


**Part No. 060378-10, Rev. A
December 2012**

OmniSwitch AOS Release 7 Data Center Switching Guide

Alcatel·Lucent 

www.alcatel-lucent.com

**This user guide documents AOS Release 7 for the OmniSwitch 10K and OmniSwitch 6900.
The functionality described in this guide is subject to change without notice.**

Copyright © 2012 by Alcatel-Lucent. All rights reserved. This document may not be reproduced in whole or in part without the express written permission of Alcatel-Lucent.

Alcatel-Lucent® and the Alcatel-Lucent logo are registered trademarks of Alcatel-Lucent. Xylan®, OmniSwitch®, OmniStack®, and Alcatel-Lucent OmniVista® are registered trademarks of Alcatel-Lucent.

OmniAccess™, Omni Switch/Router™, PolicyView™, RouterView™, SwitchManager™, VoiceView™, WebView™, X-Cell™, X-Vision™, and the Xylan logo are trademarks of Alcatel-Lucent.

This OmniSwitch product contains components which may be covered by one or more of the following U.S. Patents:

- U.S. Patent No. 6,339,830
- U.S. Patent No. 6,070,243
- U.S. Patent No. 6,061,368
- U.S. Patent No. 5,394,402
- U.S. Patent No. 6,047,024
- U.S. Patent No. 6,314,106
- U.S. Patent No. 6,542,507
- U.S. Patent No. 6,874,090



**26801 West Agoura Road
Calabasas, CA 91301
(818) 880-3500 FAX (818) 880-3505
support@ind.alcatel.com**

**US Customer Support—(800) 995-2696
International Customer Support—(818) 878-4507
Internet—service.esd.alcatel-lucent.com**

Contents

	About This Guide	xxi
	Supported Platforms	xxi
	Who Should Read this Manual?	xxi
	When Should I Read this Manual?	xxi
	What is in this Manual?	xxii
	What is Not in this Manual?	xxii
	How is the Information Organized?	xxii
	Documentation Roadmap	xxiii
	Related Documentation	xxv
	Technical Support	xxvi
Chapter 1	Understanding Data Center Switching	2-1
	In This Chapter	2-1
	Data Center Switching Components	2-2
	OmniSwitch Pod	2-2
	Data Center Mesh	2-2
	Virtual Network Profile	2-3
	Network Virtualization Technology	2-3
	Data Center Mesh Configuration Example	2-4
Chapter 2	Configuring Data Center Bridging	3-1
	In This Chapter	3-1
	DCB Specifications	3-2
	DCB Defaults	3-3
	DCB Port Defaults	3-4
	Data Center Bridging Overview	3-5
	Priority-Based Flow Control Overview	3-5
	Enhanced Transmission Selection Overview	3-7
	Data Center Bridging Exchange Overview	3-10
	Interaction with Other Features	3-13
	QoS	3-13
	Shortest Path Bridging	3-13
	SNMP	3-13
	Virtual Chassis	3-14
	Using DCB Profiles	3-15

Configuring DCB Profiles	3-21
Creating a Custom Profile	3-21
Changing the Profile Assignment	3-22
Configuring Support for Legacy Pause Frames	3-22
Configuring DCBx Port Parameters	3-23
Multicast and Unicast Traffic Distribution	3-24
Non-Default Profile	3-24
Multicast Source PFC on OmniSwitch 6900	3-26
Verifying the DCB Configuration	3-27
Chapter 3	
Configuring Shortest Path Bridging	4-1
In This Chapter	4-2
Shortest Path Bridging Specifications	4-3
SPBM Parameter Defaults	4-4
SPBM Interface Defaults	4-4
SPBM Service Defaults	4-5
Shortest Path Bridging Overview	4-6
SPBM Shortest Path Trees	4-8
SPB Services	4-12
Sample SPBM Network Topology	4-13
Interaction With Other Features	4-15
Backbone VLANs (VLAN Manager)	4-15
Dynamic Service Access Points (UNP)	4-15
Hardware	4-15
Link Aggregation	4-16
OAM	4-16
Quality of Service (QoS)	4-16
Quick Steps for Configuring SPB	4-18
Quick Steps for Configuring the SPBM Backbone	4-18
Quick Steps for Configuring SPB Services	4-19
Sample Command Configuration	4-19
Configuring SPBM	4-21
Configure the SPBM Backbone (ISIS-SPB)	4-21
Configure SPBM Services	4-21
SPB Configuration Guidelines	4-22
Configuring BVLANS	4-23
Configuring SPB Interfaces	4-25
Configuring Global ISIS-SPB Parameters	4-26
Creating a SPB Service	4-32
Configuring Service Access Points (SAPs)	4-35
Verifying the SPB Backbone and Services	4-42
Verifying the ISIS-SPB Backbone Configuration	4-42
Verifying the SPB Service Configuration	4-42

Chapter 4	Virtual Machine Classification	5-1
	In This Chapter	5-1
	UNP (vNP) Specifications	5-2
	EVB Specifications	5-2
	Server Virtualization Overview	5-3
	Classifying Virtual Machines	5-4
	UNP Overview	5-4
	Using EVB	5-8
	Tracking Virtual Machines	5-9
Appendix A	Software License and Copyright Statements	A-1
	Alcatel-Lucent License Agreement	A-1
	ALCATEL-LUCENT SOFTWARE LICENSE AGREEMENT	A-1
	Third Party Licenses and Notices	A-4
	Index	Index-1

About This Guide

This *OmniSwitch AOS Release 7 Data Center Switching Guide* describes how to set up and monitor supported protocols and software features that comprise the Alcatel-Lucent data center switching architecture. Some of the features described in this guide are purchased as an add-on package to the base switch software.

