if any of the segments fail. The switch treats the trunk as a single logical link. All physical links must have the same speed and duplex settings and must connect to the same adjacent switch including stackable switches. All interface parameters in a LAG must match, including the port tag type (tagged or untagged), the configured port speed and duplex setting, and the QoS priority.

Brocade switches support the use of static and dynamic LAGs on the same device, but can use only one type of LAG for any given port. Brocade recommends using dynamic LAGs because they simplify the configurations and help avoid possible administrative mistakes in link assignments.

Brocade recommends that the design have port members distributed across multiple stack units. If a stack unit fails or is removed, the remaining ports from the LAG continue to forward traffic.

All configuration under the primary port applies to the LAG.

In SPX architecture, client-facing LAGs are not supported over PE virtual ports.

An SPX LAG can contain from 2 to 16 ports.

VLANs

A VLAN is a group of devices on one or more LANs that are configured to communicate as if they were attached to the same wire, though they are located on a number of different LAN segments.

Because VLANs are based on logical instead of physical connections, they are extremely flexible. VLANs can span multiple switches through the Layer 2 network, and they can have more than one VLAN on each switch. Trunking helps multiple VLANs on multiple switches communicate by way of a single link.

VLAN classification helps to:

- Provide security
- Isolate broadcast domains
- Classify users in a meaningful way
- Use the network on a per-VLAN basis

The Brocade ICX product family supports various types of VLANs, and Brocade recommends deploying a Layer 2 port-based VLAN in the SPX architecture: a set of physical ports that share a common, exclusive Layer 2 broadcast domain. By default, all ports on a Brocade device are members of the default VLAN. When you configure a port-based VLAN, the device automatically reassigns the port to the configured VLAN from the default VLAN. 802.1Q tagging is an IEEE standard that allows a networking device to add information to a Layer 2 packet in order to identify the VLAN membership of the packet.

In the SPX architecture, by default, each PE port is allowed membership in a maximum of four VLANs, including the default VLAN. You can increase the number of memberships per port to as many as 16 VLANs using the **max-vlan** command.

In the SPX validation reference topology, both 4 VLANs and a maximum of 16 VLANs were tested.

Unicast Routing Design

Layer 3 routing is a critical component of the SPX network. Layer 3 routing provides network connectivity between the SPX infrastucture and the core. The primary design considerations for the Layer 3 routing design are as follows:

- Resilient, standards-based, and secure routing
- High availability

- Fast convergence for voice and video applications
- Traffic load balancing
- IPv6-ready transport
- Simplified design and easy management and troubleshooting

On Brocade ICX switches, by default, none of the unicast routing protocols are enabled. For each protocol, global- and interface-level configurations are required to bring it to an operational status. OSPF is a link-state routing protocol that supports both IPv4 and IPv6 transport. In an SPX validated topology, OSPFv2 or OSPFv3 in a single area is used for the routing between an SPX logical switch and the core. OSPF is chosen because it is widely deployed in various types of networks, such as those for enterprises, service providers, and educational institutions, and these networks are considered mature and ready for the next generation of switches.

VE Interface

A virtual Ethernet (VE) interface is a logical interface that comprises physical ports and port-channel interfaces. A VE interface is the Layer 3 counterpart of a VLAN interface. In order to create a VE interface, the **router-interface** command is issued under the VLAN configuration; this configuration defines the VE interface number as well, and the configured VLAN and VE interface are coupled together. The port membership for VE is derived from the corresponding VLAN. Deletion of the VLAN deletes the VE interface, although the converse is not true; that is, deletion of the VE interface only removes the configuration from the VE interface.

By virtue of being a Layer 3-only interface, the VE interface contains only configurations that pertain to Layer 3 (such as IP addresses and IP protocols). All Layer 2 configurations are available under the VLAN interfaces.

In the SPX infrastructure, no IP addressing is possible on the physical PE ports; however, the ports can be made a part of virtual interfaces (VEs) to enable the Layer 3 functionality on them.

In the validated design, 32 VEs were configured and tested in an SPX domain.

OSPF Routing Design

In the validated design, a single instance of OSPF is configured in a single OSPF area across the SPX infrastructure. In a typical enterprise environment, the expected route table size is a few hundred routes; this design supports efficient forwarding within a single area. OSPF supports IPv4 and IPv6 routing; however, a separate OSPF process must be configured for IPv4 and IPv6. IPv4 is enabled using an OSPFv2 process, and IPv6 is enabled using an OSPFv3 process.

For redundancy and load balancing across the links, Brocade recommends dual links (preferably LAGs) from the control bridge to the core switches.

Brocade recommends the following other key feature sets:

- High-availability features such as Non-Stop Forwarding for OSPFv2 (IPv4) and Graceful Restart for OSPFv3 (IPv6)
- Security features such as MD5 for OSPFv2 and IPsec for OSPFv3

Most external connectivity links used in the SPX infrastructure are point-to-point Ethernet links that connect the control bridge to the core devices. By default, in an OSPF network, an Ethernet link acts as a multi-access network that elects a designated router (DR) and a backup designated router (BDR) to form the OSPF adjacency. In an OSPF point-to-point network, where a direct Layer 3 connection exists between a single pair of OSPF routers, there is no need for designated routers and backup designated routers. The OSPF point-to-point network establishes adjacency and converges faster. The OSPF network type must be configured as point-to-point for all external connectivity links facing the core. The administrator must ensure that the core-side routers are also configured as OSPF network type point-to-point. For optimum performance, Brocade recommends using the default cost value and timer settings for OSPFv2 and OSPFv3.

By default, IP load sharing is enabled on all Brocade ICX switches in the forwarding plane. OSPF load sharing is enabled by default when IP load sharing is enabled. The default for OSPF load sharing is four equal-cost paths, but the user can specify from two to eight paths.

SNMP traps for OSPF must be enabled for routing management.

Multicast Routing Design

On Brocade ICX product family switches, the multicast implementation helps efficiently forward IP packets that are destined for two or more receivers, which in turn helps to save network bandwidth and data replication at the source. To support multicast forwarding, various protocols must be considered to deploy. The prime multicast applications in an enterprise environment include audio/video conferencing, video streaming, IPTV, and so on. The Protocol Independent Multicast (PIM), Internet Group Management Protocol (IGMP), and Multicast Listener Discovery (MLD) protocols are required to forward multicast packets.

The multicast design recommendation for the SPX infrastructure includes IPv4 and IPv6 multicast forwarding. PIM Sparse Mode (PIM-SM) is used in the SPX infrastructure for multicast forwarding. IGMP snooping and MLD snooping are enabled on all access layer switches. The Brocade ICX product family supports the efficient and stable multicast routing protocol PIM-SM (IPv4 and IPv6).

In an SPX domain, forwarding decisions always occur in the CB, even for incoming traffic on a PE port.

In the upstream direction (from a PE to a CB), the packets are tagged with an E-tag to identify the ingress PE or SPX port. In the downstream direction (from a CB to a PE), the packets are tagged with an E-tag to identify the egress PE ports. E-channel identifiers (E-CIDs) contained in the E-tag header are used to identify a single PE port or a set of PE ports.

PIM-SM

The Protocol Independent Multicast (PIM) protocol is a broadcast and pruning multicast protocol that delivers IP multicast datagrams. This protocol employs a reverse path lookup check and pruning to allow source-specific multicast delivery trees to reach all group members. PIM builds a different multicast tree for each source and destination host group.

The rendezvous point (RP) is the meeting point for PIM Sparse sources and receivers. A PIM Sparse domain can have multiple RPs, but each PIM Sparse multicast group address can have only one active RP. PIM Sparse switches learn the addresses of RPs. A Brocade ICX family switch uses the RP to forward only the first packet from a group source to the group's receivers. After the first packet, the switch calculates the shortest path between the receiver and the source (the Shortest Path Tree, or SPT) and uses the SPT for subsequent packets from the source to the receiver. Brocade ICX product family switches calculate a separate SPT for each source-receiver pair.

In an SPX infrastructure, the core switch is configured as the RP for the entire network. The RP is statically configured on all PIM-enabled routers. When the source is activated in a PIM RP domain, the PIM First Hop (FH) registers the source to the PIM RP. In an SPX environment, the static RP configuration is simple and easy to manage as compared to dynamic methods in a multicast network.

Brocade ICX product family switches also support IPv6 Protocol Independent Multicast (PIM) Sparse. IPv6 PIM Sparse provides multicasting that is especially suitable for widely distributed multicast environments. In an IPv6 PIM Sparse network, an IPv6 PIM Sparse router that is connected to a host that wants to receive information for a multicast group must explicitly send a join request on behalf of the receiver (host).

IGMP Snooping and MLD Snooping

When Brocade ICX product family switches process a multicast packet, by default, they broadcast the packet to all ports except the incoming port of a VLAN. This behavior causes some clients to receive unwanted traffic. IGMP snooping provides multicast containment by forwarding traffic to only the ports that have IGMP receivers for a specific multicast group (destination address). A switch maintains the IGMP group membership information by processing the IGMP reports and leave messages, so traffic can be forwarded to ports receiving IGMP reports. MLD snooping provides multicast containment by forwarding traffic to only those clients that have MLD receivers for a specific multicast group (destination address) in an IPv6 network. The switch maintains the MLD group membership information by processing MLD reports and generating messages so traffic can be forwarded to ports receiving MLD reports. This is analogous to IGMP snooping. IGMP snooping and MLD snooping must be enabled on all access switches in the SPX infrastructure.

Device Access Security

Device access must be secured with user access privileges and access control. A RADIUS-based authentication method can be used to control device access in the SPX infrastructure. The Brocade ICX and MLXe platforms support device access through Telnet, SSH, and web management.

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is a networking protocol that provides centralized Authentication, Authorization, and Accounting (AAA) management for users who connect and use a network service. RADIUS is a client/server protocol that runs in the application layer and uses the User Datagram Protocol (UDP) for transport. Network access servers, Brocade devices that control access to a network, usually contain a RADIUS client component that communicates with the RADIUS server.

When RADIUS authentication is implemented, the Brocade device prompts the user for a username and password, and the supplied credentials are matched against the username and password in the RADIUS database. If the username is found in the database, the RADIUS server validates the password, directing the device to grant access to the user with a specified privilege level.

RADIUS can be optionally configured to provide authorization wherein the Brocade device consults a list of commands supplied by the RADIUS server to determine whether a user can issue an entered command. RADIUS can also be configured to provide accounting, which causes the Brocade device to log information on a RADIUS accounting server when specified events occur on the device.

For the SPX infrastructure solution, the RADIUS server installed in the server farm connected to the core switch is employed to provide device access security including Telnet access, SSH access, web management access, access to the privileged EXEC level and configuration levels of the CLI, and network access security including 802.1X authentication and MAC authentication services.

Secure Shell

Secure Shell (SSH) is a mechanism for allowing secure remote access to management functions.

SSH provides a function similar to Telnet. Users can log in to and configure the device using a publicly or commercially available SSH client program, just as they can with Telnet. Unlike Telnet, which uses a clear-text connection, SSH provides a secure, encrypted connection to the device.

The Brocade SSH2 implementation is compatible with all versions of the SSH2 protocol (2.1, 2.2, and so on) and multiple commonly used SSH clients (SSH Secure Shell 3.2.3, VanDyke SecureCRT 5.2.2, F-Secure SSH Client 5.3 and 6.0, PuTTY 0.62). At the beginning of an SSH session, the Brocade device negotiates the version of SSH2 to be used. The latest version of SSH2 supported by both the Brocade device and the client is the version that is used for the session. Once the SSH2 version is negotiated, the encryption algorithm with the highest security ranking is selected for the session.

For SSH to work, the public-private key pair must be generated on the server, and the public key is shared with the SSH client during the start of the session (with the option to save the public key for the future sessions); the client uses the public key to encrypt all communication with the server.

User authentication and authorization are handled using AAA as in the case of Telnet and other protocols, and no separate configuration is required.

Network Access Security

Brocade ICX products support various authentication methods to control user access to the network. The Brocade Flexible Authentication (FlexAuth) feature allows the network administrator to set the sequence of the authentication methods to be attempted on a switch port. This feature supports two methods: 802.1X authentication and MAC authentication. By default, 802.1X is attempted first, and if the user is not 802.1X-capable, MAC authentication is attempted. If both of these authentication methods fail, the client is blocked. This allows each client connected to the same switch port to have a different network policy using the MAC-based VLANs. The default behavior can be changed by the administrator using the CLI.

The SPX administrator can use these security features to manage the users and wired end devices, such as servers, IP phones, wireless access points, and IP video equipment. Access for users and end devices is authenticated using 802.1X or MAC authentication. Based on the defined authentication parameters, proper authorization for available network services is provided.

After successful authentication, the RADIUS server returns the details of the VLAN where the client should be assigned based on the user profile configured on the RADIUS server. The client (MAC address of the client) is moved to the configured VLAN as a MAC VLAN member. The network administrator must create the VLANs that are assigned to clients by RADIUS. For example, RADIUS returns VLAN 200 for Client A and VLAN 201 for Client B; these two VLANs must be set up on the switch. Any additional configurations that are required, such as virtual interfaces 200 and 201, are also created in these VLANs respectively, so that Client A and Client B can use the virtual interface IP address as their gateway IP address.

If RADIUS assigns a dynamic ACL to at least one client on the interface, the maximum number of MAC sessions that can be authenticated is limited to 32 on all Brocade devices, which severely limits the maximum number of MAC sessions (default 1024); thus, for an SPX infrastructure, it is desirable to assign the network policies on the next-hop Layer 3 device on a per-VLAN basis.

By default, the number of MAC sessions that can be authenticated on a single interface is two; this number can be changed using the **authentication max-sessions** command under the port configuration.

By default, the authentication VLAN mode is "single untagged." In an SPX infrastructure network, because multiple untagged VLANs must be assigned to different users, the mode must be changed to "multiple untagged."

The network administrator can configure specific VLANs to be assigned to the user based on the authentication results, such as successful authentication, failed authentication, and RADIUS timeout scenarios. The default action is to block access to the user if the authentication fails. However, the administrator can assign the user to special VLANs, such as the Guest VLAN, Critical VLAN, or Restricted VLAN, based on the organization's security policy and can grant limited access to the authentication failures and other scenarios.

VLAN Requirements for Flexible Authentication

To deploy Flexible Authentication, VLANs, such as the Critical VLAN, Restricted VLAN, Guest VLAN, and Auth-Default VLAN, are used for various success, failure, and timeout scenarios. After authentication is enabled on the port, the port becomes a part of the Auth-Default VLAN. After successful authentication, the VLAN is assigned to the client (MAC address of the client), not to the entire port.

Critical VLAN—There may be scenarios where the RADIUS server times out or is not available, resulting in authentication failure. This can happen the first time that the client is authenticating or when it reauthenticates. In this situation, the network administrator can decide to grant some or the same access as the original instead of blocking the access. This VLAN should be configured with the desired access levels.

Restricted VLAN—When authentication fails, the client can be moved into a Restricted VLAN instead of blocking the access completely. The network administrator may decide to grant some access in this scenario, instead of blocking the access. This VLAN should be configured with the desired access levels.

Guest VLAN—Specifically used with 802.1X security authentication, the client is moved to a Guest VLAN when it does not respond to the 802.1X requests for authentication. It is possible that the client does not have the 802.1X authenticator loaded and thus needs some

way to access the network from where it can download the client software. The network administrator can configure the Guest VLAN with access methods as required.

Auth-Default VLAN—When a port is enabled for 802.1X or MAC authentication, by default, the port is assigned to this VLAN as a MAC-based VLAN member. Sometimes the RADIUS server may authenticate the client, but may not return the required VLAN information on where the client should be placed. In this scenario, the client is placed into the Auth-Default VLAN.

802.1X

802.1X is an IEEE standard for port-based network access control (PNAC) to provide an authentication mechanism for devices that want to attach to a LAN or WLAN. 802.1X authentication involves three parties: a supplicant, an authenticator, and an authentication server. The supplicant is a client device (such as a laptop) that wishes to attach to the LAN or WLAN. The authenticator is a network device, such as an Ethernet switch or wireless access point. And the authentication server is typically a host running software supporting RADIUS and EAP.

The supplicant provides credentials, such as the username and password or digital certificate, to the authenticator, and the authenticator forwards the credentials to the authentication server for verification. Based upon the scenario, the following actions are taken:

- If the authentication server determines that the credentials are valid, one of two actions is taken:
  - If the RADIUS server returns the VLAN ID, the supplicant is placed in that VLAN.
  - If the RADIUS server does not return the VLAN ID, the Auth-Default VLAN is used.
- If the supplicant is not 802.1X-capable, the authenticator either puts it in the Guest VLAN or tries MAC authentication if enabled on the port. If the Guest VLAN is not configured, the authenticator keeps on retrying the authentication.
- If the supplicant fails to get authenticated and the authentication fail action is defined as Restricted VLAN, it is placed in the Restricted VLAN.
- If the supplicant fails to get authenticated and an authentication fail action is not defined, its MAC address is blocked in the hardware (default action).
- If the authentication server (RADIUS) itself times out, and the authentication timeout action is defined as Critical VLAN, the supplicant is placed in the Critical VLAN.

MAC Authentication

MAC authentication is a mechanism by which incoming traffic originating from a specific MAC address is forwarded by the Brocade switch only if the source MAC address is successfully authenticated by a RADIUS server. The MAC address itself is used as the username and password for RADIUS authentication. If the RADIUS server cannot validate the user's MAC address, it is considered an authentication failure, and a specified authentication-failure action can be taken. MAC authentication supports the use of the Critical VLAN and the Restricted VLAN.

Based upon the scenario, the following actions are taken:

- If MAC authentication is successful for the client authentication, one of two actions is taken:
  - If the RADIUS server returns the VLAN ID, the supplicant is placed in that VLAN.
  - If the RADIUS server does not return the VLAN ID, the Auth-Default VLAN is used.
- If the RADIUS server cannot authenticate the client's MAC address, it is considered an authentication failure, and if the authentication-failure action is configured as Restricted VLAN, the client is placed in the Restricted VLAN.
- If the RADIUS server cannot authenticate the user's MAC address, it is considered an authentication failure, and if the authentication-failure action is not configured as Restricted VLAN, the client's MAC address is blocked in the hardware (default action).
- If the authentication server (RADIUS) itself times out, and the authentication-timeout action is defined as Critical VLAN, the client is placed in the Critical VLAN.

To deploy MAC authentication, the network administrator must maintain a MAC database in the RADIUS server. To ease deployment in an existing network, Brocade ICX products support three different MAC address formats for RADIUS Attribute 1 (username) and RADIUS Attribute 2 (password).

The SPX infrastructure network administrator can use the Flexible Authentication feature set for user access management for the wired and wireless clients by using either 802.1X or MAC authentication. Based upon the defined authentication parameters, proper authorization for available network services is provided. Using Flexible Authentication, the network administrator can set the sequence of the authentication methods to be attempted, resulting in better user management.

Access Control Lists

Brocade devices support rule-based Access Control Lists (ACLs), sometimes called hardware-based ACLs, wherein the decisions to permit or deny packets are processed in hardware, and all permitted packets are switched or routed in hardware. All denied packets are also dropped in hardware. Brocade ICX devices support both inbound and outbound ACLs. Two types of IP ACLs are available: standard ACLs, which permit or deny packets based on source IP address,; and extended ACLs, which permit or deny packets based on source and destination IP address and also IP information. Valid extended ACL IDs are a number from 100 to 199 or a character string.

The default action when no ACLs are configured on a device is to permit all traffic. However, once you configure an ACL and apply it to a port, the default action for that port is to deny all traffic that is not explicitly permitted on the port.

In an SPX infrastructure, it is important to apply proper security policies in order to ascertain that rogue traffic is contained and that legitimate users can access the intended network services. An ACL enables the network administrator to tailor the security policies as required, restricting or allowing user access based on parameters in the traffic headers.

In the SPX infrastructure, the application of ACLs is further optimized by programming the entries directly in the hardware of the PE units, essentially dropping the unwanted traffic on the port of entry itself.

Management Feature Sets

Power over Ethernet

Brocade Power over Ethernet (PoE) devices are compliant with the standards described in the IEEE 802.3af and 802.3at specifications for delivering inline power. PoE technology eliminates the need for an electrical outlet and dedicated UPS near IP-powered devices. With power-sourcing equipment, such as a Brocade ICX PoE device, power is consolidated and centralized in wiring closets, improving the reliability and resilience of the network.

Brocade devices support Endspan with Alternative A, wherein power is supplied through the Ethernet ports on a power-sourcing device (specifically, power is carried over the live wire pairs that deliver data) or Alternative B, the two spare pairs.

- An auto-discovery mechanism detects whether the device requires power and how much power is needed (802.3af or 802.3at).
- 802.3af (PoE) provides 15.4 watts (44 to 50 volts); 802.3at 2008 (PoE+) provides 30 watts (52 to 55 volts); Power over HDBaseT (PoH) provides 95 watts (48 volts).
- The 802.3af and 802.3at standards support PoE and PoE+ on 10/100/1000-Mbps Ethernet ports that operate over standard Category 5 unshielded twisted pair (UTP) cable.
- Voice over IP (VoIP) phones, wireless LAN access points, and IP surveillance cameras are the commonly deployed devices in an enterprise environment; the access switches can provide the PoE for these devices.

In the SPX infrastructure, the aggregation and access is collapsed and it is important to enable PoE on the end-user connecting ports, so as to provide enough power to the connected PoE devices.