Supported Platforms

The information in this guide applies only to OmniSwitch 10K and OmniSwitch 6900 switches.

Who Should Read this Manual?

The audience for this user guide are network administrators and IT support personnel who need to configure, maintain, and monitor switches and routers in a live network. However, anyone wishing to gain knowledge on how software features supporting data center switching are implemented in the OmniSwitch 10K and OmniSwitch 6900 switches will benefit from the material in this configuration guide.

When Should I Read this Manual?

Read this guide as soon as you are ready to integrate your OmniSwitch into your network and you are ready to set up the data center switching protocols and features. You should already be familiar with the basics of managing a single OmniSwitch as described in the *OmniSwitch AOS Release 7 Switch Management Guide*.

The topics and procedures in this manual assume an understanding of the OmniSwitch directory structure and basic switch administration commands and procedures.

What is in this Manual?

This switching guide includes a “Data Center Switching Introduction” chapter that provides a description of the Alcatel-Lucent OmniSwitch data center switching architecture and software features. In addition to this introduction, this guide also includes information about configuring the following features;

- Data Center Bridging (DCB) protocols.
- Shortest Path Bridging MAC (SPBM), including SPBM support of Provider Backbone Bridging (PBB) encapsulation and services.
- Universal Network Profile (UNP), specifically for configuring Virtual Network Profiles (vNP) to manage virtual machines within and between data centers.
- Edge Virtual Bridging (EVB) for managing virtual machines created and managed on servers also running the EVB protocol.

What is Not in this Manual?

The configuration procedures in this manual use Command Line Interface (CLI) commands in all examples. CLI commands are text-based commands used to manage the switch through serial (console port) connections or via Telnet sessions. Procedures for other switch management methods, such as web-based (WebView or OmniVista) or SNMP, are outside the scope of this guide.

For information on WebView and SNMP switch management methods consult the *OmniSwitch AOS Release 7 Switch Management Guide*. Information on using WebView and OmniVista can be found in the context-sensitive on-line help available with those network management applications.

This guide provides overview material on software features, how-to procedures, and application examples that will enable you to begin configuring your OmniSwitch. It is not intended as a comprehensive reference to all CLI commands available in the OmniSwitch. For such a reference to all OmniSwitch AOS Release 7 CLI commands, consult the *OmniSwitch CLI Reference Guide*.

How is the Information Organized?

Each chapter in this guide includes sections that will satisfy the information requirements of casual readers, rushed readers, serious detail-oriented readers, advanced users, and beginning users.

Quick Information. Most chapters include a *specifications table* that lists RFCs and IEEE specifications supported by the software feature. In addition, this table includes other pertinent information such as minimum and maximum values and sub-feature support. Some chapters include a *defaults table* that lists the default values for important parameters along with the CLI command used to configure the parameter. Many chapters include *Quick Steps* sections, which are procedures covering the basic steps required to get a software feature up and running.

In-Depth Information. All chapters include *overview sections* on software features as well as on selected topics of that software feature. *Topical sections* may often lead into *procedure sections* that describe how to configure the feature just described. Many chapters include *tutorials* or *application examples* that help convey how CLI commands can be used together to set up a particular feature.

Documentation Roadmap

The OmniSwitch user documentation suite was designed to supply you with information at several critical junctures of the configuration process. The following section outlines a roadmap of the manuals that will help you at each stage of the configuration process. Under each stage, we point you to the manual or manuals that will be most helpful to you.

Stage 1: Using the Switch for the First Time

Pertinent Documentation: *OmniSwitch Getting Started Guide*
Release Notes

A hard-copy *OmniSwitch 10K Getting Started Guide* is included with your switch; this guide provides all the information you need to get your switch up and running the first time. It provides information on unpacking the switch, rack mounting the switch, installing NI modules, unlocking access control, setting the switch's IP address, and setting up a password. It also includes succinct overview information on fundamental aspects of the switch, such as hardware LEDs, the software directory structure, CLI conventions, and web-based management.

At this time you should also familiarize yourself with the Release Notes that accompanied your switch. This document includes important information on feature limitations that are not included in other user guides.

Stage 2: Gaining Familiarity with Basic Switch Functions

Pertinent Documentation: *OmniSwitch Hardware Users Guide*
OmniSwitch AOS Release 7 Switch Management Guide

Once you have your switch up and running, you will want to begin investigating basic aspects of its hardware and software. Information about switch hardware is provided in the *OmniSwitch 10K Hardware Guide*. This guide provides specifications, illustrations, and descriptions of all hardware components, such as chassis, power supplies, Chassis Management Modules (CMMs), Network Interface (NI) modules, and cooling fans. It also includes steps for common procedures, such as removing and installing switch components.