Network Time Protocol

For any network, including an SPX infrastructure, it is imperative to synchronize the clocks on all devices because it helps to take a snapshot of a network at any particular time. This is especially useful to determine the time of a particular event. For example, to logically deduce the sequence of events on different network elements while accessing the syslog data, it is essential that the timestamps in all logs be in sync. This timekeeping service is provided by the industry-standard Network Time Protocol (NTP), which is configured on all network devices to synchronize the system time with the NTP server.

NTP synchronizes the system time with the NTP server, but it is also important that the time zone for the devices be maintained across the network, which is achieved by defining the local clock and time zone on the individual devices.

DHCPv4 and DHCPv6

Dynamic Host Configuration Protocol (DHCP) is a standardized network protocol used on Internet Protocol (IP) networks for dynamically distributing network configuration parameters, such as IP addresses for interfaces and services. With DHCP, computers request IP addresses and networking parameters automatically from a DHCP server, reducing the need for a network administrator or a user to configure these settings manually.

DHCP operations fall into four phases: server discovery, IP lease offer, IP request, and IP lease acknowledgment. These stages are often abbreviated as DORA for discovery, offer, request, and acknowledgment.

In the SPX infrastructure, any host machine that intends to use network services is assigned an IP address and related information dynamically using DHCP, with the centralized DHCP server hosted in the server farm. In this case, because DHCP clients are not directly served by DHCP servers, DHCP relay services are configured in the SPX domain that is under the relevant VE interface with the corresponding helper address configured. The IP address configured under the interface belongs to the same IP subnet pool from which the dynamically assigned IP address is drawn. Further, the DHCP server IP address is configured as the helper address. The DHCP client broadcasts on the local link; the relay agent (distribution switch) receives the broadcast and transmits it to one or more DHCP servers using unicast.

Similar to DHCPv4, DHCPv6 is designed to enable the distribution of IPv6 network configuration parameters.

Simple Network Management Protocol

Simple Network Management Protocol (SNMP) is a set of protocols for managing networks. SNMP sends messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.

SNMP access to the switch can be restricted to a specific IP address or a specific VLAN.

A user can use the following methods:

- A community-string match in SNMP versions 1 and 2
- A user-based model in SNMP version 3

Traffic Monitoring and sFlow

sFlow helps analyze traffic trends on various configured links, and alerts can be set to inform the network administrator when threshold levels are reached. This analysis aids in taking preventive measures to overcome unexpected congestion and also aids with network capacity planning. sFlow settings on switches can be configured from Brocade Network Advisor or from the switch CLI. In the SPX infrastructure design, Brocade Network Advisor acts like an sFlow collector to receive samples from switches to analyze traffic trends. Brocade Network Advisor provides traffic information, such as Layer 2, Layer 3, and Layer 4 information, and also the topper list of who is consuming the most network bandwidth. The sampling rate and frequency play a key role, so care should be taken while configuring these parameters; otherwise, the switch CPU is impacted.

Quality of Service (QoS)

Quality of Service (QoS) features are used to prioritize the use of bandwidth on a switch. When QoS features are enabled, traffic is classified as it arrives at the switch, and it is processed on the basis of configured priorities. Traffic can be dropped, prioritized for guaranteed delivery, or subjected to limited delivery options as configured by a number of different mechanisms.

In the SPX infrastructure design, because users may use different kinds of services for data, voice, and video reception, it is important that the different traffic be treated differently so as to provide the appropriate amount of bandwidth to each service, while not starving the other network services. For example, high jitter while using real-time applications such as VoIP, VoD, or webcast results in a less than acceptable user experience. Well-defined QoS policies can help to achieve this requirement.

One major advantage of using SPX is that the QoS policies are configured once and then attached to approximately 1700 ports; this results in ease of configuration and management.

In an SPX system, forwarding decisions for PE ports are made in the CB. When a PE unit receives frames through its extended ports, the PE appends an E-tag to each of the frames and transmits them to an upstream port toward the CB. After the CB makes forwarding decisions, the frames are forwarded to destination modules in the SPX system. If the destination module is a PE unit, the frames are sent with an E-tag. The E-tag carries an E-PCP field consisting of three bits, which allows the E-PCP to carry eight priority classifications. QoS classifications are supported for PE ports and are encoded in the E-PCP. A priority classified at an edge device is honored in both upstream and downstream directions. When congestion occurs in SPX ports, frames are dropped based on their priorities.

Classification—The process of selecting packets on which to perform QoS, reading the QoS information, and assigning a priority to the packets. The classification process assigns a priority to packets as they enter the switch. This priority can be determined on the basis of information contained within the packet, or the priority can be assigned to the packet as it arrives at the switch. Once a packet or traffic flow is classified, it is mapped to a forwarding priority queue.

Packets on Brocade devices are classified in up to eight traffic classes with values from 0 to 7. Once a packet is classified, it is mapped to a forwarding queue. Packets with higher-priority classifications are given precedence for forwarding.

Marking (optional)—The process of changing the packet QoS information (the 802.1p and DSCP information in a packet) for the next hop. For example, for traffic coming from a device that does not support Differentiated Services (DiffServ), you can change the packet IP precedence value into a DSCP value before forwarding the packet.

Scheduling—The process of mapping a packet to an internal forwarding queue based on its QoS information and servicing the queues according to a mechanism.

DSCP-based QoS is not automatically honored for switched traffic. The default is 802.1p-to-CoS mapping. To honor DSCP-based QoS, enter the **trust dscp** command at the interface level of the CLI. By default, the bandwidth scheduling mechanism is mixed weighted priority with strict priority (WRR with strict priority), which combines both the SP and WRR mechanisms. The combined method enables the Brocade device to give strict priority to delay-sensitive traffic and weighted round-robin priority to all other traffic types.

The Brocade device assigns strict priority to traffic in qosp7 and qosp6, and it assigns weighted round-robin priority to traffic in qosp0 through qosp5. Thus, the Brocade device schedules traffic in queue 7 and queue 6 first, based on the strict priority queueing method. When there is no traffic in queue 7 and queue 6, the device schedules the other queues in weighted round-robin fashion from the highest priority queue to the lowest-priority queue.

QoS Implementation Guidelines

1. Re-mark the traffic coming with COS value 6/7 to COS value 5 as COS 6/7 maps internally for control protocols.
2. Configure Trust DSCP on the ingress ports of the access switches where the end-user devices are connected.
3. Re-mark the DSCP in voice packets to 46 based on the UDP port number.
4. Re-mark the voice-control packets that normally arrive marked with DSCP 24/26 to 46.
5. Perform internal priority marking using ACL clauses.

Ruckus Campus Fabric Configuration

This section covers the aspects of provisioning and validation of the Ruckus Campus Fabric infrastructure.

SPX Design Considerations

- The Brocade ICX 7450 platform can have a maximum of 3 no. of 40-GbE or 12 no. of 10-GbE interfaces available for uplink and downlink connectivity. It is recommended to use 6 10-GbE ports in a LAG formation in both directions for achieving the effective bandwidth of 60 Gbps in both directions.
- The Brocade ICX 7250 platform can have a maximum of 8 no. of 10-GbE interfaces available for uplink and downlink connectivity. It is recommended to use 4 10-GbE ports in a LAG formation in both directions for achieving the effective bandwidth of 40 Gbps in both directions.
- In an SPX domain, no local switching occurs on the PE ports and all data and control packet processing occurs in the CB. In cases where the available bandwidth is a concern, it is better to increase the breadth of the PE chains (8 chains supported) then to increase the depth of PE chains (6 PE units per PE chain).

SPX Infrastructure Provisioning

Only the Brocade ICX 7750 platform can be used as a control bridge in the SPX architecture. To use the ICX 7750 as a control bridge, you must enable the SPX control bridge functionality on the switch.

In the validated design, an ICX 7750 stack is configured as a control bridge with a combination of ICX 7450 and ICX 7250 switches acting as port extender units.

The SPX infrastructure is validated to be provisioned in multiple ways:

- Manual provisioning of PE
- Zero Touch Provisioning
- SPX interactive-setup

Creating a Control Bridge Stack

In the SPX infrastructure, to avoid a single point of failure, it is recommended to bring up the devices marked for the control bridge role in a stack formation. The following configuration and verification steps bring up the ICX 7750 device in a stack formation.

```
CF_SW4(config)# stack enable
Enable stacking. This unit actively participates in stacking

CF_SW4(config)# exit
! Enter the stack secure-setup command.
CF_SW4# stack secure-setup
CF_SW4#Discovering the stack topology...

Current Discovered Topology - RING
Available UPSTREAM units
Hop(s)  Id      Type          Mac Address
 1        2   ICX7750-48XGF   cc4e.24d2.1d80
Available DOWNSTREAM units
Hop(s)  Id      Type          Mac Address
 1        2   ICX7750-48XGF   cc4e.24d2.1d80
No new units found...

Selected Topology:
Active  Id      Type          Mac Address
        1    ICX7750-48XGF   cc4e.24d0.5e80
Selected UPSTREAM units
Hop(s)  Id      Type        Mac Address
 1       2    ICX7750-48XGF   cc4e.24d2.1d80
Selected DOWNSTREAM units
Hop(s)  Id      Type        Mac Address
 1       2    ICX7750-48XGF   cc4e.24d2.1d80

Do you accept the unit id's (y/n)? y

Selected Topology:
Active Id      Type    MAC Address
1       ICX7750-48XGF  cc4e.24d0.5e80
Selected UPSTREAM units
Hop(s) Id      Type    MAC Address
1       2       ICX7750-48XGF  cc4e.24d2.1d80
Selected DOWNSTREAM units
Hop(s) Id      Type    MAC Address
1       2       ICX7750-48XGF  cc4e.24d2.1d80
Do you accept the unit ids (y/n)?: y
```

Issue stack enable command to enable stacking on all the stack units.

Enter the stack secure-setup command on the unit you want to configure as active unit, which discovers the stack topology.

Enter y to accept the topology. Selected topology is shown to confirm the user acceptance.

To accept the unit ID assignments, enter y.

If you accept the unit IDs, the stack is formed.

Verifying the Stack Formation

```
CF_SW4# show stack
T=9h52m55.0: alone: standalone, D: dynamic cfg, S: static
ID   Type          Role    Mac Address    Pri State   Comment
1  S ICX7750-48XGF active  cc4e.24d0.5e80 128 local   Ready
2  S ICX7750-48XGF standby cc4e.24d2.1d80  10 remote  Ready

    active       standby
     +---+        +---+
 -2/4| 1 |2/1--2/4| 2 |2/1-
  |   +---+        +---+   |
  |                        |
  |------------------------|
Standby u2 - protocols ready, can failover
Current stack management MAC is cc4e.24d0.5e80

There are 2 PEs. Please use "show spx" to display PE information.
```

Enabling Control Bridge Functionality on ICX 7750 Stack

CB functionality is enabled in global configuration mode, followed by the CB configuration under the spx cb-configure mode.

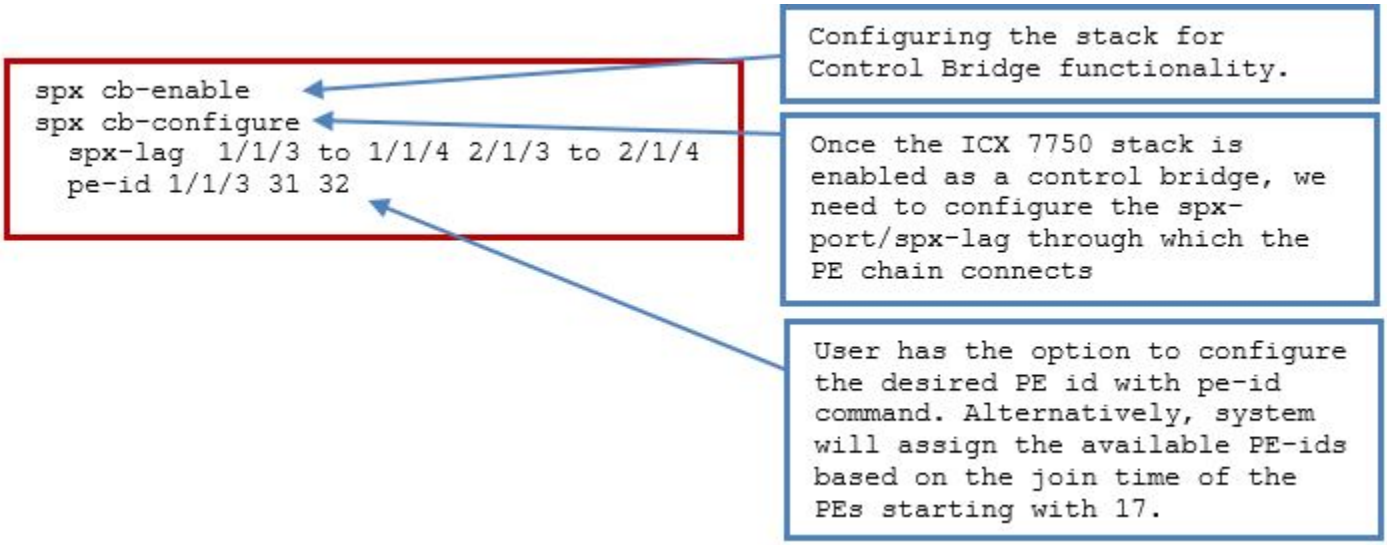

SPX Manual Provisioning

On a standalone ICX 7450, issue the **spx-enable** command to bring up the unit in provisional PE mode. When the unit is PE-enabled, the system generates two default SPX ports for the PE. Following this, the user must save the configuration and reload the PE device, enabling it to join the SPX PE chain.

The user has the option to either use the default SPX ports or configure a different set of ports as SPX ports or an SPX LAG.

On the PE-candidate unit, configure the following to bring it up as a PE.

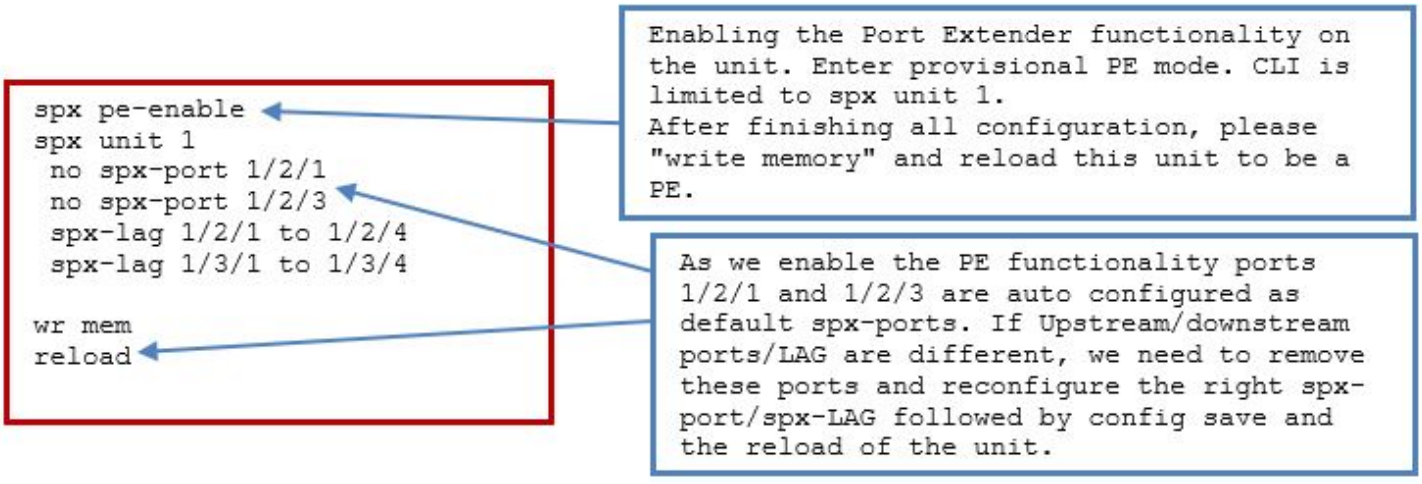

SPX Zero Touch Provisioning

Typically, with traditional SPX deployment, the deployment process is tedious and requires a lot of steps and configuration from the user end. For example, customers must prepare each PE individually and do the required configuration procedure on the CB master. Sometimes, it is possible that the user mentions a wrong connection while configuring the SPX ports or LAG and runs into difficulty while bringing up an SPX system.

To use ZTP, call out the fabric ports on the CB and when the new access switches connect to these ports, the image on the attached device is upgraded, rebooted into PE mode, and LAGs are formed on the uplinks automatically.

In the global configuration mode on the CB, configure the following to enable ZTP.

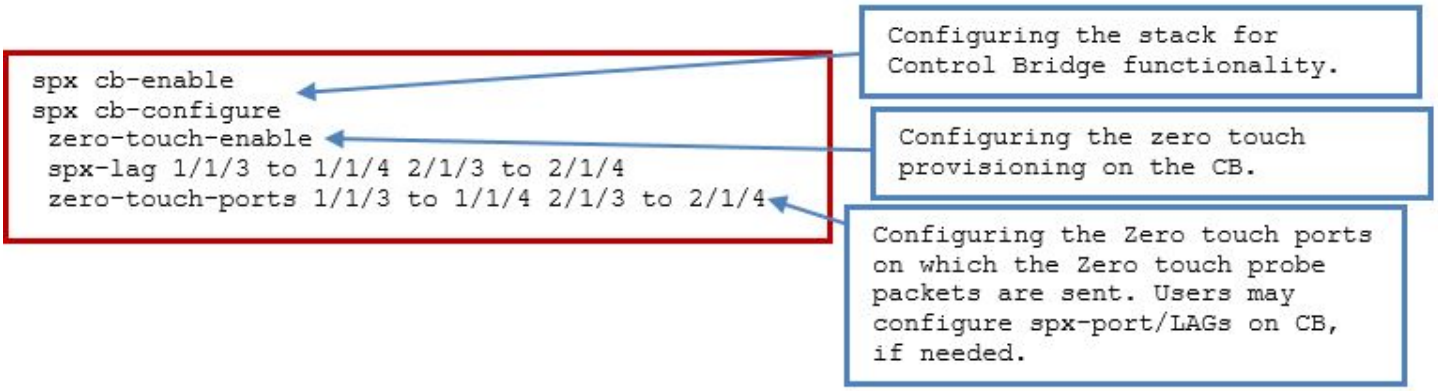

SPX Interactive-Setup

SPX interactive-setup allows users to perform tasks such as determining which units to include in the Campus Fabric domain and which IDs they will be assigned. SPX interactive-setup is a superset of ZTP features.

SPX interactive-setup can handle all the topologies supported by ZTP. In addition, SPX interactive-setup can handle some special topologies that cannot be handled by ZTP. For example, the following invalid topology can be handled by SPX interactive-setup so that the user has the option to remove the link between units 1 and 2 or the link between units 1 and 3.

FIGURE 8 Invalid topology

SPX interactive-setup is a menu-based utility in which the user is prompted for the input as required.

To use SPX interactive-setup, ZTP must be disabled. SPX interactive-setup cannot be used to bring up a CB stack. A CB stack must be provisioned before using SPX interactive-setup, as detailed in Creating a Control Bridge Stack on page 36.

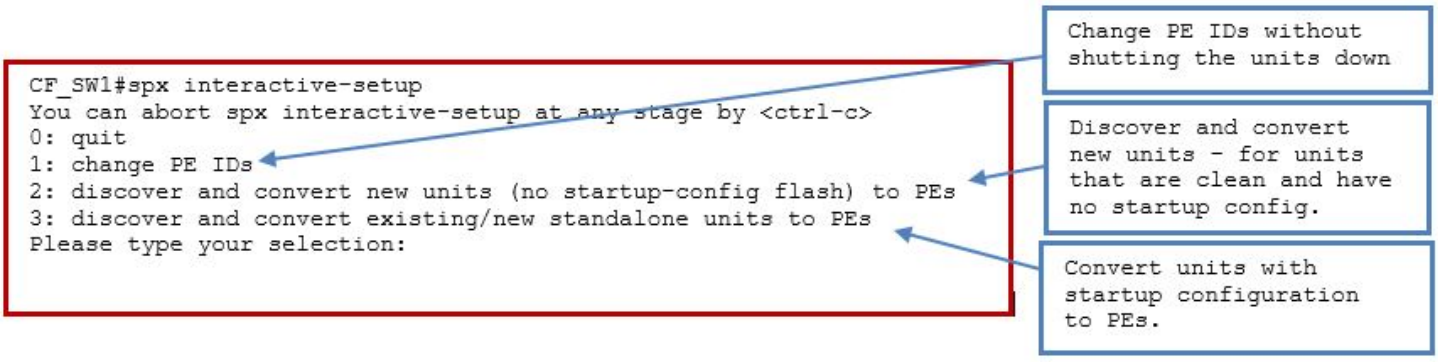

Debugging and Verification Commands

Once an SPX domain is formed, **show** commands related to SPX can be used to verify that the SPX is operational. The following**show** commands help debug issues.

SPX Verification Commands on CB

```
CF_SW4#show spx
T=2d8h18m42.4: alone: standalone, D: dynamic cfg, S: static
ID   Type          Role    Mac Address    Pri State   Comment
1  S ICX7750-48XGF active  cc4e.24d0.5e80 128 local   Ready
2  S ICX7750-48XGF standby cc4e.24d2.1d80  10 remote  Ready
31 S ICX7450-24P   spx-pe  cc4e.248d.24a8 N/A remote  Ready
32 S ICX7450-24P   spx-pe  cc4e.248b.0330 N/A remote  Ready

    active       standby
     +---+        +---+
 -2/4| 1 |2/1--2/4| 2 |2/1-
  |   +---+        +---+   |
  |                        |
  |------------------------|
            +----+       +----+
   1/1/3==2/1| 31 |3/1==4/1| 32 |
            +----+       +----+

CF_SW4#show spx csp all

PE 31 MAC: cc4e.248d.24a8
   CSP Oper: yes, Attach time: 1m17.4, up time: 2 day(s) 8 hour(s) 18 minute(s) 48 second(s)
   CB Spx Lag id: 3072, cur state up, IPC/ECP Port: 1/1/3
   PE Spx Uplink Port: 31/2/1, cur state up
   Number of Traffic Class: 8
   Priority Flow Control: no
   CSP control ECID handshake complete: yes
   CSP control ECID: 1195
   CSP Alternate control ECID: 1196
   Total number of configured ports: 33
   CSP number of allocated ECIDs (VPs created, excl Control VP): 27
   CSP last Tx Trans ID=10, last Rx Trans ID=3
   ECP txErrors=0, sequence=14 firstSeq=14 lastSeq=13 firstAckIdx=0 ackIdx=0
   Next PE: 32
   PE Spx downlink Port: 31/3/1, cur state up
   Previous PE: None
   Local CSP Major version is 1 Minor version 1
   Peer  CSP Major version is 1 Minor version 1
   Oper  CSP Major version is 1 Minor version 1

PE 32 MAC: cc4e.248b.0330
```

```
   CSP Oper: yes, Attach time: 3m15.4, up time: 2 day(s) 8 hour(s) 16 minute(s) 50 second(s)
   CB Spx Lag id: 3072, cur state up, IPC/ECP Port: 1/1/3
   PE Spx Uplink Port: 32/4/1, cur state up
   Number of Traffic Class: 8
   Priority Flow Control: no
   CSP control ECID handshake complete: yes
   CSP control ECID: 1275
   Total number of configured ports: 33
   CSP number of allocated ECIDs (VPs created, excl Control VP): 30
   CSP last Tx Trans ID=2, last Rx Trans ID=7
   ECP txErrors=0, sequence=6 firstSeq=6 lastSeq=5 firstAckIdx=0 ackIdx=0
   Next PE: None
   Previous PE: 31
   Local CSP Major version is 1 Minor version 1
   Peer  CSP Major version is 1 Minor version 1
   Oper  CSP Major version is 1 Minor version 1
```

```
CF_SW4#show stack
T=2d8h29m57.8: alone: standalone, D: dynamic cfg, S: static
ID   Type          Role    Mac Address    Pri State   Comment
1  S ICX7750-48XGF active  cc4e.24d0.5e80 128 local   Ready
2  S ICX7750-48XGF standby cc4e.24d2.1d80  10 remote  Ready

    active       standby
     +---+        +---+
 -2/4| 1 |2/1--2/4| 2 |2/1-
  |   +---+        +---+   |
  |                        |
  |------------------------|
Standby u2 - protocols ready, can failover
Current stack management MAC is cc4e.24d0.5e80

There are 2 PEs. Please use "show spx" to display PE information.
```

This line conveys that this switch is SPX capable and has 2 PEs.

SPX Verification Commands on PE

```
[PE]local-31@ICX7450-24P Router#show running-config
Current configuration:
!
ver 08.0.50T213
!
spx pe-enable
spx unit 31
  module 1 icx7450-24p-poe-port-management-module
  module 2 icx7400-xgf-4port-40g-module
  module 3 icx7400-xgf-4port-40g-module
  module 4 icx7400-qsfp-1port-40g-module
  spx-lag 31/2/1 to 31/2/4
  spx-lag 31/3/1 to 31/3/4
!
end
[PE]local-31@ICX7450-24P Router#show spx csp all

   CSP Oper: yes, Attach time: 33.4, up time: 2 day(s) 8 hour(s) 26 minute(s) 43 second(s)
   PE Spx Lag id: 1, cur state up, IPC/ECP Port: 31/2/1
   Number of Traffic Class: 8
   Priority Flow Control: no
   CSP control ECID handshake complete: yes
   CSP control ECID: 1195
   CSP Alternate control ECID: 1196
```

```
    Total number of configured ports: 33
    CSP number of create port requests sent: 33
    CSP last Tx Trans ID=3, last Rx Trans ID=10
    ECP txErrors=0, sequence=14 firstSeq=14 lastSeq=13 firstAckIdx=0 ackIdx=0
    Next PE: None
    Previous PE: None
    Local CSP Major version is 1 Minor version 1
    Peer  CSP Major version is 1 Minor version 1
    Oper  CSP Major version is 1 Minor version 1
```

Link Aggregation Group (LAG) Formation

FIGURE 9 LAG Between the Control Bridge and Core Devices

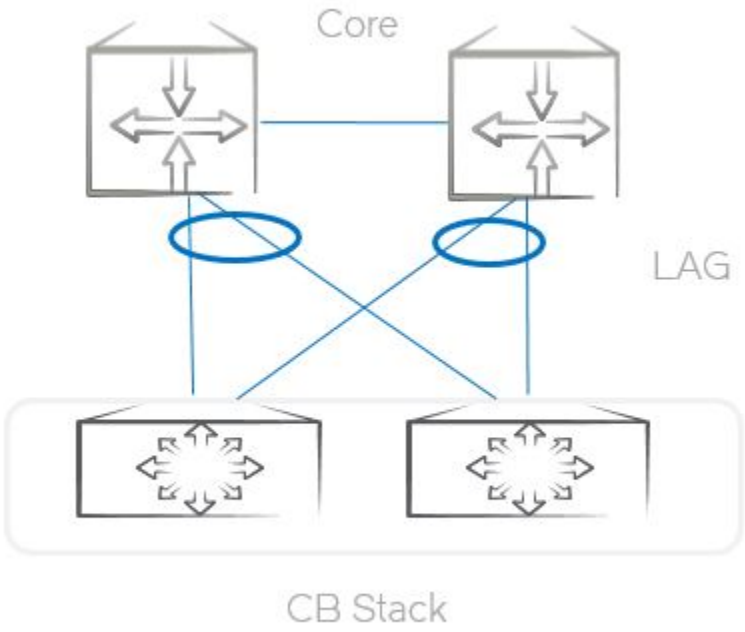

It is a good practice to bundle multiple links between two devices into a LAG wherever possible because this results in high bandwidth along with the benefits of fault tolerance. The following example outlines the commands required to provision the LAG between the CB and the core device as detailed in the topology above.

Configure the following on the CB stack to bring up the LAG.

```
lag dynCB2to301 dynamic id 2
 ports ethernet 1/1/12 ethernet 2/1/12
 primary-port 1/1/12
 deploy
```

Configure the following on the core side to bring up the LAG on the other end.

```
lag "dyn301toCB2" dynamic id 5
 ports ethernet 1/7 to 1/8
```

```
primary-port 1/7
deploy
```

Verification

```
CF_SW4#show lag brief
Total number of LAGs:           3
Total number of deployed LAGs: 3
Total number of trunks created:3 (253 available)
LACP System Priority / ID:      1 / cc4e.24d0.5e80
LACP Long timeout:              120, default: 120
LACP Short timeout:             3, default: 3
LAG             Type    Deploy Trunk Primary  Port List
dynCB2to301     dynamic  Y     2     1/1/12   e 1/1/12 e 2/1/12
dynCB2to302     dynamic  Y     1     1/1/11   e 1/1/11 e 2/1/11
spx_lag_3072    spx      Y     3072  1/1/3    e 1/1/3 to 1/1/4 e 2/1/3 to 2/1/4

CF_SW4#sh lag dynCB2to301
Total number of LAGs:           3
Total number of deployed LAGs: 3
Total number of trunks created:3 (253 available)
LACP System Priority / ID:      1 / cc4e.24d0.5e80
LACP Long timeout:              120, default: 120
LACP Short timeout:             3, default: 3
=== LAG "dynCB2to301" ID 2 (dynamic Deployed) ===
LAG Configuration:
   Ports:          e 1/1/12 e 2/1/12
   Port Count:     2
   Primary Port:   1/1/12
   Trunk Type:     hash-based
   LACP Key:       20002
Deployment: HW Trunk ID 2
Port       Link    State    Dupl Speed Trunk Tag Pvid Pri MAC              Name
1/1/12     Up      Forward  Full 1G    2     No  2301 0   cc4e.24d0.5e80
2/1/12     Up      Forward  Full 1G    2     No  2301 0   cc4e.24d0.5e80
Port       [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp][Ope]
1/1/12          1        1   20002   Yes   L   Agg  Syn  Col  Dis  No   No   Ope
2/1/12          1        1   20002   Yes   L   Agg  Syn  Col  Dis  No   No   Ope
 Partner Info and PDU Statistics
Port          Partner          Partner     LACP      LACP
              System ID         Key      Rx Count  Tx Count
1/1/12   1-001b.ed38.0ec1      103     20261     20229
2/1/12   1-001b.ed38.0ec1      103     20242     20218
```

Each LAG can be verified, using the LAG name, to be in deployed state with links in operational state

VLAN Configuration

The following configuration outlines the procedure to create a VLAN and assign the ports (on the CB and PE) to the VLAN as tagged and untagged members.

```
vlan 2202
 tagged ethe 1/1/48

vlan 2201
 tagged ethe 1/1/48

vlan 2301
 untagged ethe 1/1/12 ethe 2/1/12

vlan 2302
 untagged ethe 1/1/11 ethe 2/1/11

interface eth 31/1/10
 max-vlan 16
```

Assigning tagged and untagged ports under the VLAN as per the configuration requirements

Under the PE port configuration by default maximum 4 numbers of VLANs are supported which can be increased to 16 using "max-vlan" command.

```
Verification

CF_SW4#sh vlan 2201
Total PORT-VLAN entries: 26
Maximum PORT-VLAN entries: 64

Legend: [Stk=Stack-Id, S=Slot]

PORT-VLAN 2201, Name [None], Priority level0,
Spanning tree Off
 Untagged Ports: None
   Tagged Ports: (U1/M1)  48
   Uplink Ports: None
 DualMode Ports: None
 Mac-Vlan Ports: None
     Monitoring: Disabled

CF_SW4#sh vlan 2301
Total PORT-VLAN entries: 26
Maximum PORT-VLAN entries: 64

Legend: [Stk=Stack-Id, S=Slot]

PORT-VLAN 2301, Name [None], Priority level0,
Spanning tree Off
 Untagged Ports: (U1/M1)  12
 Untagged Ports: (U2/M1)  12
   Tagged Ports: None
   Uplink Ports: None
 DualMode Ports: None
 Mac-Vlan Ports: None
     Monitoring: Disabled
```

VLANs are verified for port assignment as tagged/untagged as per the configuration applied

Enabling Unicast Routing

VE Interface Configuration

In the SPX infrastructure, no IP addressing is possible on the physical PE ports; however, the ports can be made a part of virtual interfaces (VEs) to circumvent this limitation.

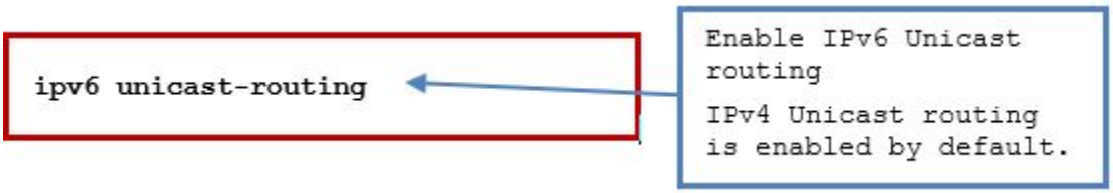

Enable the VE interface under the desired VLAN.

```
vlan 2301
 router-interface ve 2301
```

Under the VE interface created, configure the IPv4 and IPv6 addresses.

```
interface ve 2301
 ip address 172.25.231.1/24
 ipv6 address fdd7:b215:19cb:4552:172:25:231:1/112
```

Loopback Interface Configuration

Each device in the Campus Fabric needs one loopback interface with a unique IPv4 address for the purpose of the router ID.

```
interface loopback 1
 ip address 22.1.1.1 255.255.255.255
 ipv6 address fdd7:b215:19cb:4552:22:1:1:1/128
```

OSPFv2 (IPv4) Configuration Details

OSPFv2 is configured in a single OSPF area (area 0) across the SPX infrastructure.

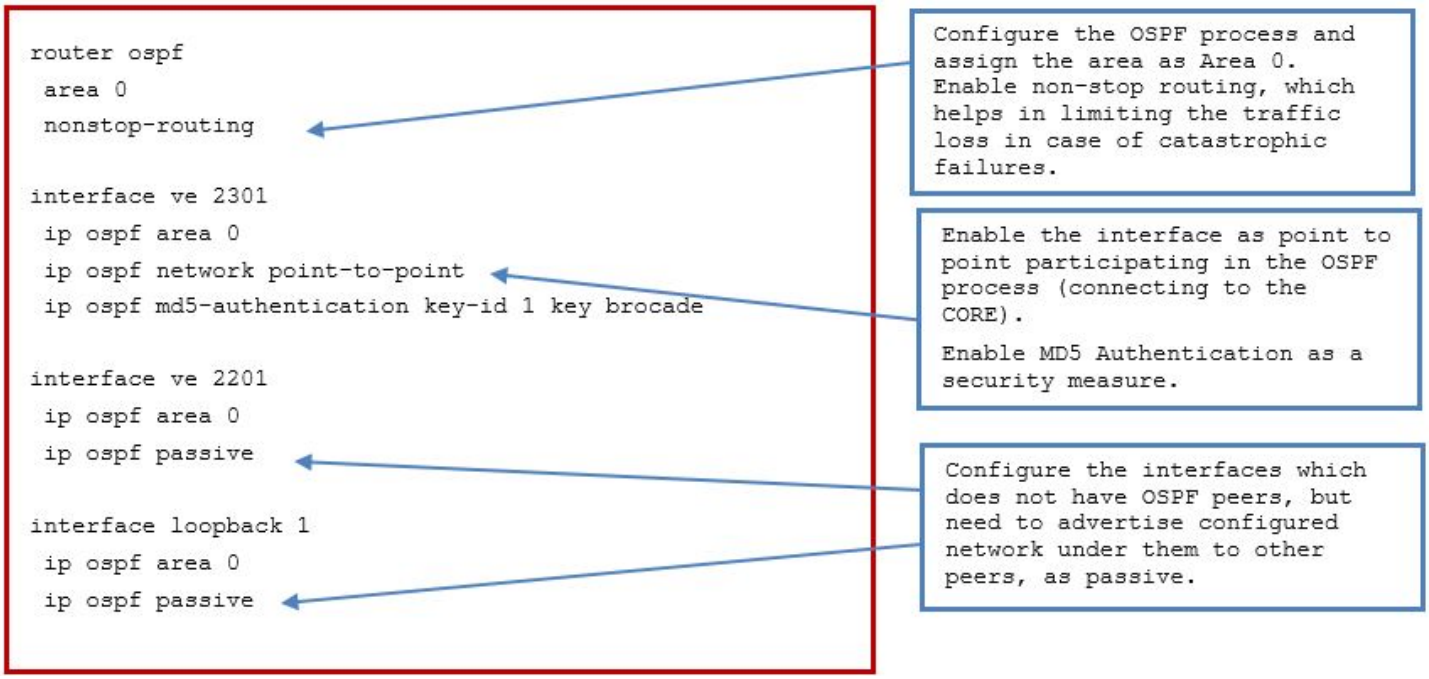

OSPFv3 (IPv6) Configuration Details

OSPFv3 is configured in a single OSPF area (area 0) across the SPX infrastructure.

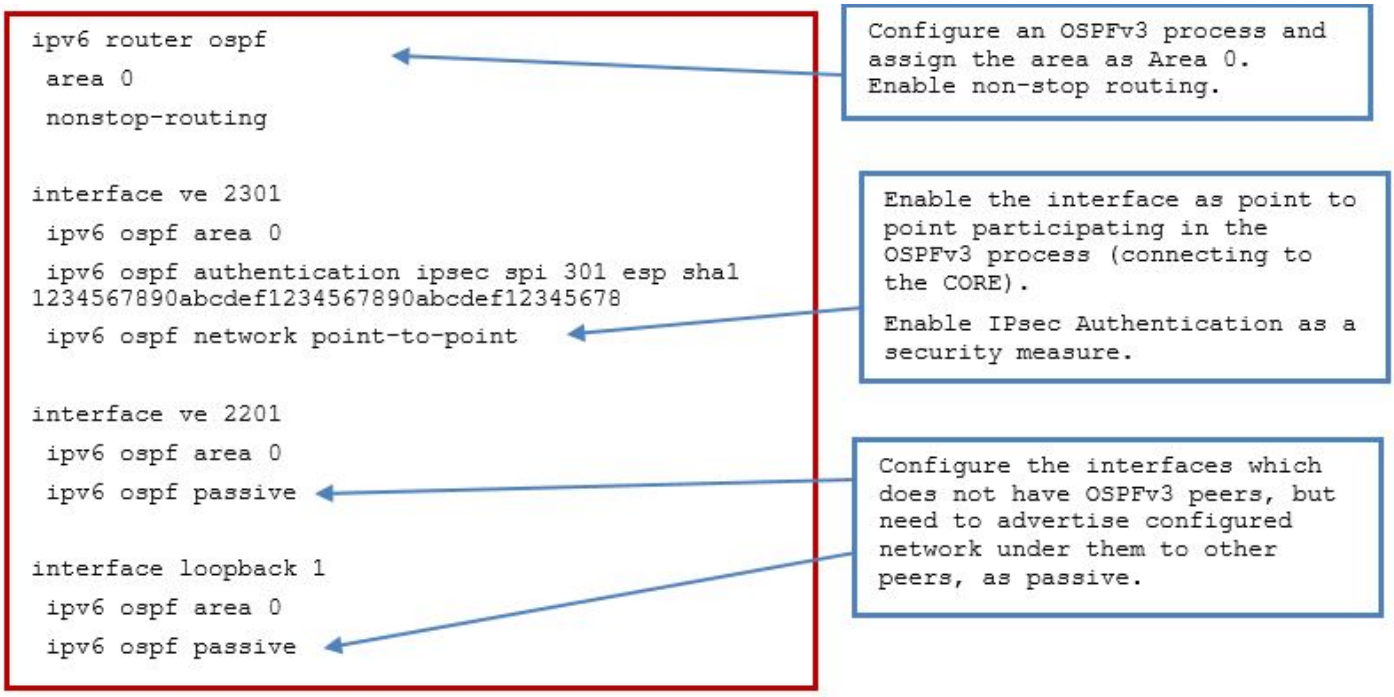

In the SPX infrastructure, OSPFv2 and OSPFv3 share neighborship between the CB devices and core devices, as verified in the following example. In each case, there are two OSPF neighbors.

```
Verification :

CF_SW4#sh ip ospf neigh
Number of Neighbors is 2, in FULL state 2

Port          Address          Pri State      Neigh Address    Neigh ID         Ev Opt Cnt
v2301          172.25.231.1     1   FULL/OTHER 172.25.231.2     10.10.10.1       7  2   0
v2302          172.25.232.1     1   FULL/OTHER 172.25.232.2     10.10.10.2       5  2   0

CF_SW4#sh ipv6 ospf neigh

Total number of neighbors in all states: 2
Number of neighbors in state Full      : 2

RouterID         Pri State     DR              BDR             Interface      [State]
10.10.10.1         1 Full      0.0.0.0         0.0.0.0         ve 2301        [P2P]
10.10.10.2         1 Full      0.0.0.0         0.0.0.0         ve 2302        [P2P]
```

Enabling Multicast

Enabling IGMP Snooping and MLD Snooping

Configure IGMP snooping and MLD snooping on a per-VLAN basis wherever multicast receivers are connected to the CB and PE ports (access side).

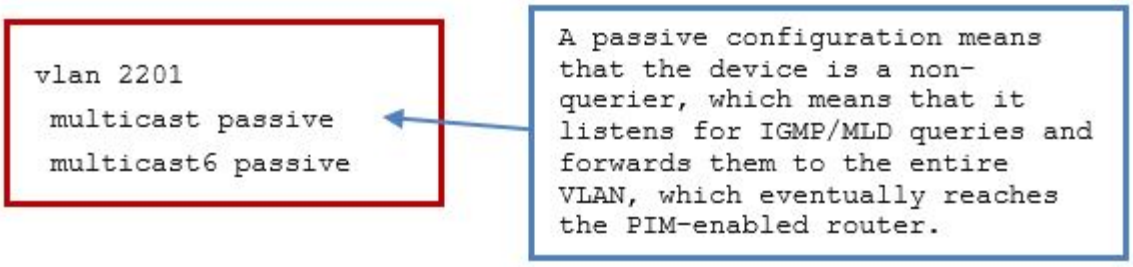

Enabling IPv4 and IPv6 PIM Sparse Protocols on CB

The following configuration is required to bring up the IPv4 and IPv6 PIM neighborship between the CB and the core devices.