This guide is the primary users guide for the basic software features on a single switch. This guide contains information on the switch directory structure, basic file and directory utilities, switch access security, SNMP, and web-based management. It is recommended that you read this guide before connecting your switch to the network.

Stage 3: Integrating the Switch Into a Network

Pertinent Documentation: *OmniSwitch AOS Release 7 Network Configuration Guide*
OmniSwitch AOS Release 7 Advanced Routing Configuration Guide
OmniSwitch AOS Release 7 Data Center Switching Guide

When you are ready to connect your switch to the network, you will need to learn how the OmniSwitch implements fundamental software features, such as 802.1Q, VLANs, Spanning Tree, and network routing protocols. The *OmniSwitch AOS Release 6 Network Configuration Guide* contains overview information, procedures, and examples on how standard networking technologies are configured on the OmniSwitch.

The *OmniSwitch AOS Release 6 Advanced Routing Configuration Guide* includes configuration information for networks using advanced routing technologies (OSPF and BGP) and multicast routing protocols (DVMRP and PIM-SM).

The *OmniSwitch AOS Release 7 Data Center Switching Guide* includes configuration information for data center networks using virtualization technologies (SPBM and UNP) and Data Center Bridging protocols (PFC, ETC, and DCBX).

Anytime

The *OmniSwitch CLI Reference Guide* contains comprehensive information on all CLI commands supported by the switch. This guide includes syntax, default, usage, example, related CLI command, and CLI-to-MIB variable mapping information for all CLI commands supported by the switch. This guide can be consulted anytime during the configuration process to find detailed and specific information on each CLI command.

Related Documentation

The following are the titles and descriptions of all the related OmniSwitch user manuals:

- *OmniSwitch 10K and OmniSwitch 6900 Getting Started Guides*

Describes the hardware and software procedures for getting an OmniSwitch up and running. Also provides information on fundamental aspects of OmniSwitch software architecture.
- *OmniSwitch 10K and OmniSwitch 6900 Getting Started Guides*

Complete technical specifications and procedures for all OmniSwitch chassis, power supplies, fans, and Network Interface (NI) modules.
- *OmniSwitch CLI Reference Guide*

Complete reference to all CLI commands supported on the OmniSwitch. Includes syntax definitions, default values, examples, usage guidelines and CLI-to-MIB variable mappings.
- *OmniSwitch AOS Release 7 Switch Management Guide*

Includes procedures for readying an individual switch for integration into a network. Topics include the software directory architecture, image rollback protections, authenticated switch access, managing switch files, system configuration, using SNMP, and using web management software (WebView).
- *OmniSwitch AOS Release 7 Network Configuration Guide*

Includes network configuration procedures and descriptive information on all the major software features and protocols included in the base software package. Chapters cover Layer 2 information (Ethernet and VLAN configuration), Layer 3 information (routing protocols, such as RIP and IPX), security options (authenticated VLANs), Quality of Service (QoS), link aggregation, and server load balancing.
- *OmniSwitch AOS Release 7 Advanced Routing Configuration Guide*

Includes network configuration procedures and descriptive information on all the software features and protocols included in the advanced routing software package. Chapters cover multicast routing (DVMRP and PIM-SM), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP).
- *OmniSwitch AOS Release 7 Data Center Switching Guide*

Includes and introduction to the OmniSwitch data center switching architecture as well as network configuration procedures and descriptive information on all the software features and protocols that support this architecture. Chapters cover Shortest Path Bridging MAC (SPBM), Data Center Bridging (DCB) protocols, Virtual Network Profile (vNP), and the Edge Virtual Bridging (EVB) protocol.
- *OmniSwitch Transceivers Guide*

Includes SFP and XFP transceiver specifications and product compatibility information.
- *Technical Tips, Field Notices*

Includes information published by Alcatel's Customer Support group.
- *Release Notes*

Includes critical Open Problem Reports, feature exceptions, and other important information on the features supported in the current release and any limitations to their support.

Technical Support

An Alcatel-Lucent service agreement brings your company the assurance of 7x24 no-excuses technical support. You'll also receive regular software updates to maintain and maximize your Alcatel-Lucent product's features and functionality and on-site hardware replacement through our global network of highly qualified service delivery partners.

With 24-hour access to Alcatel-Lucent's Service and Support web page, you'll be able to view and update any case (open or closed) that you have reported to Alcatel-Lucent's technical support, open a new case or access helpful release notes, technical bulletins, and manuals.