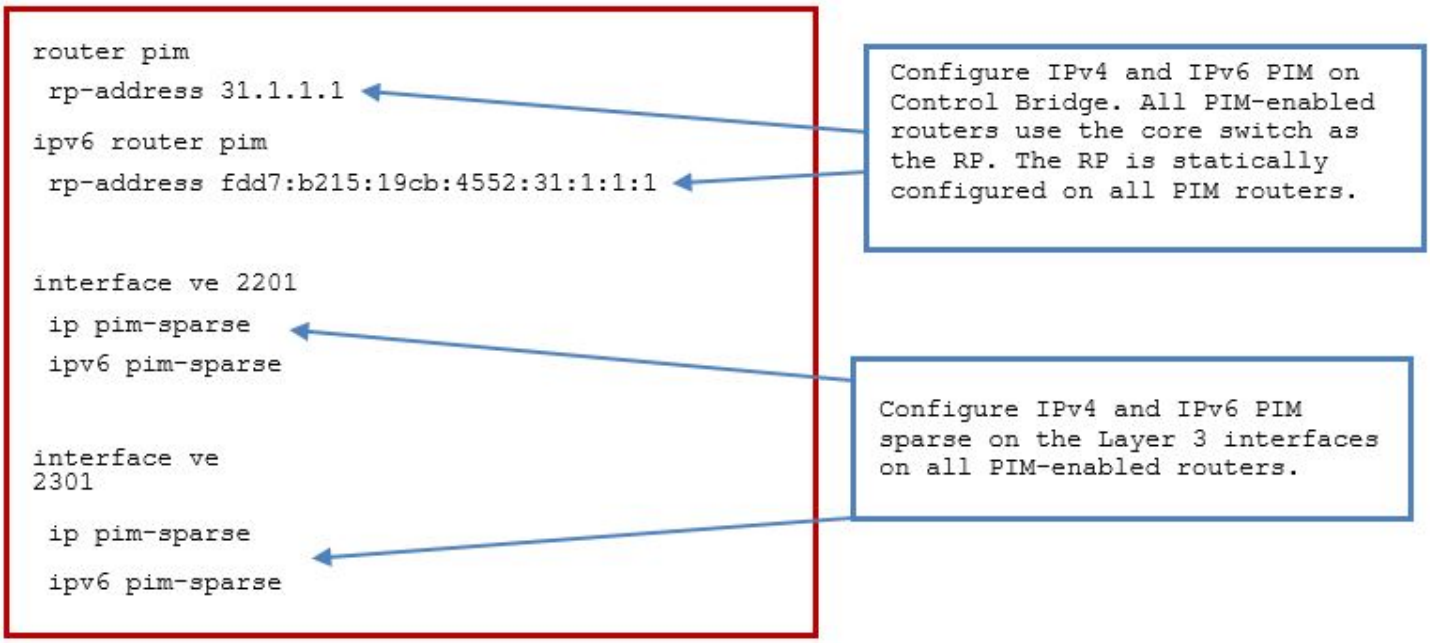

Enabling IPv4 and IPv6 PIM Sparse Protocols on Core Devices

Configure the loopback address on the core device to be used as a static RP.

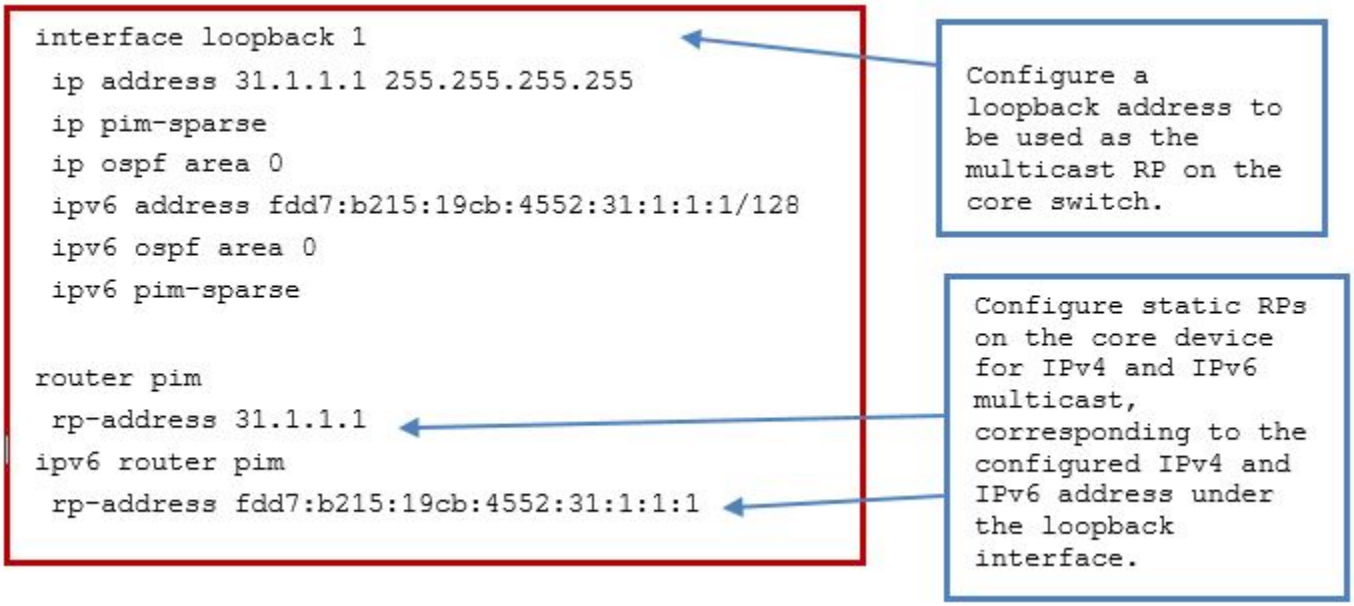

Verifying PIM Neighbors and Multicast Mcache Table

```
! Show command to verify the RP configuration.
CORE-301# show ip pim rp-set
Static RP
```

```
---------
Static RP count: 1
31.1.1.1
Number of group prefixes Learnt from BSR: 0
No BSR RP-Set present.

CORE-301# show ipv6 pim rp-set
Static RP
---------
Static RP count: 1
fdd7:b215:19cb:4552:31:1:1:1
Number of group prefixes Learnt from BSR: 0
No BSR RP-Set present.
```

```
! Show command to verify PIM neighbors.
CORE-301# show ip pim neighbor

--------+--------+---------------+--------+---+---------+---------+-----+----------------+-----------+----+
Port    |PhyPort |Neighbor       |Holdtime|T  |PropDelay|Override |Age  |UpTime          |VRF        |Prio
        |        |               |sec     |Bit|msec     |msec     |sec  |                |           |
--------+--------+---------------+--------+---+---------+---------+-----+----------------+-----------+----+
v131     e1/5     172.25.131.1    105      1   500       3000      19    1d 14:59:31      default     1
v231     e1/7     172.25.231.1    105      1   500       3000      30    2d 10:12:11      default     1

Total Number of Neighbors : 2

! Show command to verify multicast mcache.
CORE-301# show ip pim mcache
IP Multicast Mcache Table
Total entries in mcache: 2

1      (*, 225.1.1.1) RP 31.1.1.1, in NIL (NIL), Uptime 08:50:07 (SM)
       No upstream neighbor because RP 31.1.1.1 is itself
       Flags (0x00220480) SM RPT
       slow ports: ethe 1/5 ethe 1/7
       AgeSltMsk: 00000000,  FID: NotReq,  DIT: NotReq, RegPkt: 0, profile: none
       Forwarding_oif: 2, Immediate_oif: 2, Blocked_oif: 0
       L3 (SW) 2:
           TR(e1/5,e1/5)(VL1301), 08:26:34/180, Flags: IM
           TR(e1/7,e1/7)(VL2301), 08:50:07/172, Flags: IM

2      (172.25.151.2, 225.1.1.1) in v131 (e1/5), Uptime 09:19:47 (SM)
       upstream neighbor 172.25.131.1
       Flags (0x7002c0c1) SM SPT HW FAST MSDPADV
       fast ports:
       AgeSltMsk: 00000000,  FID: 0xffff (D),  DIT: NotReq, RegPkt: 0, profile: none, KAT Timer value: 160
       Forwarding_oif: 0, Immediate_oif: 0, Blocked_oif: 2
       Blocked OIF 2:
           TR(e1/5,e1/5)(VL1301), 08:26:34/180, Flags: IH BR
           TR(e1/7,e1/7)(VL2301), 08:50:07/172, Flags: IH BR

Number of matching entries: 2
```

Configuring Device Access Security

The following configuration is used to provision the redundant RADIUS server and AAA services on the CB device.

Configure redundant RADIUS server details and enable Telnet access.

```
radius-server host 10.20.12.33 auth-port 1812 acct-port 1813 default key brocade123 dot1x mac-auth web-auth

radius-server host 10.32.12.33 auth-port 1812 acct-port 1813 default key brocade123 dot1x
```

```
enable telnet authentication
CB_SW4(config)# web-management https
```

In the following configuration, RADIUS is the primary authentication method for securing access to the device. If authentication fails due to an error with the RADIUS server, authentication is performed using local user accounts. Local user accounts must be created on individual switches before using the local authentication method as fallback authentication. Local user accounts can be created using the **username xxxx password xxxx** command.

```
aaa authentication login default radius local
aaa authentication enable default radius local
aaa authentication web-server default radius local
```

Configure SSH2 on all devices. The first time while provisioning the device, generate a host DSA or RSA and a private key pair.

```
CB_SW4(config)# crypto key generate rsa modulus 2048

Creating RSA key pair, please wait...

Download request from active unit 2 mac = cc4e.24d0.7580
Downloading - $$ssh8rsahost.key

stack done ssh rsa host key sync Done.


RSA Key pair is successfully created

CB_SW4(config)# show ip ssh

Connection      Version Encryption      Username        HMAC    Server
Hostkey IP Address

Inbound:

1       SSH-2   aes256-ctr      brocade hmac-sha1       ssh-rsa
10.37.224.233

Outbound:

 SSH-v2.0 enabled; hostkey: RSA(2048)
```

Now any SSH client can be used to securely log in (over the encrypted session) to the device after the user provides the required credentials.

Configuring Network Access Security

The following configuration is required to enable network access policies on the SPX infrastructure. For more details on network access security, refer to Network Access Security on page 31.

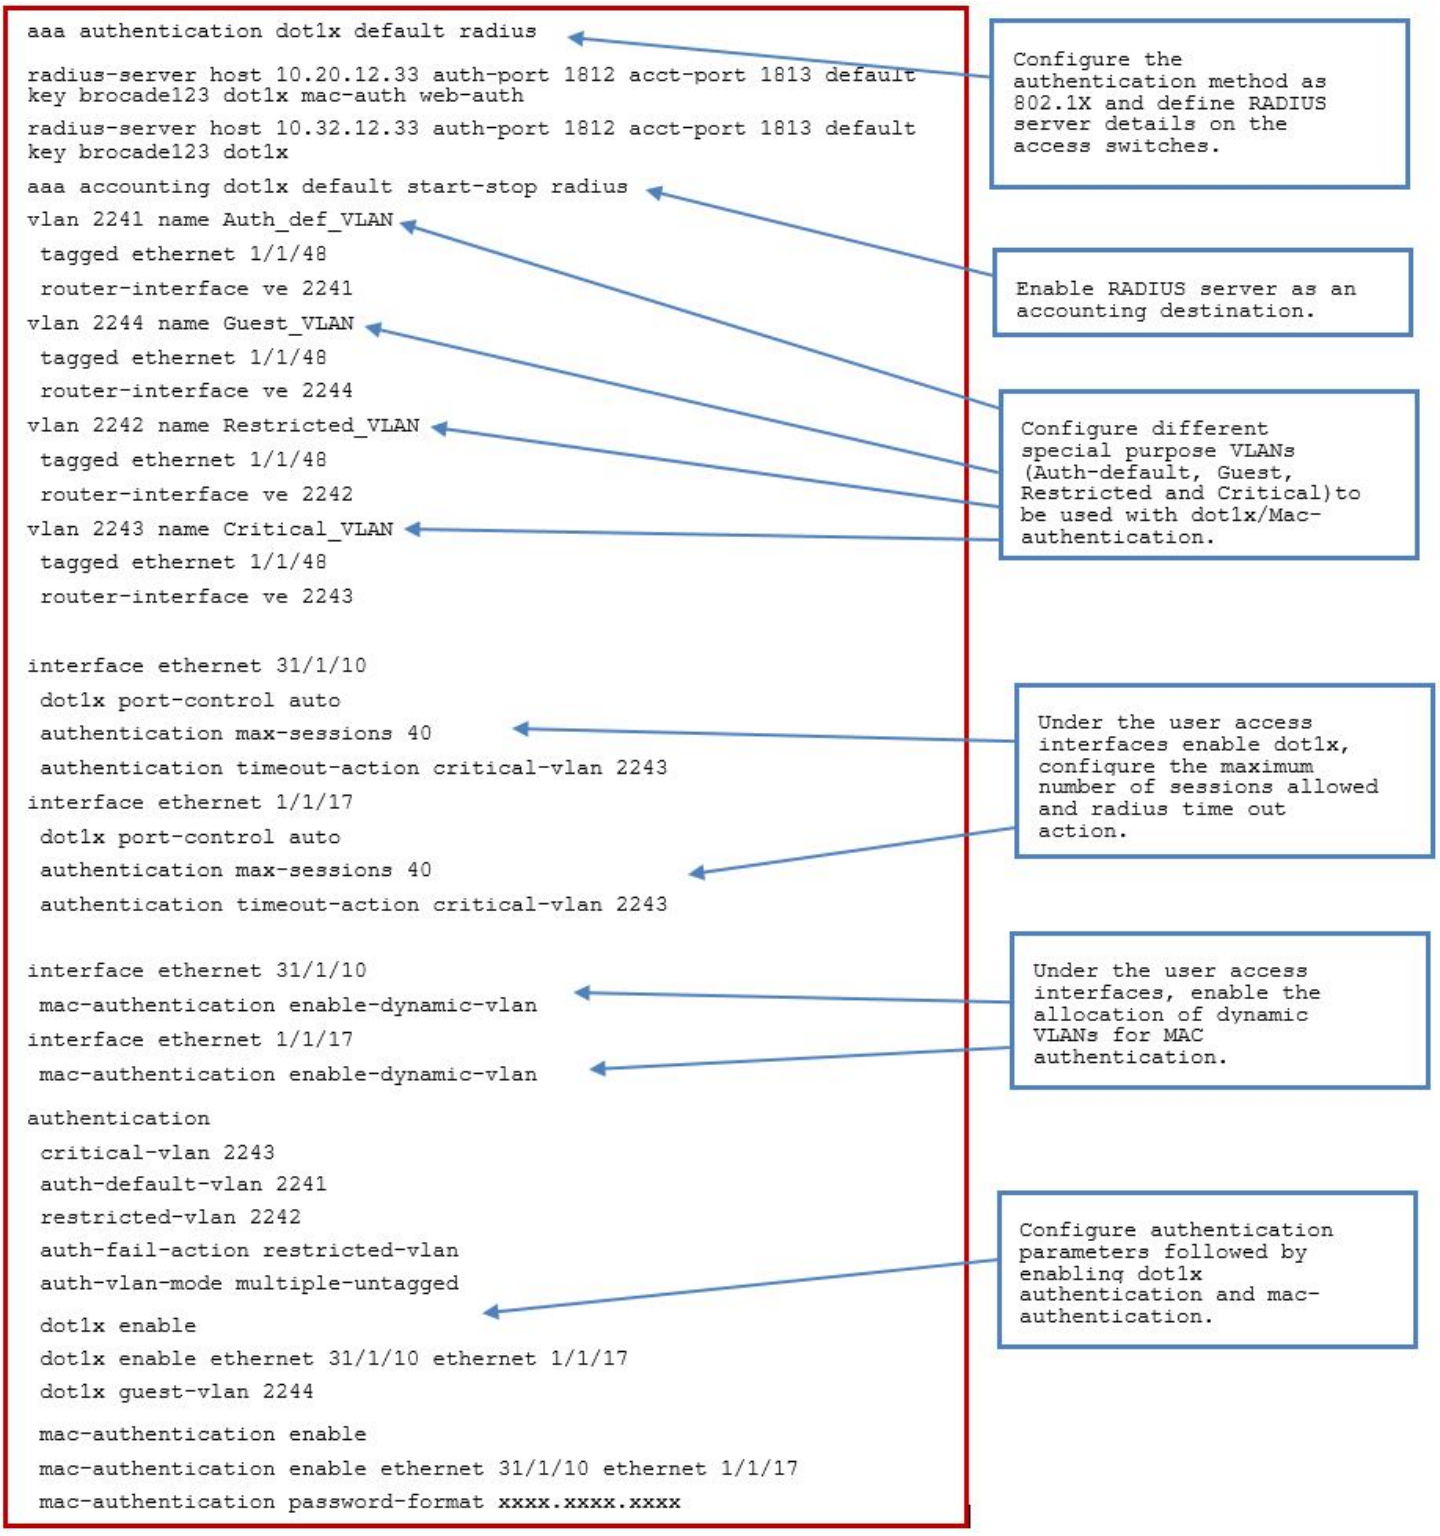

Access Control Lists (ACLs) can be configured and applied under the desired interfaces using the following configurations.

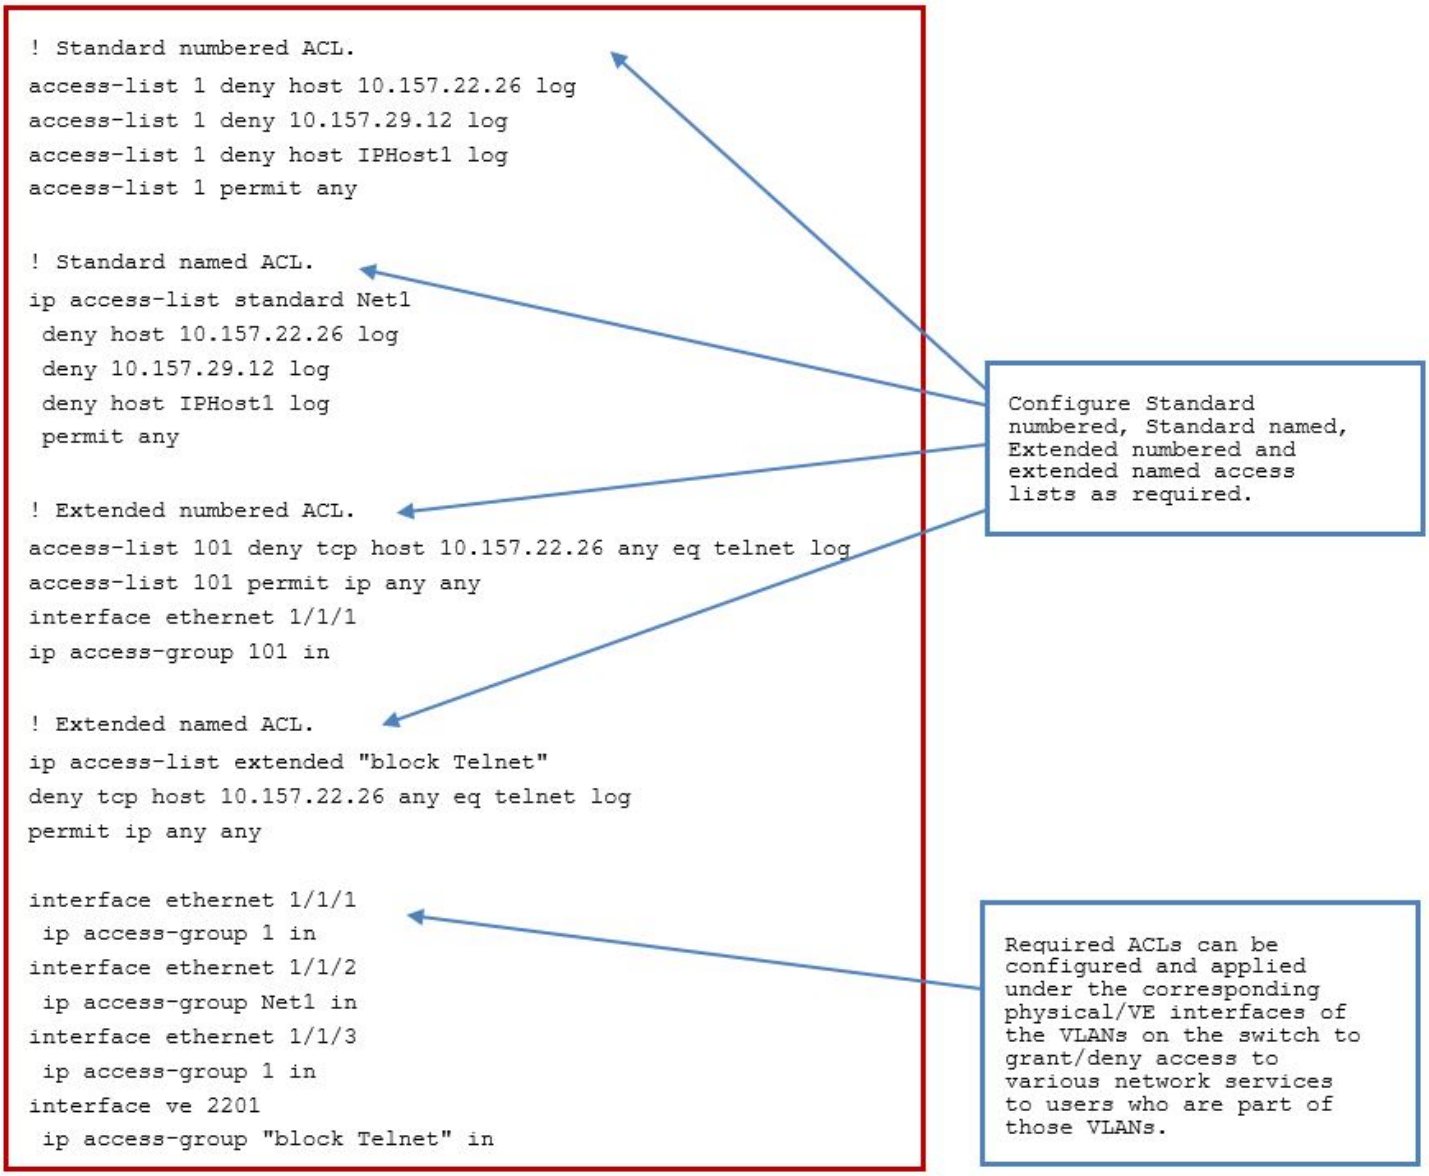

Enabling Management Features

Configuring Power over Ethernet (PoE)