Access additional information on Alcatel-Lucent's Service Programs:

Web: service.esd.alcatel-lucent.com

Phone: 1-800-995-2696

Email: esd.support@alcatel-lucent.com

1 Understanding Data Center Switching

Alcatel-Lucent helps enterprises address the challenges and ongoing transformation of data center networks while delivering a high-quality user experience for new real-time applications, greater agility in deploying new applications, seamless integration of public cloud services, and reduced data center costs. To deliver on these capabilities, Alcatel-Lucent provides a unique blueprint for data center switching that brings together three core innovations:

- **OmniSwitch Pod.** A unique architecture concept for top-of-rack switches that can provide server-to-server connectivity without the need for a core switch to carry traffic. The pod is a highly dense architecture that ensures low latency and high performance between servers connected to the same pod.
- **Data Center Mesh.** The mesh consists of pods connected to each other and to core switches to bring together thousands of server-facing ports with low aggregate end-to-end latency. Implementing the data center mesh architecture allows enterprises to create virtual data centers supporting defined work-groups or departments.
- **Virtual Network Profile (vNP).** A type of Universal Network Profile (UNP) that the administrator configures specifically for virtual machine classification. A virtual machine is bound to a vNP to ensure consistent application of network access controls and QoS policies when a virtual machine is initially detected or moves to a different network location.

In This Chapter

This chapter contains an overview of the OmniSwitch components and features of the Alcatel-Lucent Data Center Switching solution. It provides a general example for configuring the related data center software applications through the Command Line Interface (CLI). CLI commands are used in the configuration examples; for more details about the syntax of commands, see the *OmniSwitch CLI Reference Guide*.

The following topics are included in this chapter:

- [“Data Center Switching Components” on page 1-2.](#)
- [“Data Center Mesh Configuration Example” on page 1-4.](#)

Data Center Switching Components

Key components of the OmniSwitch data center switching framework include an OmniSwitch pod, the data center mesh, and virtual network profiles (vNP). The OmniSwitch also supports the use of other network virtualization technology features, such as Shortest Path Bridging (SPB), Data Center Bridging (DCB), and Virtual Chassis.

OmniSwitch Pod

The Alcatel-Lucent OmniSwitch pod is an architecture concept for top-of-rack switches that provides server-to-server (east-west and north-south) connectivity without the need for a core switch to carry traffic. The pod is a highly dense architecture that ensures low latency and high performance between servers connected to the same pod.

The architecture of the OmniSwitch pod provides the scalability necessary to handle changing data center demands. A pod may initially only consist of one or two such switches, but as the demand to handle more and more traffic grows, additional switches are easily added into the pod configuration.

Figure 1 shows some examples of OmniSwitch pod architectures.




Figure 1: OmniSwitch Pod Examples

Each pod is a switching fabric where every switch could be connected to every other switch. When pods are then connected together with other pods and core switches, a data center mesh is formed.

Data Center Mesh

The Alcatel-Lucent data center mesh consists of pods connected to each other and to core switches to bring together thousands of server-facing ports with low aggregate end-to-end latency and high resiliency. Implementing the data center mesh architecture allows enterprises to create virtual data centers supporting virtualized workgroups or departments. Figure 2 shows an example of a data center mesh that consists of four OmniSwitch 6900 pods and two OmniSwitch 10K core switches:




Figure 2: Alcatel-Lucent Data Center Mesh

Virtual Network Profile

The Virtual Network Profile (vNP) feature facilitates the discovery and movement of virtual machines (VMs). A vNP is a type of Universal Network Profile (UNP) that is configured specifically for machine classification, especially VMs. A virtual machine is bound to a vNP to ensure consistent application of network access controls and QoS policies when a virtual machine is initially detected or moved to a different network location.

A vNP classifies VMs in the same manner as any other device connected to a UNP port. Once a VM is assigned to a vNP, the VM traffic is bound to the VLAN or service defined in the profile. In addition, any optional QoS policies associated with the profile are also applied to the VM traffic.

Network Virtualization Technology

OmniSwitch data center switching leverages the use of the following AOS Release 7 software features to facilitate a network virtualization solution:

- **Shortest Path Bridging MAC (SPBM)** to virtualize the data center mesh.

SPBM extends Layer 2 across the data center mesh while maintaining a loop-free network. All connections between all switches in the topology remain active (no blocking of redundant links).

SPBM uses the Provider Backbone Bridge (PBB) network model to encapsulate (using IEEE 802.1ah headers) and tunnel customer traffic through the network backbone. The shortest path trees upon which the PBB network infrastructure operates are determined using a version of the Intermediate System-to-Intermediate System (IS-IS) link state protocol that supports TLV extensions for SPB (ISIS-SPB).

For more information about SPBM, see [Chapter 3, “Configuring Shortest Path Bridging.”](#)

- **Data Center Bridging (DCB)** protocols to convert Ethernet into a lossless transport to support a reliable storage area network fabric within the data center mesh. DCB protocols are implemented using embedded profiles in the same manner that QoS QSet profiles are applied.

The DCB profiles are based on the 802.1Q-REV/D1-5 standard to define how the switch classifies different traffic types and priority mappings and then groups those types into traffic classes. Profiles also specify the Traffic Class Flow (TCF), which is LL (lossless; PFC initiated upstream) or nL (lossy; PFC not initiated upstream).

For more information about DCB, see [Chapter 2, “Configuring Data Center Bridging.”](#)

- **Virtual Chassis (VC)** is a group of chassis managed through a single management IP address.

VC provides both node level and link level redundancy for devices connecting to the aggregation layer via standard 802.3ad link aggregation mechanisms. See [Chapter 9, “Configuring Virtual Chassis.”](#) for more information.

Data Center Mesh Configuration Example

Traditional data center networks consist of a 3-tier approach: access, aggregation, and core. The high-speed, low-latency capability of the OmniSwitch 6900 provides the ability to combine the access and aggregation tiers.

The application example in this section consists of two data centers that incorporate the use of the OmniSwitch data center switching mesh architecture. Each pod shown in Figure 5 interconnects six OS6900 switches delivering 240 server-facing ports while maintaining low latency between servers in the same pod.

In this example, each OmniSwitch 6900 pod is connected to two OmniSwitch 10K core switches to form the data center mesh. In turn, the OmniSwitch 10K core switches connect to Alcatel-Lucent Service Routers (7750) platforms to provide interconnectivity between the two data centers through an MPLS cloud over Layer 2 virtual LAN services.

The following diagram illustrates the example data center switching network that implements the OmniSwitch pod/mesh architecture combined with supported virtualization technologies.

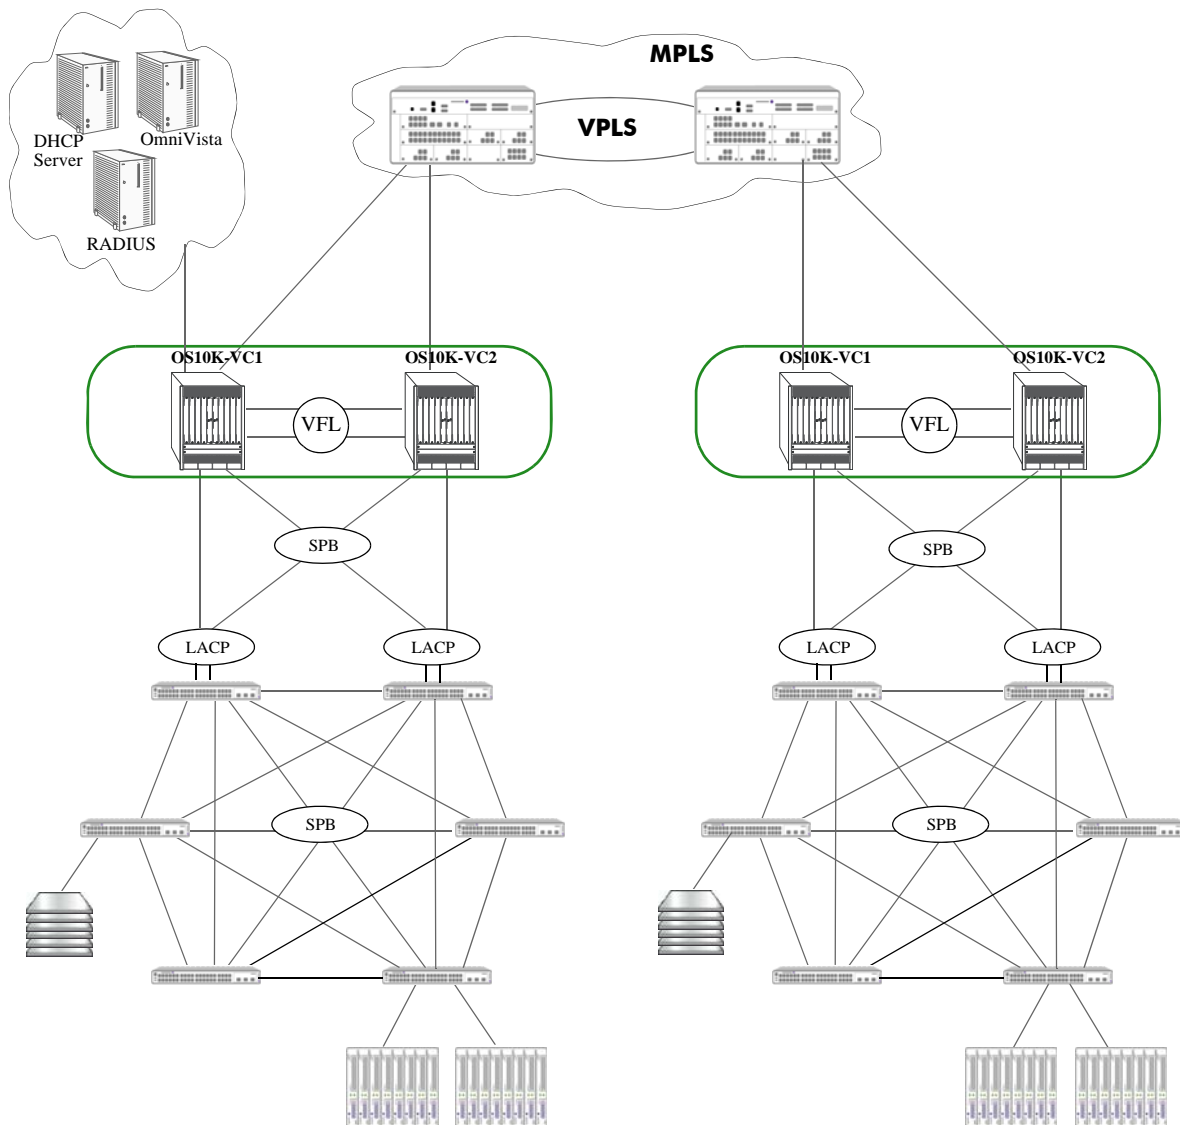


Figure 5: Data Center Mesh Example

In this data center mesh configuration example:

- Two separate data center configurations are depicted, each using a single OS6900 pod and two OS10K switches for the core.
- Each pod consists of six switches. The pod is formed by connecting each switch to every one of the other six switches. In other words, each switch is configured with five links that connect the switch to the pod.
- The core tier for each data center is made up of two OS10K switches. Virtual Chassis (VC) is configured for each pair of OS10Ks to create a chassis group that is managed through a single management IP address. VC provides both node level and link level redundancy for the 10K chassis group connecting to the aggregation (pod) layer via standard 802.3ad link aggregation mechanisms.