Access devices such as VoIP phones, Wireless APs, and so on connect to the ports on the PE. These ports are provisioned to supply the required power. The following configuration enables the inline power on a PE port.

```
interface ethernet 31/1/1
 inline power
interface ethernet 31/1/2
 inline power
```

Verifying PoE on the device.

```
CF_SW4# show inline power
Power Capacity:    Total is 748000 mWatts. Current Free is 735400 mWatts. Power Allocations:    Requests
Honored 9 times
Port    Admin    Oper    ---Power(mWatts)---  PD Type  PD Class  Pri  Fault/ State    State
Consumed  Allocated    Error
-----------------------------------------------------------------------------
31/1/1    On    On    3500    6300  802.3af  Class 2    3    n/a
31/1/2    On    On    3600    6300  802.3af  Class 2    3    n/a
Total    7100    12600
```

Configuring and Verifying Network Time Protocol (NTP)

```
CF_SW4(config)# ntp
CF_SW4(config-ntp)# server 10.31.2.80
CF_SW4(config-ntp)# end
! Configuring the local clock and timezone on the device.
CF_SW4(config)# clock summer-time
CF_SW4(config)# clock timezone us Pacific
! Show command to verify NTP operation.
CF_SW4# show ntp status
Clock is synchronized, stratum 5, reference clock is 10.31.2.80
 precision is 2**-16
 reference time is 3689689157.504644503 (09:39:17.504644503 Pacific Fri Dec 02 2016)
 clock offset is 0.3389 msec, root delay is 265.1191 msec
 root dispersion is 200.5500 msec,  peer dispersion is 197.2204 msec
 system poll interval is 64,  last clock update was 410 sec ago
 NTP server mode is enabled, NTP client mode is enabled
 NTP master mode is disabled, NTP master stratum is 8
 NTP is not in panic mode
```

Configuring DHCPv4 and DHCPv6

The DHCPv4 or DHCPv6 client that initiates the DHCPv4 or DHCPv6 transaction connects to the Layer 2 switch on a port assigned to a specific VLAN. The IP address of the DHCP server is configured as the IP helper address and IPv6 DHCP-relay destination under the corresponding Layer 3 interface.

```
interface ve 2202
 ip helper-address 1 172.25.10.32
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32

! Note: It is important that the DHCPv4 and DHCPv6 servers be running and be reachable. Also, the IP
addresses assigned to Layer 3 interfaces should be included in the exclusion list in DHCP servers.
```

Configuring and Verifying SNMP

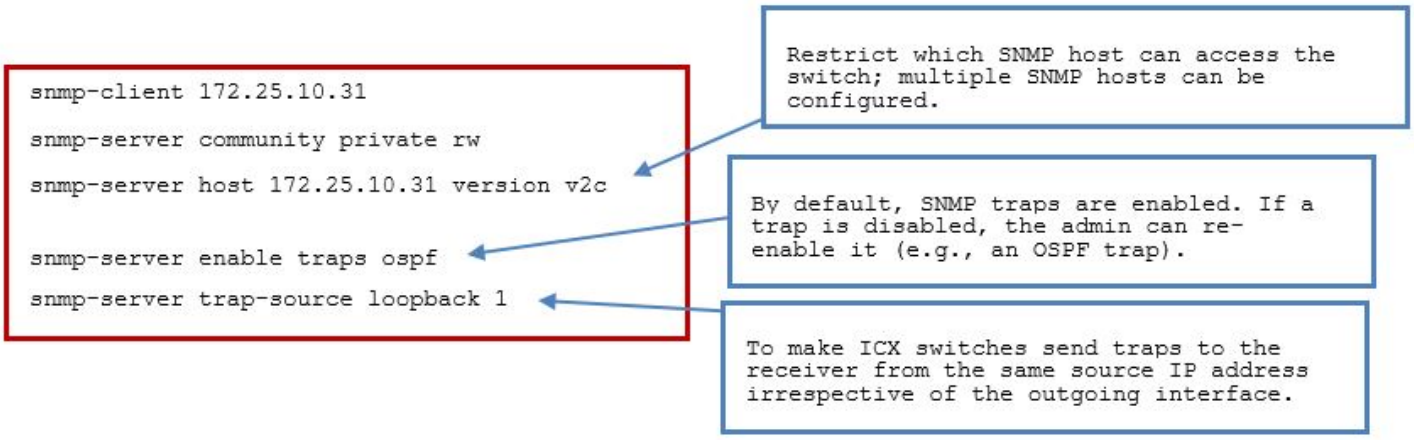


```
CF_SW4# show snmp server

       Status: Enabled

      Contact:
     Location:
Community(ro): .....
Community(rw): .....

Max Ifindex per module: 64

Traps
                    Cold start: Enable
                       Link up: Enable
                     Link down: Enable
                Authentication: Enable
          Power supply failure: Enable
                    Fan failure: Enable
               Fan speed change: Enable
                Module inserted: Enable
                Module removed: Enable
Redundant module state change: Enable
            Temperature warning: Enable
                  STP new root: Enable
             STP topology change: Enable
                MAC notification: Enable
          MAC-AUTH notification: Enable
                           OSPF: Enable
                            BGP: Enable
                           VRRP: Enable
                           VSRP: Enable
                            MRP: Enable
                           UDLD: Enable
                      link-oam: Enable
                            cfm: Enable
                         syslog: Disable
              entity-cfg-change: Enable

 Total Trap-Receiver Entries: 2

Trap-Receiver IP-Address                Version    Port-Number Comm-or-Security
    1         10.32.12.31               v2c            162     .....                                2
172.25.10.31              v2c           162      .....
```

Configuring and Verifying sFlow

The following configuration enables and verifies sFlow on a CB device.

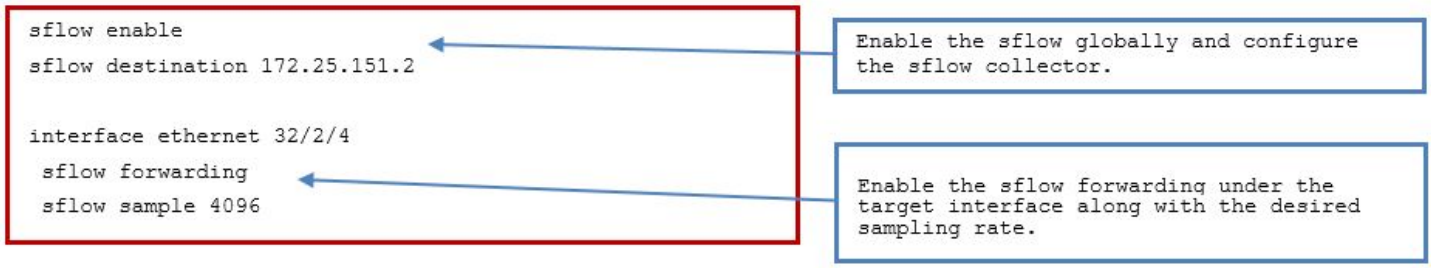


```
CF_SW4# show sflow
sFlow version: 5
sFlow services are enabled.

sFlow agent IP address: 22.1.1.1
sFlow source IP address: unspecified
sFlow source IPv6 address: unspecified
Collector IP 172.25.151.2, UDP 6343, Configured VRF: None, Using VRF: default-vrf
UDP source port: 8888 (Default)
Polling interval is 20 seconds.
Configured default sampling rate: 1 per 4096 packets.
Actual default sampling rate: 1 per 4096 packets.
The maximum sFlow sample size: 128.
Sample mode: All packets including dropped packets
sFlow exporting cpu-traffic is disabled.
258 UDP packets exported
0 sFlow flow samples collected.
sFlow ports: ethe 32/2/4
Module Sampling Rates
---------------------
U32:M2 configured rate=4096, actual rate=4096
Port Sampling Rates
-------------------
Port=32/2/4, configured rate=4096, actual rate=4096
```

Enabling Quality of Service (QoS)

```
interface ethernet 1/2/4
 trust dscp
ip access-list extended MARK-DVLAN-TRAFFIC

Clause to match the UDP port for Lync services.

 permit udp any any range 30000 30100 dscp-matching 46 dscp-
marking 46 802.1p-and-internal-marking 5
Clause to match the TCP port of Skinny client (IP phone, SCCP)
service.
 permit tcp any any eq 2000 dscp-marking 46 802.1p- and-
internal-marking 5
Clauses to match the TCP/UDP port of H.323 services.
 permit tcp any any range 1719 1720 dscp-marking 46 802.1p-and-
internal-marking 5
 permit udp any any range 1719 1720 dscp-marking 46 802.1p-and-
internal-marking 5
Clauses to match the TCP/UDP port of MS-ICCP (Audio Call Control
Protocol, NetMeeting) services.
 permit tcp any any eq 1731 dscp-marking 46 802.1p- and-
internal-marking 5
 permit udp any any eq 1731 dscp-marking 46 802.1p- and-
internal-marking 5
Clauses to match the TCP/UDP port for SIP services.
 permit tcp any any eq 5060 dscp-marking 46 802.1p- and-
internal-marking 5
 permit udp any any eq 5060 dscp-marking 46 802.1p- and-
internal-marking 5
Clauses to re-mark the voice-control/video-control packets that
normally arrive marked with DSCP 24 to 46.
 permit udp any any dscp-matching 24 dscp-marking 46 802.1p-and-
internal-marking 5
 permit tcp any any dscp-matching 24 dscp-marking 46 802.1p-and-
internal-marking 5
Clause to match the UDP port for Cisco unified services.
 permit udp any any eq 5445 dscp-marking 34 802.1p- and-
internal-marking 4
Clauses to match the TCP/UDP port for remote file transfer services.
 permit tcp any any dscp-matching 34 dscp-marking 34 802.1p-and-
internal-marking 4
 permit udp any any dscp-matching 34 dscp-marking 34 802.1p-and-
internal-marking 4
Clauses to match the TCP/UDP port for Polycom voice and video
services.
 permit tcp any any range 3230 3247 dscp-marking 34 802.1p-and-
internal-marking 4
CF SW4(config-ext-nacl)# permit udp any any range 3230 3247
dscp-marking 34 802.1p-and-internal-marking 4
```

Configure Trust DSCP on the ingress ports and Configure an extended access list for traffic marking

```
Clauses to re-mark the traffic arriving with different COS values with
appropriate DSCP values.
 permit ip any any 802.1p-priority-matching 0 dscp- marking 0 802.1p-
and-internal-marking 0
 permit ip any any 802.1p- priority-matching 1 dscp- marking 8 802.1p-
and-internal-marking 1
 permit ip any any 802.1p-priority-matching 2 dscp- marking 16 802.1p-
and-internal-marking 2
 permit ip any any 802.1p-priority-matching 3 dscp- marking 24 802.1p-
and-internal-marking 3
 permit ip any any 802.1p-priority-matching 4 dscp- marking 34 802.1p-
and-internal-marking 4
 permit ip any any 802.1p-priority-matching 5 dscp- marking 46 802.1p-
and-internal-marking 5


Clauses to re-mark the traffic arriving with COS value 6/7 to COS value
5 since COS 6/7 maps internally for control protocols.
 permit ip any any 802.1p-priority-matching 6 dscp- marking 46 802.1p-
and-internal-marking 5
 permit ip any any 802.1p-priority-matching 7 dscp- marking 46 802.1p-
and-internal-marking 5


Permit clause for any other IP traffic.
 permit ip any any


interface ethernet 32/2/4
 per-vlan 2201
  ip access-group Tplay_Trfc in
interface ethernet 32/2/4
 per-vlan 2202
  ip access-group Tplay_Trfc in
```

Apply QoS marking on the ingress interface.

Alternatively the acl can be applied under the Vlan dynamically as it is returned by RADIUS server as a result of authentication parameters.

```
CF_SW1# sh dot1x se all
--------------------------------------------------------------------------------------------------------
Port    MAC              IP(v4/v6)          User          VLAN  Auth      ACL    Session    Age    PAE
        Addr             Addr               Name                State           Time              State
--------------------------------------------------------------------------------------------------------
32/2/4  0010.b232.8001   172.25.204.2       a2204         2204  permit    Yes    30259      Ena
AUTHENTICATED

CF_SW1# sh dot1x ip-acl all
-------------------------------------------------------------------------
Port      MAC Address       V4 Ingress   V4 Egress    V6 Ingress   V6 Egress
-------------------------------------------------------------------------
32/2/4   0010.b119.8001    Tplay_Trfc   -            -            -
```

Illustration Examples

The following examples illustrate the use cases by using sections of the validated design network topology as appropriate. This will help you to further understand the deployment scenarios.

Example 1—Network Admission Control (NAC)

The end user connects to a port on a PE, and is authenticated using dot1x authentictation followed by the assignment of IPv4/IPv6 from DHCP.

FIGURE 10 Dot1x Authentication

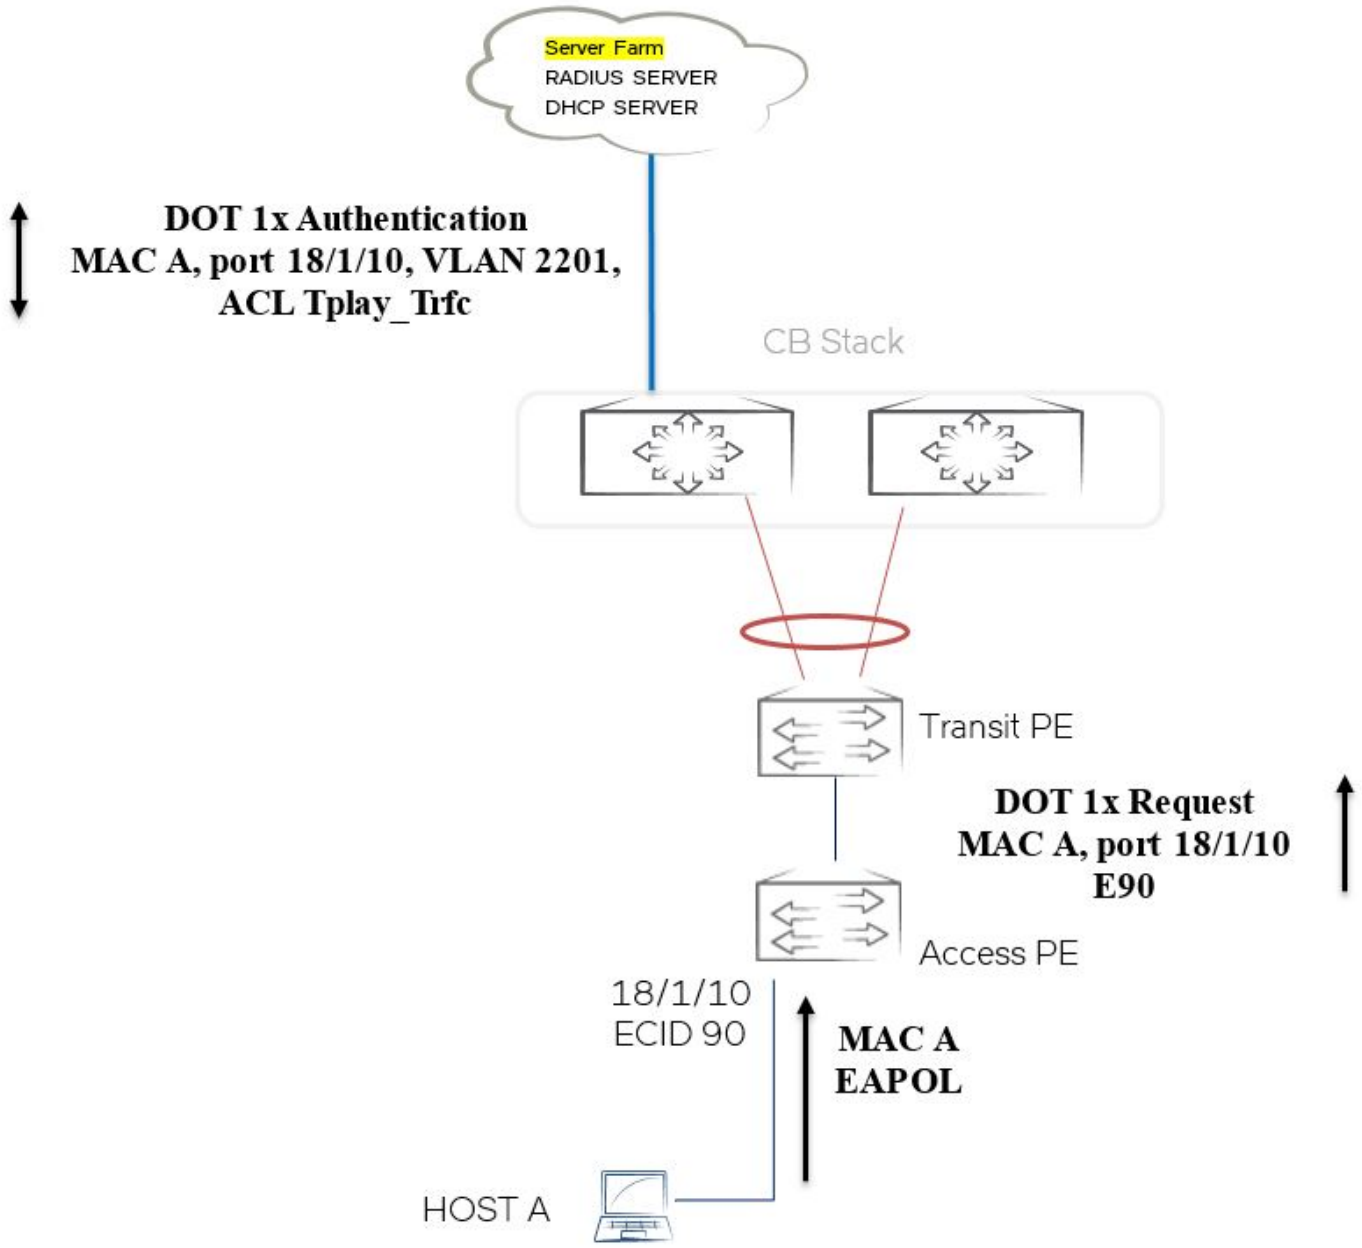

1. The user connects to the dot1x-enabled port of the PE.

```
Global configuration :
authentication
 critical-vlan 2242
 auth-default-vlan 2241
 restricted-vlan 2243
 auth-fail-action restricted-vlan
 auth-vlan-mode multiple-untagged
```

```
 dot1x enable
 dot1x enable ethe 18/1/10 ethe 23/2/4
 dot1x guest-vlan 2244
 mac-authentication enable
 mac-authentication password-format xxxx.xxxx.xxxx
!
aaa authentication dot1x default radius
radius-server host 10.32.12.33 auth-port 1812 acct-port 1813 default key 2 $b24tbz1nIlp8OA== dot1x
mac-auth web-auth


Configuration under the Interface :
interface ethernet 18/1/10
 authentication max-sessions 1024
 authentication timeout-action critical-vlan
 dot1x port-control auto
 max-vlan 5
!
```

2. Dot1x authentication is carried out and the user is authenticated.
   - The user is put into a dynamic VLAN returned by RADIUS.
   - ACL is applied dynamically for the user as returned by RADIUS.

```
CF_SW4#sh dot1x se all
--------------------------------------------------------------------------------------------------------------
Port    MAC              IP(v4/v6)           User          VLAN  Auth      ACL    Session   Age   PAE
        Addr             Addr                Name                State           Time            State
--------------------------------------------------------------------------------------------------------------
18/1/10 0010.9411.1101   N/A                 a2201         2201  permit    Yes    523       Ena
AUTHENTICATED

CF_SW4#show dot1x ip-acl  all
-------------------------------------------------------------------------
Port      MAC Address       V4 Ingress   V4 Egress    V6 Ingress   V6 Egress
-------------------------------------------------------------------------
18/1/10   0010.9411.1101    Tplay_Trfc   -            -            -
```