The following configuration snapshot shows a sample VC configuration for the core switches:

```
OS10K-VC1-> show configuration vcm-snapshot chassis-id 1
! Virtual Chassis Manager:
virtual-chassis chassis-id 1 configured-chassis-id 1
virtual-chassis chassis-id 1 vf-link 0 create
virtual-chassis chassis-id 1 vf-link 0 member-port 1/2/1
```

- Shortest Path Bridging MAC (SPBM) is used in each data center mesh (pod plus core) to define sets of loop-free shortest path trees (SPTs) through the network. Each switch serves as the SPT root for all traffic entering the switch, thus allowing the switch to provide the shortest path to every other switch.

The bridging methodology that allows each switch to serve as its own root is enforced through the use of SPBM backbone VLANs (BVLANS). The BVLAN is a transport VLAN for SPBM services and SPT calculations.

The following configuration snapshot shows a sample SPBM bridging configuration. Each switch that will participate in the SPBM domain (pod and core) must have the same BVLAN configuration. SPBM merges the core and pod together into the same virtual mesh domain.

```
OS6900-1-> show configuration snapshot spb-isis
! SPB-ISIS:
spb isis bvlan 4001 ect-id 1
spb isis bvlan 4002 ect-id 2
spb isis control-bvlan 4001
spb isis interface port 1/1/2-5
spb isis interface port 1/1/7
spb isis interface port 2/1/2-5
spb isis interface linkagg 2
spb isis admin-state enable
```

- SPBM Ethernet services are bound to the SPT bridging configuration and are used to encapsulate and tunnel data through the data center mesh. SPBM services are configured on any OmniSwitch in the mesh that will originate or terminate a service. This usually is on switches that are connected to host devices or that connect to other networks where traffic will enter or leave the mesh.

The OmniSwitch Service Manager feature is used to build the SPBM service architecture layer within the SPBM mesh domain. Provisioning a service requires the configuration of three logical entities: a service instance identifier (I-SID), service access port, and a service access point (SAP).

The following configuration snapshot shows a sample SPB service configuration:

```
OS6900-1-> show configuration snapshot svcmgr
! SVCMgr:
service access port 1/11
service access port 1/12
service access port 1/13
service spb 1001 isid 1001 bvlan 4001 admin-state enable
service spb 1002 isid 1002 bvlan 4001 admin-state enable
service spb 1001 sap port 1/11:1001 admin-state enable
service spb 1001 sap port 1/12:1001 admin-state enable
service spb 1001 sap port 1/13:1001 admin-state enable
service spb 1002 sap port 1/11:1002 admin-state enable
service spb 1002 sap port 1/12:1002 admin-state enable
service spb 1002 sap port 1/13:1002 admin-state enable
```

- When a physical switch port comes up, a QSet instance (a set of eight queues) is automatically associated with the port for unicast traffic. In addition, a default Data Center Bridging profile (DCP 8) is automatically assigned to the QSI.

The DCP 8 profile applies strict priority scheduling with lossy traffic class management to each of the

eight egress port queues. However, ports connected to devices running critical applications that require lossless Ethernet, such as the storage and server hosts shown in Figure 5, can be assigned to one of the other 10 pre-defined profiles or a user-configured profile that will provide the necessary class of service for that device.

The following configuration snapshot shows a sample DCB configuration. In this sample, a custom profile (DCP 11) is created with custom settings for the traffic classes. DCP 11 is then assigned to ports 1/1 and 1/9-10, replacing the default DCP 8 assignment. In addition, ports 1/12 and 1/34-35 are assigned to DCP 7 and 9. All other ports on the switch are using default DCP 8.

```
OS6900-1-> show configuration snapshot vfc
! Virtual Flow Control:
qos qsp dcb 11 import qsp dcb "dcp-9"
qos qsp dcb "dcp-11" tc 1 pfc flow-type nLL
qos qsp dcb "dcp-11" tc 0 pfc flow-type nLL
qos qsp dcb "dcp-11" tc 2 pfc flow-type nLL
qos qsp dcb "dcp-11" tc 4 pfc flow-type nLL
qos qsp dcb "dcp-11" tc 5 pfc flow-type nLL
qos qsp dcb "dcp-11" tc 6 pfc flow-type nLL
qos qsp dcb "dcp-11" tc 7 pfc flow-type nLL
qos qsi port 1/1 qsp dcb "dcp-11"
qos qsi port 1/9-10 qsp dcb "dcp-11"
qos qsi port 1/12 qsp dcb "dcp-7"
qos qsi port 1/34-35 qsp dcb "dcp-9"
```

- Virtual Network Profile (vNPs) are configured on each switch that is connected to a server to facilitate VM discovery and mobility within the local data center or mobility between the local and remote data center. In addition, Universal Network Profile (UNP) functionality is enabled on each of the server connections to activate vNP classification of VMs.

The vNPs are used to assign VMs to SPBM services that span the data center mesh. When a VM is discovered or migrates to a new location, the assigned vNP applies the necessary access controls and any QoS policies specifically defined for the VM.

The following configuration snapshot shows a sample vNP configuration.

```
OS6900-1-> show configuration snapshot da-unn
! DA-UNP:
unn customer-domain 1 description "AT&T"
unn customer-domain 2 description "Sprint"
unn port 1/15 port-type spb-access
unn port 1/15 unn-customer-domain 1
unn port 1/16 port-type spb-access
unn port 1/16 unn-customer-domain 1
unn port 1/17 port-type spb-access
unn port 1/17 unn-customer-domain 2
unn port 1/18 port-type spb-access
unn port 1/18 unn-customer-domain 2
unn spb-profile unn_for_spb1 tag-value 115:150 isid 6000 bvlan 4001
unn classification vlan-tag 125 spb-profile unn_for_spb1
unn classification vlan-tag 225 spb-profile unn_for_spb1
```


2 Configuring Data Center Bridging

Data Center Bridging (DCB) is a set of standards that extend Ethernet capabilities to support the convergence of storage and data in today's virtualized networks. The Alcatel-Lucent implementation of DCB for the OmniSwitch supports the following DCB standards:

- **Priority-Based Flow Control (PFC)**—based on the IEEE 802.1Qbb standard, PFC pauses traffic based on congestion priority instead of blocking the entire link when congestion occurs. Allows lossless and lossy traffic with different priorities on the same physical port.
- **Enhanced Transmission Selection (ETS)**—based on the IEEE 802.1Qaz standard, ETS groups related traffic into priority groups (traffic classes) to which bandwidth guarantees and scheduling are applied.
- **Data Center Bridging Exchange (DCBx)**—exchanges and negotiates PFC and ETS information between two directly connected peer devices.

This implementation of the DCB enhanced Ethernet protocols uses embedded profiles to apply the PFC, ETS, and DCBx configuration to traffic flows. This approach is similar to how QSet profiles (QSPs) are used to apply the QoS configuration for bandwidth management and egress port queue scheduling.

In This Chapter

This chapter describes DCB in general and how DCB profiles and port configurations are applied to the switch. It provides information about configuring DCB through the Command Line Interface (CLI). CLI commands are used in the configuration examples; for more details about the syntax of commands, see the *OmniSwitch CLI Reference Guide*.

The following topics and procedures are included in this chapter:

- [“DCB Specifications” on page 2-2.](#)
- [“DCB Defaults” on page 2-3.](#)
- [“Data Center Bridging Overview” on page 2-5.](#)
- [“Interaction with Other Features” on page 2-13](#)
- [“Using DCB Profiles” on page 2-15](#)
- [“Configuring DCB Profiles” on page 2-21](#)
- [“Configuring DCBx Port Parameters” on page 2-23.](#)
- [“Multicast and Unicast Traffic Distribution” on page 2-24.](#)

DCB Specifications

The DCB functionality described in this chapter is supported on the OmniSwitch 10K and OmniSwitch 6900 switches, unless otherwise stated in the following DCB Specifications table or specifically noted within any other section of this chapter. Note that any maximum limits provided in the DCB Specifications table are subject to available system resources.

Platforms Supported	OmniSwitch 6900 and the following OmniSwitch 10K modules: <ul style="list-style-type: none"> • OS10K-QNI-U8 (8 x 40G) • OS10K-QNI-U4 (4 x 40G) • OS10K-XNI-U32E (32 x 10G) • OS10K-XNI-U16E (16 x 10G) • OS10K-XNI-U16L (8 x 10G, 8 x 1G)
OmniSwitch Software License	Data Center
IEEE Standards Supported	802.1Qbb— <i>Priority-based Flow Control</i> 802.1Qaz D2.5— <i>Enhanced Transmission Selection</i> 802.1Qaz D2.5— <i>Data Center Bridging Exchange</i> 802.1Q-REV/D1.5— <i>Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks</i>
Maximum number of DCB profiles	128 profiles: <ul style="list-style-type: none"> • Profiles 1–10 are predefined, with profile 8 serving as the default profile for all ports. • Profiles 11–128 are reserved for user-defined (custom) profiles.
Maximum number of lossless queues (priorities)	110 total per switch (OmniSwitch 6900) 8 per-port (OmniSwitch 10K)
DCB TLVs supported	ETS Configuration ETS Recommendation PFC Configuration (Application Priority TLV not supported)

DCB Defaults

Traffic class management and related QoS functions are implemented using a Queue Set (QSet) framework. Each port and link aggregate is associated with a set of eight egress queues, referred to as a Queue Set Instance (QSI). Each QSI is associated with DCB profile 8 (DCP 8) by default.