FIGURE 11 DHCP Transactions

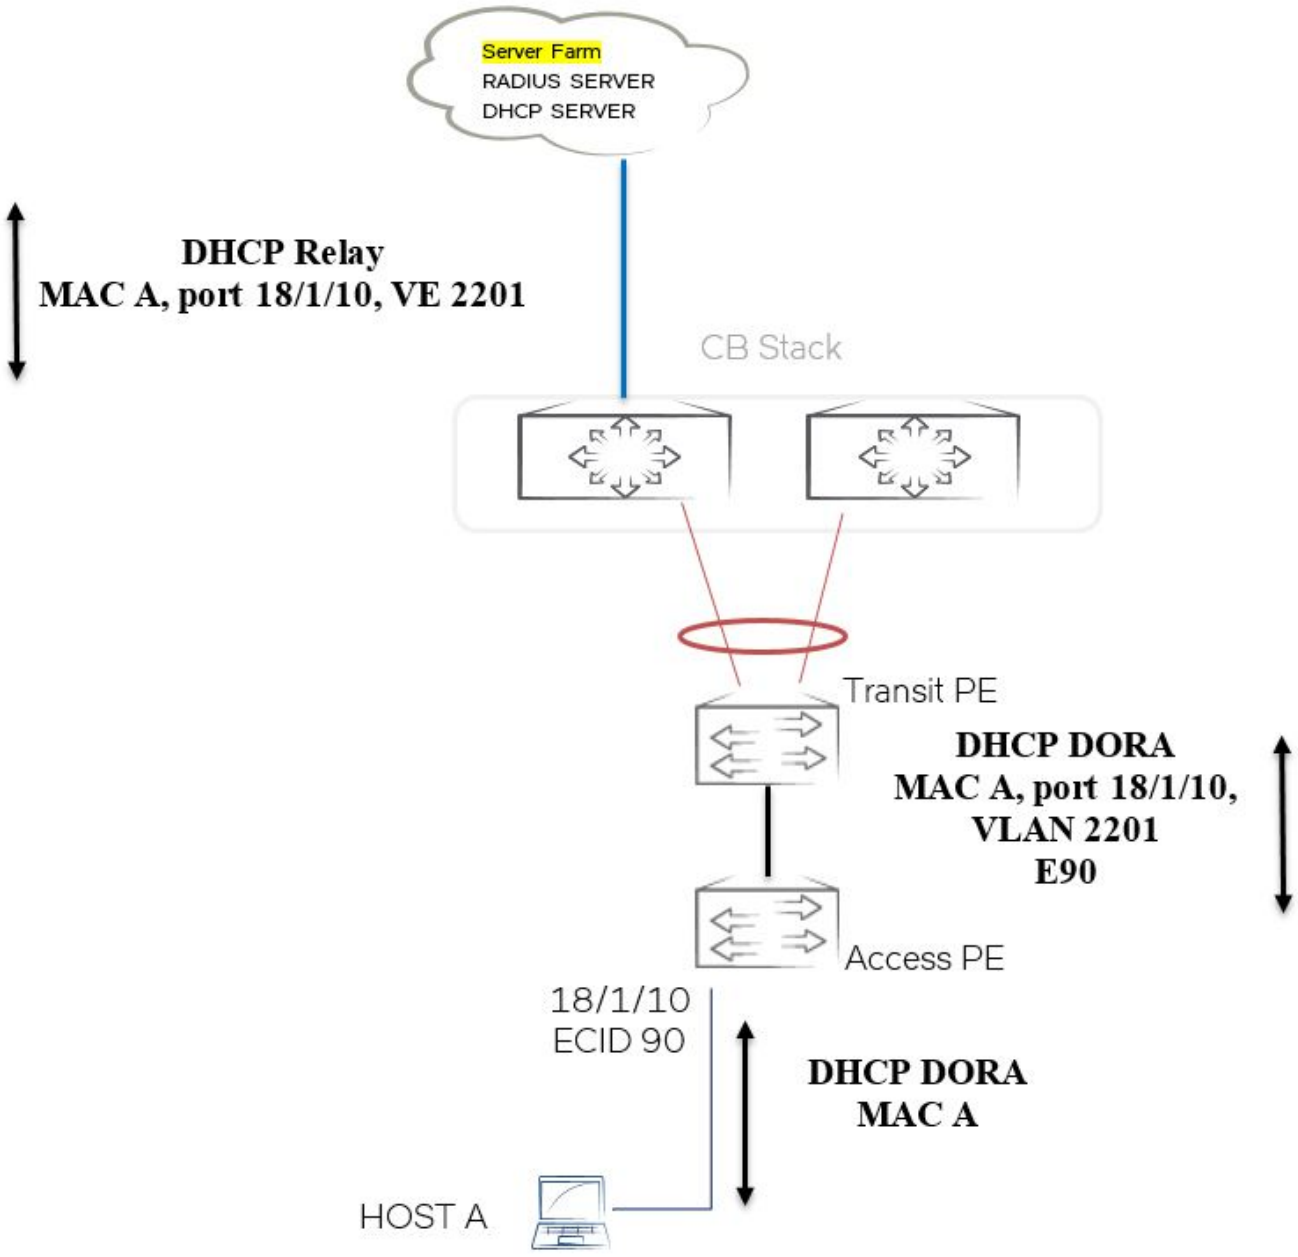

3. The DHCP process is started and, using the DHCP-relay functionality, the DHCP control packets are relayed between the client and the DHCP server. The user is assigned the IP address/IPv6 address.

```
CF_SW4#sh run vlan 2201
vlan 2201 by port
 tagged ethe 1/1/48
 router-interface ve 2201
!
!
CF_SW4#sh run int ve 2201
interface ve 2201
 ip address 172.25.201.1 255.255.255.0
 ip helper-address 1 172.25.10.32
 ipv6 pim-sparse
!
```

4. The user is provided the network services and is ready to communicate over the network.

Example 2—Layer 2 Communication Between End Users

FIGURE 12 Layer 2 Communication

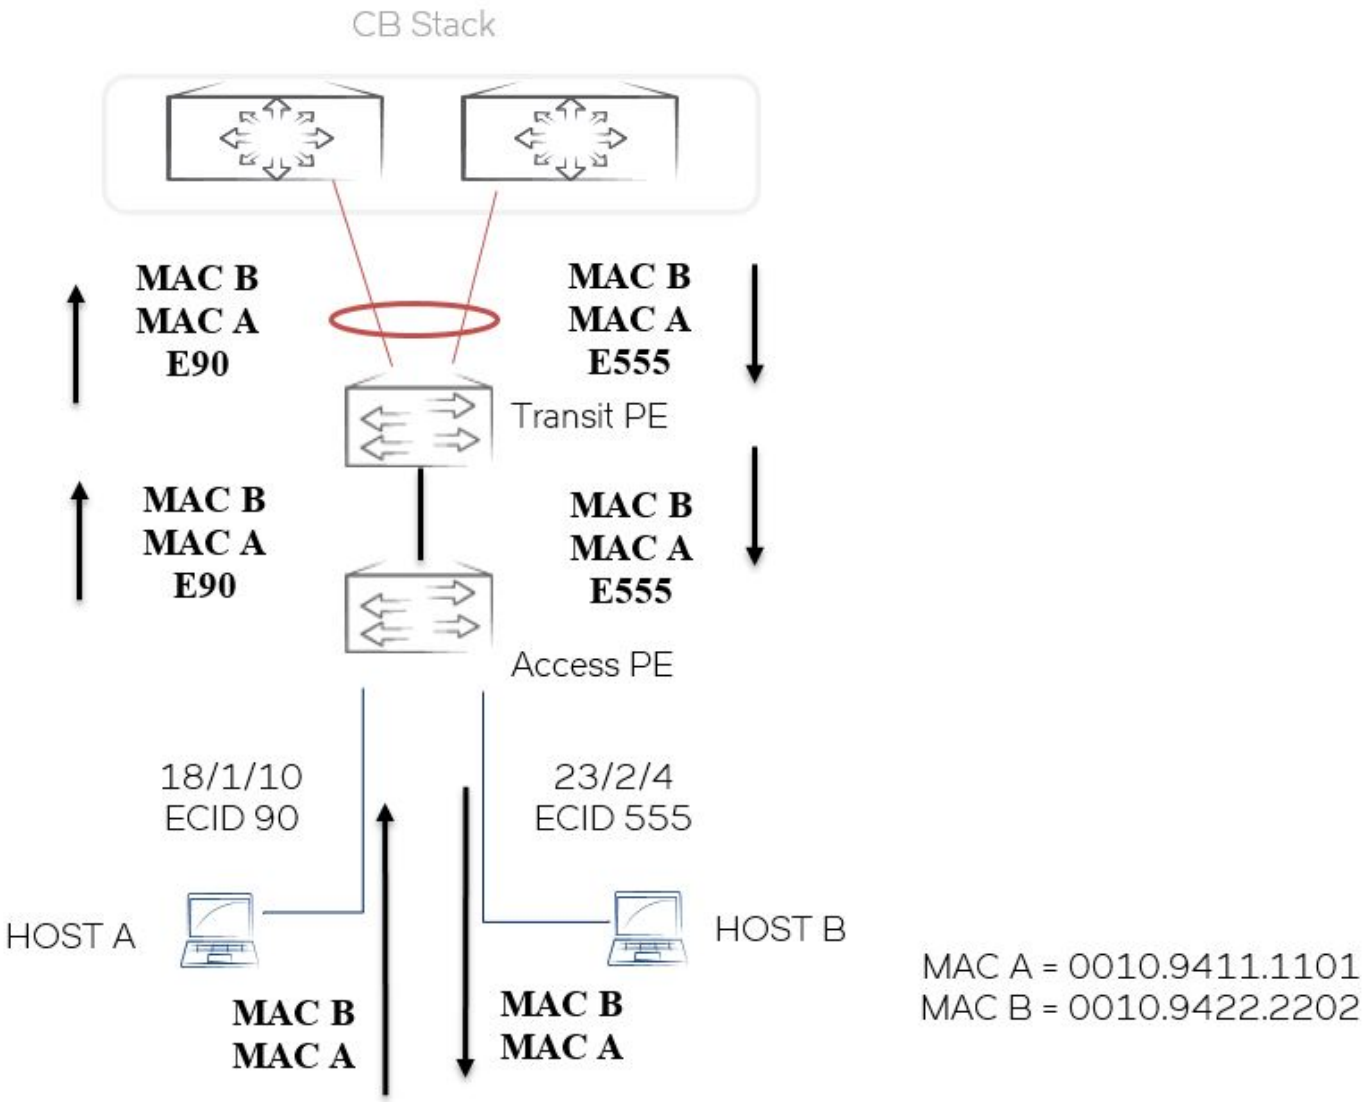

1. The user sends the Layer 2 packet with source MAC address0010.9411.1101 and destination MAC address0010.9422.2202.

```
CF_SW4#sh mac-address
Total active entries from all ports = 17
MAC-Address      Port                Type          VLAN
0010.9411.1101  18/1/10            Dynamic       2201
0010.9422.2202  23/2/4              Dynamic       2201

CF_SW4#sh vlan 2201
Total PORT-VLAN entries: 26
Maximum PORT-VLAN entries: 64
```

```
Legend: [Stk=Stack-Id, S=Slot]

PORT-VLAN 2201, Name [None], Priority level0, Spanning tree Off
 Untagged Ports: (U18/M1) 10
   Tagged Ports: (U1/M1)  48
   Uplink Ports: None
 DualMode Ports: None
 Mac-Vlan Ports: (U18/M1) 10
 Mac-Vlan Ports: (U23/M2) 4
     Monitoring: Disabled
```

2. Because the destination MAC address is learned on port 23/2/4 of the PE, the packet is sent out of port 23/2/4.

Example 3—Layer 3 Communication Between End Users

FIGURE 13 Layer 3 Communication

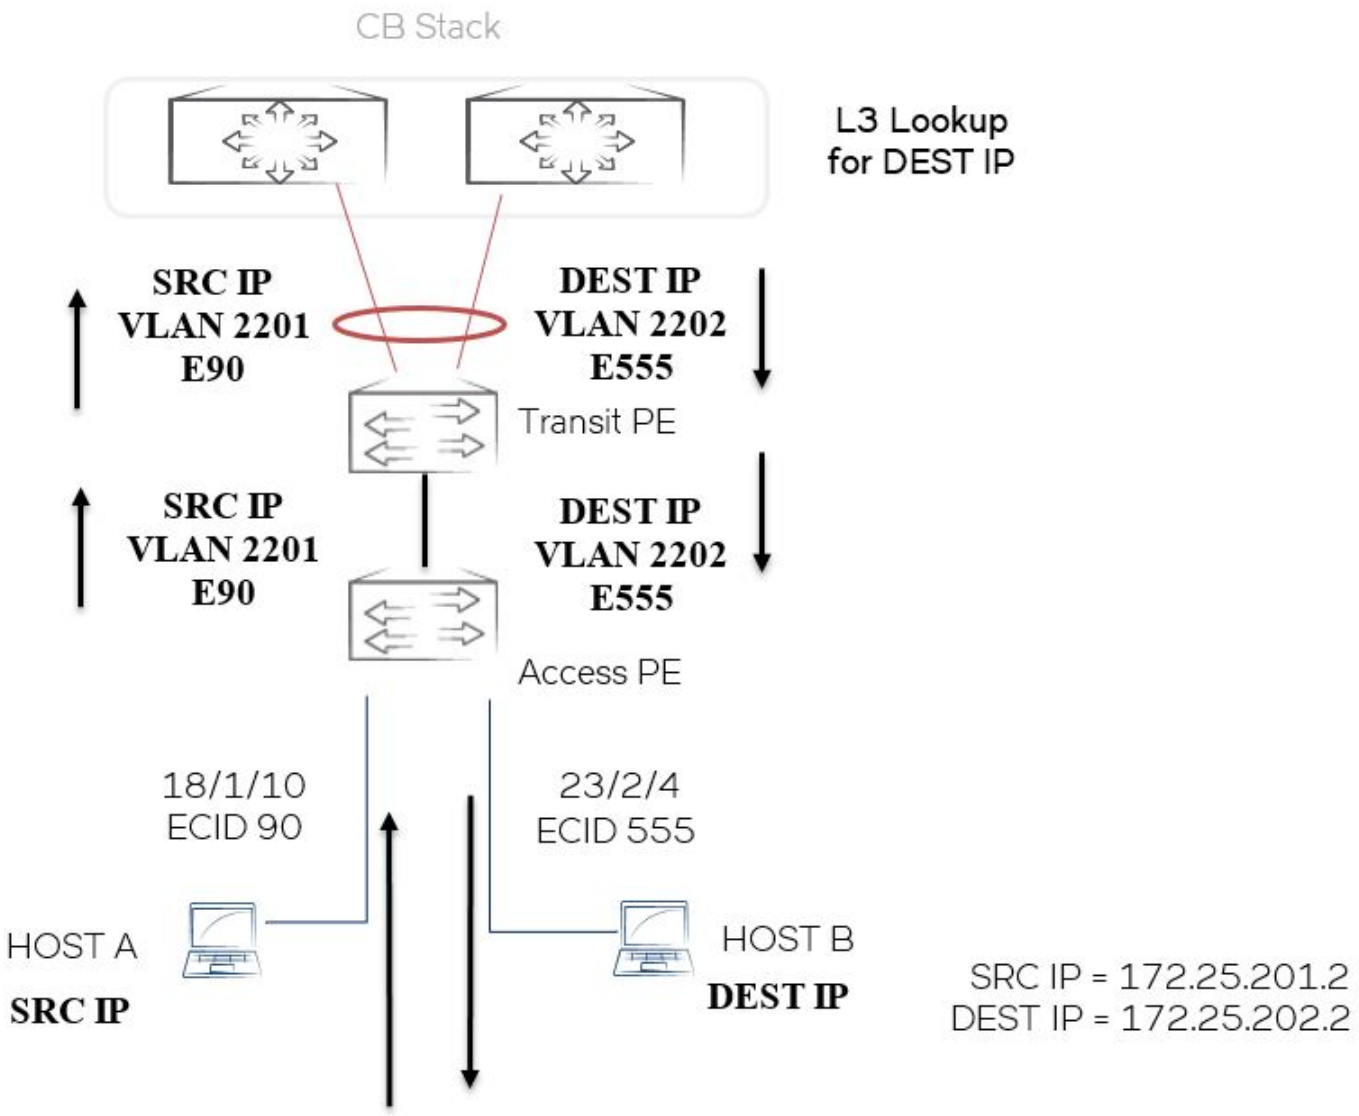

1. The user sends Layer 3 packet with source IP address172.25.201.2 to destination IP address172.25.202.2.

```
CF_SW4#sh run int ve 2201
interface ve 2201
 ip address 172.25.201.1 255.255.255.0
 ip helper-address 1 172.25.10.32
!
CF_SW4#sh run int ve 2202
interface ve 2202
 ip address 172.25.202.1 255.255.255.0
 ip helper-address 1 172.25.10.32
!
```

2. Layer 3 routing lookup is performed for the destination IP address on the CB.

```
CF_SW4#sh ip route 172.25.202.2
Type Codes - B:BGP D:Connected O:OSPF R:RIP S:Static; Cost - Dist/Metric
BGP  Codes - i:iBGP e:eBGP
OSPF Codes - i:Inter Area 1:External Type 1 2:External Type 2
        Destination        Gateway         Port          Cost          Type Uptime
1       172.25.102.0/24    DIRECT          ve 2202       0/0           D    1m6s
```

3. Inter-VLAN routing is done and the packets are switched out the appropriate port to reach the destination subnet.

Example 4—End Users Accessing the Server Farm

FIGURE 14 End Users Accessing Server Farm

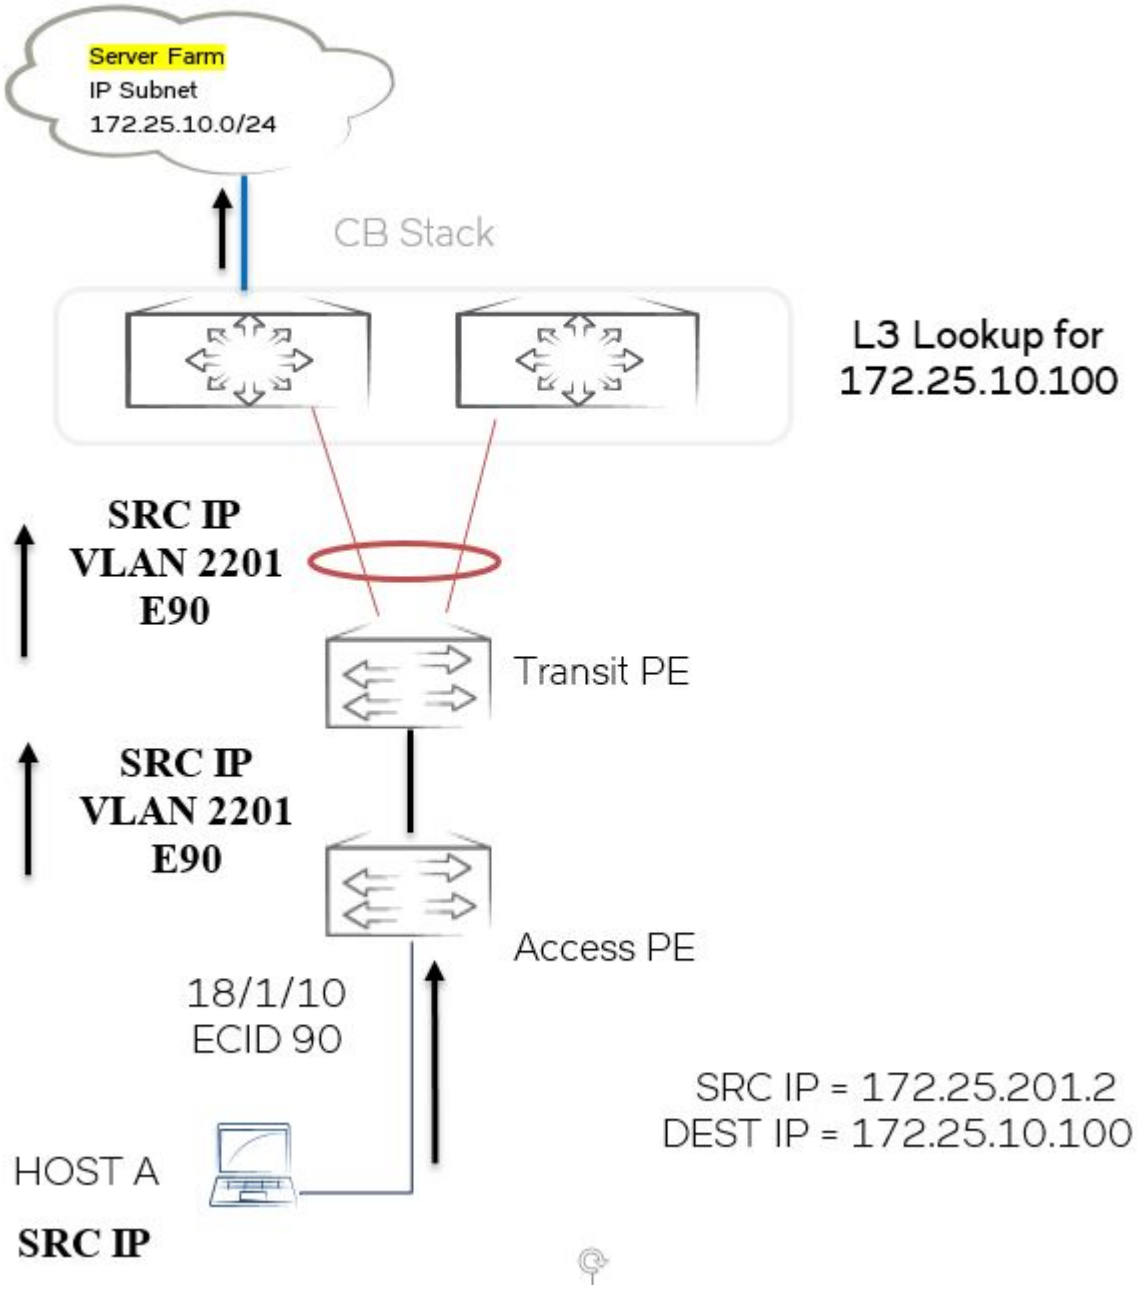

1. The user intends to access the server farm connected to the CB in another PE chain. The user is a part of VLAN 2201 and the server farm is reachable through the core.

```
CF_SW4#sh run vlan 2201
vlan 2201 by port
```

```
 tagged ethe 1/1/48
 router-interface ve 2201
!
!
CF_SW4#sh run int ve 2201
interface ve 2201
 ip address 172.25.201.1 255.255.255.0
 ip helper-address 1 172.25.10.32
 ipv6 pim-sparse
!
CF_SW4#sh ip route 172.25.10.100
Type Codes - B:BGP D:Connected O:OSPF R:RIP S:Static; Cost - Dist/Metric
BGP  Codes - i:iBGP e:eBGP
OSPF Codes - i:Inter Area 1:External Type 1 2:External Type 2
        Destination        Gateway         Port          Cost          Type Uptime
1       172.25.10.0/24     172.25.231.2    ve 2301       110/3         O    1d7h
        172.25.10.0/24     172.25.232.2    ve 2302       110/3         O    1d7h
```

2. Layer 3 communication is achieved using inter-VLAN routing as detailed in Example 3—Layer 3 Communication Between End Users on page 66.

Example 5—End Users Subscribing to Multicast Stream

FIGURE 15 Multicast Communication

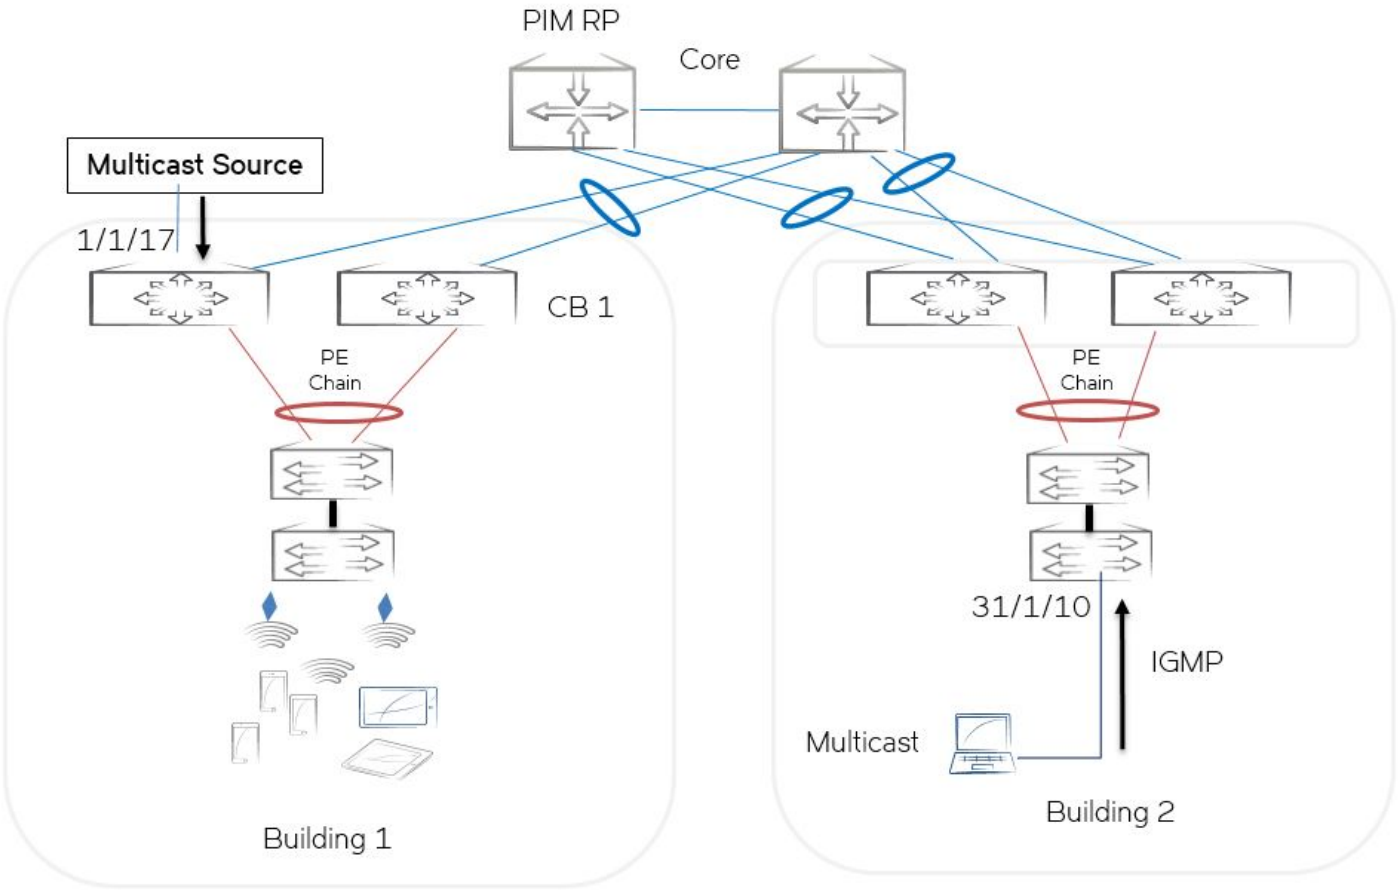

1. The multicast source is connected to port 1/1/17 on the CB, which resides in VLAN 2251 and is part of virtual interface ve 2251, which is PIM-enabled.

2. Core devices are configured as the PIM RP.

```
vlan 2251 by port
 untagged ethe 1/1/17
 router-interface ve 2251
!
!
interface ve 2251
 ip address 172.25.251.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:251::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
router pim
 rp-address 31.1.1.1
!
!
ipv6 router pim
 rp-address fdd7:b215:19cb:4552:31:1:1:1
!
```

3. Multicast receivers reside on the PE port in a different PE chain, which is reachable by way of the core. The core is configured to be the PIM RP.

4. The mulitcast stream is replicated on the CB and sent out the PE ports to which the receivers are connected.

```
CF_SW1#sh ip igmp group
Total 1 groups
--------------------------------------------------------------------
Idx  Group Address    Port      Intf      GrpCmpV Mode    Timer Srcs
----+---------------+---------+---------+-------+-------+-----+-----
   1 225.1.1.1        e24/2/1   v1104     Ver2    exclude 235   0
                      e19/2/1   v1104     Ver2    exclude 250   0

Total number of groups 1

CF_SW1#sh ip pim mcache
IP Multicast Mcache Table
Entry Flags    : SM  - Sparse Mode, SSM - Source Specific Multicast, DM - Dense Mode
                 RPT    - RPT Bit, SPT - SPT Bit, LSRC - Local Source, LRCV - Local Receiver
                 HW - HW Forwarding Enabled, FAST - Resource Allocated, TAG - Need For Replication
Entry
                 REG - Source Registration, L2REG - Source Registration with PIM Neighbors on Source
VLAN
                 REGPROB - Register In Progress, REGSUPP - Register Suppression Timer
                 MSDPADV - Advertise MSDP, NEEDRTE - Route Required for Src/RP,  PRUN - DM Prune
Upstream
Interface Flags: IM - Immediate, IH - Inherited, WA - Won Assert, MCTPEERF - Traffic Forw By Cluster
Peer CCEP
                 MJ - Membership Join, MI - Membership Include, ME - Membership Exclude
                 BR - Blocked RPT, BA - Blocked Assert, BF - Blocked Filter, BI - Blocked IIF
Total entries in mcache: 2

1     (*, 225.1.1.1) RP 31.1.1.1, in v1301 (NIL), Uptime 00:01:21 (SM)
      upstream neighbor 172.25.131.2
      Flags (0x002604a0) SM RPT LRCV TAG
      slow ports: ethe 19/2/1 ethe 24/2/1
      AgeSltMsk: 0, IPMC: NotReq
      Forwarding_oif: 2, Immediate_oif: 2, Blocked_oif: 0
      L3 (SW) 2:
          e19/2/1(VL1104), 00:01:21/0, Flags: MJ
          e24/2/1(VL1104), 00:01:02/0, Flags: MJ

2     (172.25.251.2, 225.1.1.1) in v1302 (e1/1/11), Uptime 00:01:21 (SM)
      upstream neighbor 172.25.132.2
      Flags (0x600680e1) SM SPT LRCV HW FAST TAG
      fast ports: ethe 19/2/1 ethe 24/2/1
      AgeSltMsk: 1, IPMC:   367
      Forwarding_oif: 2, Immediate_oif: 0, Blocked_oif: 0
      L3 (HW) 2:
          e19/2/1(VL1104), 00:01:21/0, Flags: MJ
          e24/2/1(VL1104), 00:01:02/0, Flags: MJ
      Src-Vlan: 1302

Number of matching entries: 2
```

Appendix—Configuration of the Nodes

This appendix includes the relevant configurations of a few nodes in the fabric.

SPX CB1

```
CF_SW1# sh runn
Current configuration:
!
ver 08.0.50B14T203
```

```
!
stack unit 1
  module 1 icx7750-48-xgf-port-management-module
  module 2 icx7750-qsfp-6port-qsfp-240g-module
  priority 128
  stack-port 1/2/1 1/2/4
stack unit 2
  module 1 icx7750-48-xgf-port-management-module
  module 2 icx7750-qsfp-6port-qsfp-240g-module
  stack-port 2/2/1 2/2/4
stack unit 3
  module 1 icx7750-48-xgf-port-management-module
  module 2 icx7750-qsfp-6port-qsfp-240g-module
  stack-port 3/2/1 3/2/4
spx unit 17
  module 1 icx7450-24p-poe-port-management-module
  module 2 icx7400-qsfp-1port-40g-module
  module 3 icx7400-qsfp-1port-40g-module
  module 4 icx7400-qsfp-1port-40g-module
  spx-lag 17/2/1 17/3/1
  spx-port 17/4/1
spx unit 18
  module 1 icx7450-24p-poe-port-management-module
  module 2 icx7400-xgf-4port-40g-module
  module 3 icx7400-qsfp-1port-40g-module
  module 4 icx7400-qsfp-1port-40g-module
  spx-port 18/3/1
  spx-port 18/4/1
spx unit 19
  module 1 icx7450-24p-poe-port-management-module
  module 2 icx7400-xgf-4port-40g-module
  module 3 icx7400-qsfp-1port-40g-module
  module 4 icx7400-qsfp-1port-40g-module
  spx-port 19/3/1
  spx-port 19/4/1
spx unit 20
  module 1 icx7450-24p-poe-port-management-module
  module 2 icx7400-xgf-4port-40g-module
  module 4 icx7400-qsfp-1port-40g-module
  spx-lag 20/2/1 to 20/2/4
  spx-port 20/4/1
spx unit 21
  module 1 icx7250-24p-poe-port-management-module
  module 2 icx7250-sfp-plus-8port-80g-module
  spx-lag 21/2/1 to 21/2/4
  spx-lag 21/2/5 to 21/2/8
spx unit 23
  module 1 icx7450-24p-poe-port-management-module
  module 2 icx7400-qsfp-1port-40g-module
  module 3 icx7400-qsfp-1port-40g-module
  module 4 icx7400-qsfp-1port-40g-module
  spx-lag 23/2/1 23/3/1
  spx-port 23/4/1
spx unit 24
  module 1 icx7450-48p-poe-management-module
  module 2 icx7400-xgf-4port-40g-module
  module 3 icx7400-qsfp-1port-40g-module
  module 4 icx7400-qsfp-1port-40g-module
  spx-port 24/3/1
spx unit 29
  module 1 icx7450-24p-poe-port-management-module
  module 2 icx7400-xgf-4port-40g-module
  module 3 icx7400-qsfp-1port-40g-module
  module 4 icx7400-qsfp-1port-40g-module
  spx-lag 29/2/1 to 29/2/4
  spx-lag 29/3/1 29/4/1
stack enable
stack rconsole-off
stack mac cc4e.24d1.ba00
!
```

```
!
!
spx cb-enable
spx cb-configure
  spx-lag  1/2/2 2/2/2
  spx-lag  1/2/6 2/2/6
  spx-lag  2/2/5 3/2/5
  pe-id 1/2/2 17 18 19 20 21 29 1/2/6
  pe-id 2/2/5 23 24
!
!
!
!
rate-limit-arp 100
lag dynCB1to301 dynamic id 2
 ports ethernet 2/1/12 ethernet 3/1/11
 primary-port 2/1/12
 deploy
!
lag dynCB1to302 dynamic id 1
 ports ethernet 1/1/11 ethernet 2/1/11
 primary-port 1/1/11
 deploy
!
!
vlan 1 name DEFAULT-VLAN by port
!
vlan 11 by port
 untagged ethe 1/1/40 ethe 2/1/40
 router-interface ve 11
!
vlan 111 name VLANFromPE17 by port
 untagged ethe 17/1/11
 router-interface ve 111
!
vlan 121 name VLANFromPE18 by port
 untagged ethe 18/1/11
 router-interface ve 121
!
vlan 131 name VLANFromPE19 by port
 untagged ethe 19/1/11
 router-interface ve 131
!
vlan 141 name VLANFromPE20 by port
 untagged ethe 20/1/11
 router-interface ve 141
!
vlan 151 name VLANFromPE21 by port
 untagged ethe 21/1/11
 router-interface ve 151
!
vlan 161 name VLANFromPE29 by port
 untagged ethe 29/1/11
 router-interface ve 161
!
vlan 171 name VLANFromPE23 by port
 untagged ethe 23/1/11
 router-interface ve 171
!
vlan 181 name VLANFromPE24 by port
 untagged ethe 24/1/11
 router-interface ve 181
!
!
!
vlan 1101 by port
 tagged ethe 1/1/48
 router-interface ve 1101
!
vlan 1102 by port
```

```
 tagged ethe 1/1/48
 router-interface ve 1102
!
vlan 1103 by port
 tagged ethe 1/1/48
 router-interface ve 1103
!
vlan 1104 by port
 tagged ethe 1/1/48
 router-interface ve 1104
!
vlan 1105 by port
 tagged ethe 1/1/48
 router-interface ve 1105
!
vlan 1106 by port
 tagged ethe 1/1/48
 router-interface ve 1106
!
vlan 1107 by port
 tagged ethe 1/1/48
 router-interface ve 1107
!
vlan 1108 by port
 tagged ethe 1/1/48
 router-interface ve 1108
!
vlan 1111 by port
 tagged ethe 1/1/48
 router-interface ve 1111
!
vlan 1112 by port
 tagged ethe 1/1/48
 router-interface ve 1112
!
vlan 1113 by port
 tagged ethe 1/1/48
 router-interface ve 1113
!
vlan 1114 by port
 tagged ethe 1/1/48
 router-interface ve 1114
!
vlan 1115 by port
 tagged ethe 1/1/48
 router-interface ve 1115
!
vlan 1116 by port
 tagged ethe 1/1/48
 router-interface ve 1116
!
vlan 1117 by port
 tagged ethe 1/1/48
 router-interface ve 1117
!
vlan 1118 by port
 tagged ethe 1/1/48
 router-interface ve 1118
!
vlan 1141 by port
 tagged ethe 1/1/48
 router-interface ve 1141
!
vlan 1142 by port
 tagged ethe 1/1/48
 router-interface ve 1142
!
vlan 1143 by port
tagged ethe 1/1/48
 router-interface ve 1143
```

```
!
vlan 1144 by port
 tagged ethe 1/1/48
 router-interface ve 1144
!
vlan 1151 by port
 untagged ethe 2/1/17
 router-interface ve 1151
!
vlan 1301 by port
 untagged ethe 2/1/12 ethe 3/1/11
 router-interface ve 1301
!
vlan 1302 by port
 untagged ethe 1/1/11 ethe 2/1/11
 router-interface ve 1302
!
!
!
!
!
!
!
!
!
!
authentication
 critical-vlan 1142
 auth-default-vlan 1141
 restricted-vlan 1143
 auth-fail-action restricted-vlan
 re-authentication
 auth-vlan-mode multiple-untagged
 dot1x enable
 dot1x enable ethe 19/2/1 ethe 24/2/1
 dot1x guest-vlan 1144
 mac-authentication enable
 mac-authentication enable ethe 19/2/1 ethe 24/2/1
 mac-authentication dot1x-override
!
aaa authentication dot1x default radius
aaa authentication login default radius
aaa accounting dot1x default start-stop radius
hostname CF_SW1
ip route 0.0.0.0/0 10.37.234.1
!
ipv6 unicast-routing
radius-server host 10.32.12.33 auth-port 1812 acct-port 1813 default key 2 $b24tbz1nIlp8OA== dot1x
radius-server host 172.22.108.156 auth-port 1812 acct-port 1813 default key 2 $UituWnw4 dot1x mac-auth web-
auth
snmp-server community ..... ro
snmp-server community ..... rw
!
!
clock summer-time
clock timezone us Pacific
!
!
ntp
 server 10.31.2.80
!
!
hitless-failover enable
ip multicast-routing
!
router ospf
 area 0
 no graceful-restart
 nonstop-routing
 redistribute connected
```

```
!
!
!
router pim
 rp-address 31.1.1.1
!
!
ipv6 router pim
 rp-address fdd7:b215:19cb:4552:31:1:1:1
!
ipv6 router ospf
 area 0
 nonstop-routing
 redistribute connected
!
interface loopback 1
 ip address 11.1.1.1 255.255.255.255
 ip pim-sparse
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:4552:11:1:1:1/128
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
!
interface management 1
 ip address 10.37.234.29 255.255.255.128
!
interface ethernet 1/1/11
 speed-duplex 1000-full
!
interface ethernet 1/1/40
 speed-duplex 1000-full
!
interface ethernet 2/1/12
 speed-duplex 1000-full
!
interface ethernet 19/2/1
 authentication max-sessions 1024
 authentication timeout-action critical-vlan
 dot1x port-control auto
 trust dscp
 sflow forwarding
 sflow sample 4096
!
interface ethernet 24/2/1
 authentication max-sessions 1024
 authentication timeout-action critical-vlan
 dot1x port-control auto
 trust dscp
 sflow forwarding
 sflow sample 4096
!
interface ve 11
 ip address 172.25.10.100 255.255.255.0
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:4552:172:25:10:100/112
!
interface ve 111
!
interface ve 121
!
interface ve 131
!
interface ve 141
!
interface ve 151
!
interface ve 161
```

```
!
interface ve 171
!
interface ve 181
!
interface ve 1101
 ip address 172.25.101.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:101::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1102
 ip address 172.25.102.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:102::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1103
 ip address 172.25.103.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:103::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1104
 ip address 172.25.104.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:104::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1105
 ip address 172.25.105.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:105::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1106
 ip address 172.25.106.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
```

```
 ip ospf passive
 ipv6 address fdd7:b215:19cb:106::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1107
ip address 172.25.107.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:107::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1108
 ip address 172.25.108.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:108::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1111
 ip address 172.25.111.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:111::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1112
 ip address 172.25.112.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:112::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1113
 ip address 172.25.113.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:113::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1114
 ip address 172.25.114.1 255.255.255.0
 ip pim-sparse
```

```
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:114::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1115
 ip address 172.25.115.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:115::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1116
 ip address 172.25.116.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:116::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1117
 ip address 172.25.117.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:117::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1118
 ip address 172.25.118.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:118::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1141
 ip address 172.25.141.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:141::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1142
```

```
 ip address 172.25.142.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:142::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1143
 ip address 172.25.143.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
ipv6 address fdd7:b215:19cb:143::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1144
 ip address 172.25.144.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:144::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1151
 ip address 172.25.151.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:1511::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 1301
 ip address 172.25.131.1 255.255.255.0
 ip pim-sparse
 ip ospf area 0
 ip ospf md5-authentication key-id 1 key 2 $MlVzZCFAbg==
 ip ospf network point-to-point
 ipv6 address fdd7:b215:19cb:4552:172:25:131:1/112
 ipv6 ospf area 0
 ipv6 ospf authentication ipsec spi 301 esp sha1
encryptb64 $Wnw4M09tWVd7USEyZEBuR1p8ODNPbVlXe1EhMmRAbkdafDgzT21ZVw==
 ipv6 ospf network point-to-point
 ipv6 pim-sparse
!
interface ve 1302
 ip address 172.25.132.1 255.255.255.0
 ip pim-sparse
 ip ospf area 0
 ip ospf md5-authentication key-id 1 key 2 $MlVzZCFAbg==
 ip ospf network point-to-point
 ipv6 address fdd7:b215:19cb:4552:172:25:132:1/112
 ipv6 ospf area 0
 ipv6 ospf authentication ipsec spi 302 esp sha1
encryptb64 $Wnw4M09tWVd7USEyZEBuR1p8ODNPbVlXe1EhMmRAbkdafDgzT21ZVw==
```

```
 ipv6 ospf network point-to-point
 ipv6 pim-sparse
!
!
!
ip access-list extended Tplay_Trfc
 permit udp any any range 30000 30100 dscp-matching 46 dscp-marking 46 802.1p-and-internal-marking 5
 permit tcp any any eq 2000 dscp-marking 46 802.1p-and-internal-marking 5
 permit tcp any any range 1719 1720 dscp-marking 46 802.1p-and-internal-marking 5
 permit udp any any range 1719 1720 dscp-marking 46 802.1p-and-internal-marking 5
 permit tcp any any eq 1731 dscp-marking 46 802.1p-and-internal-marking 5
 permit udp any any eq 1731 dscp-marking 46 802.1p-and-internal-marking 5
 permit tcp any any eq 5060 dscp-marking 46 802.1p-and-internal-marking 5
 permit udp any any eq 5060 dscp-marking 46 802.1p-and-internal-marking 5
 permit udp any any dscp-matching 24 dscp-marking 46 802.1p-and-internal-marking 5
 permit tcp any any dscp-matching 24 dscp-marking 46 802.1p-and-internal-marking 5
 permit udp any any eq 5445 dscp-marking 34 802.1p-and-internal-marking 4
 permit udp any any dscp-matching 34 dscp-marking 34 802.1p-and-internal-marking 4
 permit tcp any any dscp-matching 34 dscp-marking 34 802.1p-and-internal-marking 4
 permit tcp any any range 3230 3247 dscp-marking 34 802.1p-and-internal-marking 4
 permit udp any any range 3230 3247 dscp-marking 34 802.1p-and-internal-marking 4
 permit ip any any 802.1p-priority-matching 0 dscp-marking 0 802.1p-and-internal-marking 0
permit ip any any 802.1p-priority-matching 1 dscp-marking 8 802.1p-and-internal-marking 1
 permit ip any any 802.1p-priority-matching 2 dscp-marking 16 802.1p-and-internal-marking 2
 permit ip any any 802.1p-priority-matching 3 dscp-marking 24 802.1p-and-internal-marking 3
 permit ip any any 802.1p-priority-matching 4 dscp-marking 34 802.1p-and-internal-marking 4
 permit ip any any 802.1p-priority-matching 5 dscp-marking 46 802.1p-and-internal-marking 5
 permit ip any any 802.1p-priority-matching 6 dscp-marking 48 802.1p-and-internal-marking 6
 permit ip any any 802.1p-priority-matching 7 dscp-marking 56 802.1p-and-internal-marking 7
 permit ip any any
!
!
sflow destination 172.25.151.2
sflow enable
!
lldp run
!
!
!
!
end
```

SPX CB2

```
CF_SW4# sh runn
Current configuration:
!
ver 08.0.50T203
!
stack unit 1
  module 1 icx7750-48-xgf-port-management-module
  module 2 icx7750-qsfp-6port-qsfp-240g-module
  priority 128
  stack-port 1/2/1 1/2/4
stack unit 2
  module 1 icx7750-48-xgf-port-management-module
  module 2 icx7750-qsfp-6port-qsfp-240g-module
  priority 10
  stack-port 2/2/1 2/2/4
spx unit 31
  module 1 icx7450-24p-poe-port-management-module
  module 2 icx7400-xgf-4port-40g-module
  module 3 icx7400-xgf-4port-40g-module
  module 4 icx7400-qsfp-1port-40g-module
  spx-lag 31/2/1 to 31/2/4
```

```
   spx-lag 31/3/1 to 31/3/4
spx unit 32
   module 1 icx7450-24p-poe-port-management-module
   module 2 icx7400-xgf-4port-40g-module
   module 3 icx7400-qsfp-1port-40g-module
   module 4 icx7400-xgf-4port-40g-module
   spx-lag 32/4/1 to 32/4/4
stack enable
stack rconsole-off
stack mac cc4e.24d0.5e80
!
!
!
spx cb-enable
spx cb-configure
   spx-lag  1/1/3 to 1/1/4 2/1/3 to 2/1/4
   pe-id 1/1/3 31 32
!
!
!
!
rate-limit-arp 100
lag dynCB2to301 dynamic id 2
 ports ethernet 1/1/12 ethernet 2/1/12
 primary-port 1/1/12
 deploy
!
lag dynCB2to302 dynamic id 1
 ports ethernet 1/1/11 ethernet 2/1/11
 primary-port 1/1/11
 deploy
!
!
vlan 1 name DEFAULT-VLAN by port
!
vlan 221 name VLANFromPE23_2 by port
 tagged ethe 31/1/9
 router-interface ve 221
!
vlan 256 by port
!
vlan 257 by port
!
vlan 258 by port
!
vlan 259 by port
!
vlan 260 by port
!
!
!
!
vlan 1888 by port
 tagged ethe 31/1/10
!
!
vlan 2201 by port
 tagged ethe 1/1/48
 router-interface ve 2201
!
vlan 2202 by port
 tagged ethe 1/1/48
 router-interface ve 2202
!
vlan 2203 by port
 tagged ethe 1/1/48
 router-interface ve 2203
!
vlan 2204 by port
 tagged ethe 1/1/48
```

```
 router-interface ve 2204
!
vlan 2205 by port
 tagged ethe 1/1/48
 router-interface ve 2205
!
vlan 2206 by port
 tagged ethe 1/1/48
 router-interface ve 2206
!
vlan 2207 by port
 tagged ethe 1/1/48
 router-interface ve 2207
!
vlan 2208 by port
 tagged ethe 1/1/48
 router-interface ve 2208
!
vlan 2211 by port
 tagged ethe 1/1/48
 router-interface ve 2211
!
vlan 2212 by port
 tagged ethe 1/1/48
 router-interface ve 2212
!
vlan 2213 by port
 tagged ethe 1/1/48
 router-interface ve 2213
!
vlan 2214 by port
 tagged ethe 1/1/48
 router-interface ve 2214
!
vlan 2215 by port
 tagged ethe 1/1/48
 router-interface ve 2215
!
vlan 2216 by port
 tagged ethe 1/1/48
 router-interface ve 2216
!
vlan 2217 by port
 tagged ethe 1/1/48
 router-interface ve 2217
!
vlan 2218 by port
 tagged ethe 1/1/48
 router-interface ve 2218
!
vlan 2241 by port
 tagged ethe 1/1/48
 router-interface ve 2241
!
vlan 2242 by port
 tagged ethe 1/1/48
 router-interface ve 2242
!
vlan 2243 by port
 tagged ethe 1/1/48
 router-interface ve 2243
!
vlan 2244 by port
 tagged ethe 1/1/48
 router-interface ve 2244
!
vlan 2251 by port
 untagged ethe 1/1/17
 router-interface ve 2251
!
```

```
vlan 2301 by port
 untagged ethe 1/1/12 ethe 2/1/12
 router-interface ve 2301
!
vlan 2302 by port
 untagged ethe 1/1/11 ethe 2/1/11
 router-interface ve 2302
!
!
!
!
!
!
!
!
authentication
 critical-vlan 2242
 auth-default-vlan 2241
 restricted-vlan 2243
 auth-fail-action restricted-vlan
 auth-vlan-mode multiple-untagged
 dot1x enable
 dot1x enable ethe 31/1/10 ethe 32/2/4
 dot1x guest-vlan 2244
 mac-authentication enable
 mac-authentication password-format xxxx.xxxx.xxxx
!
aaa authentication dot1x default radius
aaa authentication login default radius
aaa accounting dot1x default start-stop radius
hostname CF_SW4
ip route 0.0.0.0/0 10.37.234.1
ip multicast passive
!
ipv6 unicast-routing
logging console
radius-server host 10.32.12.33 auth-port 1812 acct-port 1813 default key 2 $b24tbz1nIlp8OA== dot1x mac-auth
web-auth
snmp-server community ..... ro
snmp-server community ..... rw
snmp-server enable mib np-qos-stat
snmp-server trap-source loopback 1
snmp-server host 10.32.12.31 version v2c .....
snmp-server host 172.25.10.31 version v2c .....
!
!
clock summer-time
clock timezone us Pacific
!
!
ntp
 server 10.31.2.80
!
!
hitless-failover enable
ip multicast-routing
!
router ospf
 area 0
 no graceful-restart
 nonstop-routing
 redistribute connected
!
!
!
!
!
!
router pim
 rp-address 31.1.1.1
```

```
!
!
ipv6 router pim
 rp-address fdd7:b215:19cb:4552:31:1:1:1
!
ipv6 router ospf
 area 0
 nonstop-routing
 redistribute connected
!
interface loopback 1
 ip address 22.1.1.1 255.255.255.255
 ip pim-sparse
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:4552:22:1:1:1/128
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
!
interface management 1
 ip address 10.37.234.32 255.255.255.128
!
interface ethernet 1/1/1
 disable
 ip access-group "block Telnet" in
!
interface ethernet 1/1/2
 disable
!
interface ethernet 1/1/11
 speed-duplex 1000-full
!
interface ethernet 1/1/12
 speed-duplex 1000-full
!
interface ethernet 1/1/19
 port-name Connected to CER1-2/3
!
interface ethernet 1/1/20
 speed-duplex 1000-full
!
interface ethernet 2/1/1
 disable
!
interface ethernet 2/1/2
disable
!
interface ethernet 31/1/1
 inline power power-limit 15400
!
interface ethernet 31/1/2
 inline power power-by-class 3
!
interface ethernet 31/1/3
 inline power power-by-class 3
!
interface ethernet 31/1/4
 inline power power-by-class 3
!
interface ethernet 31/1/5
 inline power power-by-class 3
!
interface ethernet 31/1/6
 inline power power-by-class 3
!
interface ethernet 31/1/7
 inline power power-by-class 3
!
interface ethernet 31/1/8
```

```
 inline power power-by-class 3
!
interface ethernet 31/1/10
 authentication max-sessions 1024
 authentication timeout-action critical-vlan
 dot1x port-control auto
 max-vlan 12
!
interface ethernet 31/1/13
 max-vlan 16
!
interface ethernet 31/1/17
 max-vlan 16
!
interface ethernet 32/1/1
 inline power power-limit 15400
!
interface ethernet 32/1/13
 max-vlan 16
!
interface ethernet 32/2/4
 authentication max-sessions 1024
 authentication timeout-action critical-vlan
 dot1x port-control auto
 trust dscp
 sflow forwarding
!
interface ve 221
!
interface ve 2201
 ip address 172.25.201.1 255.255.255.0
 ip helper-address 1 172.25.10.32
 ipv6 pim-sparse
!
interface ve 2202
 ip address 172.25.202.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:202::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2203
ip address 172.25.203.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:203::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2204
 ip address 172.25.204.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:204::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
```

```
interface ve 2205
 ip address 172.25.205.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:205::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2206
 ip address 172.25.206.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:206::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2207
 ip address 172.25.207.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:207::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2208
 ip address 172.25.208.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:208::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2211
 ip address 172.25.211.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:211::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2212
 ip address 172.25.212.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:212::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
```

```
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2213
 ip address 172.25.213.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:213::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2214
 ip address 172.25.214.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:214::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2215
 ip address 172.25.215.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:215::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2216
 ip address 172.25.216.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:216::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2217
 ip address 172.25.217.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:217::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2218
 ip address 172.25.218.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:218::1/64
 ipv6 ospf area 0
```

```
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2241
 ip address 172.25.241.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:241::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2242
 ip address 172.25.242.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:242::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2243
 ip address 172.25.243.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:243::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2244
 ip address 172.25.244.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
ipv6 address fdd7:b215:19cb:244::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2251
 ip address 172.25.251.1 255.255.255.0
 ip pim-sparse
 ip helper-address 1 172.25.10.32
 ip ospf area 0
 ip ospf passive
 ipv6 address fdd7:b215:19cb:251::1/64
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
 ipv6 dhcp-relay destination fdd7:b215:19cb:4552:172:25:10:32
!
interface ve 2301
 ip address 172.25.231.1 255.255.255.0
 ip pim-sparse
 ip ospf area 0
 ip ospf md5-authentication key-id 1 key 2 $MlVzZCFAbg==
 ip ospf network point-to-point
```

```
 ipv6 address fdd7:b215:19cb:4552:172:25:231:1/112
 ipv6 ospf area 0
 ipv6 ospf authentication ipsec spi 301 esp sha1
encryptb64 $Wnw4M09tWVd7USEyZEBuR1p8ODNPbVlXe1EhMmRAbkdafDgzT21ZVw==
 ipv6 ospf network point-to-point
 ipv6 pim-sparse
!
interface ve 2302
 ip address 172.25.232.1 255.255.255.0
 ip pim-sparse
 ip ospf area 0
 ip ospf md5-authentication key-id 1 key 2 $MlVzZCFAbg==
 ip ospf network point-to-point
 ipv6 address fdd7:b215:19cb:4552:172:25:232:1/112
 ipv6 ospf area 0
 ipv6 ospf authentication ipsec spi 302 esp sha1
encryptb64 $Wnw4M09tWVd7USEyZEBuR1p8ODNPbVlXe1EhMmRAbkdafDgzT21ZVw==
 ipv6 ospf network point-to-point
 ipv6 pim-sparse
!
!
!
ip access-list standard 1
 deny host 10.157.22.26 log
 deny host 10.157.29.12 log
 permit any
!
ip access-list extended 101
 deny tcp host 10.157.22.26 any eq telnet log
 permit ip any any
!
ip access-list standard Net1
 deny host 10.157.22.26 log
 deny host 10.157.29.12 log
 permit any
!
ip access-list extended dynamic_acl_1
 sequence 10 deny tcp host 10.1.1.1 any
 permit ip any any
!
ip access-list extended Tplay_Trfc
 permit udp any any range 30000 30100 dscp-matching 46 dscp-marking 46 802.1p-and-internal-marking 5
 permit tcp any any eq 2000 dscp-marking 46 802.1p-and-internal-marking 5
 permit tcp any any range 1719 1720 dscp-marking 46 802.1p-and-internal-marking 5
 permit udp any any range 1719 1720 dscp-marking 46 802.1p-and-internal-marking 5
 permit tcp any any eq 1731 dscp-marking 46 802.1p-and-internal-marking 5
permit udp any any eq 1731 dscp-marking 46 802.1p-and-internal-marking 5
 permit tcp any any eq 5060 dscp-marking 46 802.1p-and-internal-marking 5
 permit udp any any eq 5060 dscp-marking 46 802.1p-and-internal-marking 5
 permit udp any any dscp-matching 24 dscp-marking 46 802.1p-and-internal-marking 5
 permit tcp any any dscp-matching 24 dscp-marking 46 802.1p-and-internal-marking 5
 permit udp any any eq 5445 dscp-marking 34 802.1p-and-internal-marking 4
 permit udp any any dscp-matching 34 dscp-marking 34 802.1p-and-internal-marking 4
 permit tcp any any dscp-matching 34 dscp-marking 34 802.1p-and-internal-marking 4
 permit tcp any any range 3230 3247 dscp-marking 34 802.1p-and-internal-marking 4
 permit udp any any range 3230 3247 dscp-marking 34 802.1p-and-internal-marking 4
 permit ip any any 802.1p-priority-matching 0 dscp-marking 0 802.1p-and-internal-marking 0
 permit ip any any 802.1p-priority-matching 1 dscp-marking 8 802.1p-and-internal-marking 1
 permit ip any any 802.1p-priority-matching 2 dscp-marking 16 802.1p-and-internal-marking 2
 permit ip any any 802.1p-priority-matching 3 dscp-marking 24 802.1p-and-internal-marking 3
 permit ip any any 802.1p-priority-matching 4 dscp-marking 34 802.1p-and-internal-marking 4
 permit ip any any 802.1p-priority-matching 5 dscp-marking 46 802.1p-and-internal-marking 5
 permit ip any any 802.1p-priority-matching 6 dscp-marking 48 802.1p-and-internal-marking 6
 permit ip any any 802.1p-priority-matching 7 dscp-marking 56 802.1p-and-internal-marking 7
 permit ip any any
!
ip access-list extended "block Telnet"
 deny tcp host 10.157.22.26 any eq telnet log
 permit ip any any
!
```

```
!
sflow destination 172.25.151.2
sflow enable
!
lldp run
!
!
!
!
end
```

SPX Core Switch 1

```
CORE-301#sh runn
!Current configuration:
!
ver V5.9.0pT183
!
!
!
lag "TransitTo301" dynamic id 6
 ports ethernet 1/15 to 1/16
 primary-port 1/15
 deploy
!
lag "dyn301toCB1" dynamic id 4
 ports ethernet 1/5 to 1/6
 primary-port 1/5
 deploy
!
lag "dyn301toCB2" dynamic id 5
 ports ethernet 1/7 to 1/8
 primary-port 1/7
 deploy
!
!
!

!
no spanning-tree
!
!
vlan 1301
 untagged ethe 1/5 to 1/6
 router-interface ve 131
!
vlan 2301
 untagged ethe 1/7 to 1/8
 router-interface ve 231
!

!
system-max ip-subnet-port 128
system-max ip-vrf 5
!
!
default-max-frame-size 5000
!
!
clock timezone us Pacific
!
!
ntp
 server 10.37.6.213
!
```

```
!
!
!
!
!
!
hostname CORE-301
ip multicast-routing
!
router ospf
 area 0
 redistribute connected
!
!
!
router pim
rp-address 31.1.1.1
!
ipv6 router ospf
 area 0
 redistribute connected
!
!
ipv6 router pim
 rp-address fdd7:b215:19cb:4552:31:1:1:1
!
!
!
interface loopback 1
 ip ospf area 0
 ip ospf passive
 ip address 31.1.1.1/32
 ip pim-sparse
 ipv6 address fdd7:b215:19cb:4552:31:1:1:1/128
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
!
!
interface management 1
 ip address 10.37.238.12/24
 enable
!
interface ethernet 1/1
 enable
 speed-duplex 1000-full
!
interface ethernet 1/5
 enable
!
interface ethernet 1/7
 enable
!
interface ethernet 1/11
 enable
!
interface ethernet 1/12
 enable
!
interface ethernet 1/15
 enable
!
interface ethernet 1/21
 enable
!
interface ethernet 1/23
 enable
!
interface ve 101
 ip address 10.10.10.1/24
```

```
!
interface ve 131
 ip ospf area 0
 ip ospf md5-authentication key-id 1 key 2 $MlVzZCFAbg==
 ip ospf network point-to-point
 ip address 172.25.131.2/24
 ip pim-sparse
 ipv6 address fdd7:b215:19cb:4552:172:25:131:2/112
 ipv6 ospf area 0
 ipv6 ospf authentication ipsec spi 301 esp sha1
encryptb64 $Wnw4M09tWVd7USEyZEBuR1p8ODNPbVlXe1EhMmRAbkdafDgzT21ZVw==
 ipv6 ospf network point-to-point
 ipv6 pim-sparse
!
interface ve 231
ip ospf area 0
 ip ospf md5-authentication key-id 1 key 2 $MlVzZCFAbg==
 ip ospf network point-to-point
 ip address 172.25.231.2/24
 ip pim-sparse
 ipv6 address fdd7:b215:19cb:4552:172:25:231:2/112
 ipv6 ospf area 0
 ipv6 ospf authentication ipsec spi 301 esp sha1
encryptb64 $Wnw4M09tWVd7USEyZEBuR1p8ODNPbVlXe1EhMmRAbkdafDgzT21ZVw==
 ipv6 ospf network point-to-point
 ipv6 pim-sparse
!
!
!
!
lldp enable ports ethe 1/1 to 1/5 ethe 1/7 ethe 1/9 to 1/12 ethe 1/15 ethe 1/17 to 1/23
lldp run
!
!
!
!
end
```

SPX Core Switch 2

```
CORE-302#sh runn
!Current configuration:
!
ver V5.7.0aT183
!
!
!
lag "TransitTo302" dynamic id 6
 ports ethernet 1/15 to 1/16
 primary-port 1/15
 deploy
!
lag "dyn302toCB1" dynamic id 4
 ports ethernet 1/5 to 1/6
 primary-port 1/5
 deploy
!
lag "dyn302toCB2" dynamic id 5
 ports ethernet 1/7 to 1/8
 primary-port 1/7
 deploy
!
!
!
!
```

```
!
no spanning-tree
!
!
!
vlan 1302
 untagged ethe 1/5 to 1/6
 router-interface ve 132
!
vlan 2302
 untagged ethe 1/7 to 1/8
 router-interface ve 232
!

!
system-max ip-subnet-port 128
system-max ip-vrf 5
!
!
default-max-frame-size 5000
!
!
clock timezone us Pacific
!
!
ntp
 server 10.37.6.213
!
!
!
!
!
!
!
hostname CORE-302
ip multicast-routing
!
router ospf
 area 0
 redistribute connected
!
!
!
router pim
 rp-address 31.1.1.1
!
ipv6 router ospf
 area 0
 redistribute connected
!
!
ipv6 router pim
 rp-address fdd7:b215:19cb:4552:31:1:1:1
!
!
!
interface loopback 1
 ip ospf area 0
 ip ospf passive
 ip address 32.1.1.1/32
 ip pim-sparse
 ipv6 address fdd7:b215:19cb:4552:32:1:1:1/128
 ipv6 ospf area 0
 ipv6 ospf passive
 ipv6 pim-sparse
!
!
interface management 1
 ip address 10.37.238.13/24
```

```
 enable
!
interface ethernet 1/1
 enable
 speed-duplex 1000-full
!
interface ethernet 1/2
 enable
 speed-duplex 1000-full
!
interface ethernet 1/5
 enable
!
interface ethernet 1/7
 enable
!
interface ethernet 1/11
 enable
!
interface ethernet 1/12
 enable
!
interface ethernet 1/15
 enable
!
interface ethernet 1/21
 enable
!
interface ethernet 1/23
 enable
!
interface ve 101
 ip address 10.10.10.2/24
!
interface ve 132
 ip ospf area 0
 ip ospf md5-authentication key-id 1 key 2 $MlVzZCFAbg==
 ip ospf network point-to-point
 ip address 172.25.132.2/24
 ip pim-sparse
 ipv6 address fdd7:b215:19cb:4552:172:25:132:2/112
 ipv6 ospf area 0
 ipv6 ospf authentication ipsec spi 302 esp sha1
encryptb64 $Wnw4M09tWVd7USEyZEBuR1p8ODNPbVlXe1EhMmRAbkdafDgzT21ZVw==
 ipv6 ospf network point-to-point
 ipv6 pim-sparse
!
interface ve 232
 ip ospf area 0
 ip ospf md5-authentication key-id 1 key 2 $MlVzZCFAbg==
 ip ospf network point-to-point
 ip address 172.25.232.2/24
 ip pim-sparse
 ipv6 address fdd7:b215:19cb:4552:172:25:232:2/112
 ipv6 ospf area 0
 ipv6 ospf authentication ipsec spi 302 esp sha1
encryptb64 $Wnw4M09tWVd7USEyZEBuR1p8ODNPbVlXe1EhMmRAbkdafDgzT21ZVw==
 ipv6 ospf network point-to-point
 ipv6 pim-sparse
!
!
!
!
lldp enable ports ethe 1/1 to 1/5 ethe 1/7 ethe 1/9 to 1/12 ethe 1/15 ethe 1/17 to 1/23
lldp run
!
!
!
```

```
!
end
```

References

1. *802.1BR - Bridge Port Extension*

   http://www.ieee802.org/1/pages/802.1br.html

2. *Brocade FastIron Campus Fabric Configuration Guide, 08.0.50*

   http://www.brocade.com/en/backend-content/pdf-page.html?/content/dam/secure-external/product-guides/brocade-fastiron-os/08-0-50/fastiron-08050-campusfabricguide.pdf