A DCP defines both the Data Center Bridging Exchange (DCBX) control parameters and the traffic class management parameters that are applied to the eight queues associated with the QSI. See [“Configuring DCB Profiles” on page 2-21](#) for more information.

The following are the default DCBX and traffic class management settings applied with DCP 8:

DCB Profile 8 (Default Profile)							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 8							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%)		TC Flow Control (TF)	TC Scheduler (TS)
0	BE	Best Effort	0	0	100	Lossy	Strict Priority
1	BK	Background	1	0	100	Lossy	Strict Priority
2	EE	Excellent Effort	2	0	100	Lossy	Strict Priority
3	CA	Critical Applications	3	0	100	Lossy	Strict Priority
4	VI	Video (<100ms latency)	4	0	100	Lossy	Strict Priority
5	VO	Voice (<10ms latency)	5	0	100	Lossy	Strict Priority
6	IC	Internetwork Control	6	0	100	Lossy	Strict Priority
7	NC	Network Control	7	0	100	Lossy	Strict Priority

Note. QSet profiles and DCB profiles are mutually exclusive on the same port. If the OmniSwitch Data Center software license is installed, then DCB profiles are used by default. If this license is not installed, then QSet profiles are used by default. See [“Using DCB Profiles” on page 2-15](#) for more information.

DCB Port Defaults

Parameter Description	Command	Default
The DCBx administrative status for the port.	qos qsi dcb dcbx admin-state	Enabled
The transmission status for the PFC configuration TLV	qos qsi dcb dcbx pfc tlv	Enabled
The PFC defense mode status	qos qsi dcb dcbx pfc defense	Enabled
The PFC “willing” to negotiate status	qos qsi dcb dcbx pfc willing	Yes
The transmission status for the ETS configuration TLV	qos qsi dcb dcbx ets config-tlv	Enabled
The transmission status for the ETS recommended TLV	qos qsi dcb dcbx ets recommended TLV	Enabled
The ETS “willing” to negotiate status	qos qsi dcb dcbx ets willing	Yes

Data Center Bridging Overview

Convergence of data and storage into Ethernet has become more prevalent due to the combined data and storage bandwidth requirements of virtualized networks. Data Center Bridging (DCB) technology combined with the Alcatel-Lucent Data Center Mesh architecture provides a unified framework upon which different types of network traffic are converged into a single transport layer.

For example, each of the following traffic types can coexist in a converged network without imposing serious restrictions on the performance of the other traffic types:

- **Storage Area Network (SAN) traffic**—A SAN provides access to a shared storage solutions. Requires transmission of traffic with no packet loss (Lossless).
- **Local Area Network (LAN) traffic**—A LAN provides end-user access to server-based network applications, such as email, database, and web servers. Requires best effort transmission of traffic for important applications. Can tolerate some traffic loss (Lossy).
- **Inter-Process Communication (IPC) traffic**—IPC refers to the mechanisms or activities that provide communications and data sharing between applications. Requires the highest priority for low latency transmission.

The Alcatel-Lucent OmniSwitch supports the following Data Center Bridging (DCB) protocols to provide a single, shared infrastructure for reliable delivery of data and storage over Ethernet:

- **Priority-Based Flow Control (PFC)**—Based on the IEEE 802.1Qbb standard, PFC pauses traffic based on congestion priority instead of blocking the entire link when congestion occurs. Allows lossless and lossy traffic with different priorities on the same physical port.
- **Enhanced Transmission Selection (ETS)**—Based on the IEEE 802.1Qaz standard, ETS provides a common framework for dynamic bandwidth management. ETS groups related traffic into priority groups (Traffic Classes) to which bandwidth guarantees and scheduling are applied.
- **Data Center Bridging Exchange (DCBX)**—Based on the IEEE 802.1Qaz standard, DCBX uses the Link Layer Discovery Protocol (LLDP) to exchange and negotiate PFC and ETS information between two directly connected peer switches. Enabled by default, DCBX is responsible for auto-negotiation and auto-configuration of link parameters for DCB functions.

Priority-Based Flow Control Overview

Priority-based Flow Control (PFC) is the evolution of the IEEE 802.3 Ethernet PAUSE mechanism that uses pause frames to signal the other end of a link to stop sending traffic when buffer congestion is detected on any one of the ingress port queues. The problem with the 802.3x approach is that when congestion occurs for only one type of traffic on a specific queue, the transmission of all traffic on the port is stopped.

PFC offers a more granular approach that is needed to support the convergence of various traffic types over the same link. Instead of pausing all traffic when ingress port congestion occurs, PFC only pauses traffic tagged with a specific Class of Service (CoS) priority value while allowing traffic assigned to other priority values to continue (or get dropped when congestion occurs) on the same link.

In essence, using PFC divides an Ethernet link into eight virtual lanes. Each virtual lane represents one of the CoS priority values (0–7). When buffer congestion occurs for traffic marked with one of these priority values, PFC initiates pause frames back to the source of that specific traffic. The source then holds back the traffic until the congestion clears (bandwidth allocated for that priority becomes available). As shown in Figure 1, PFC detected congestion for traffic marked with priority 4, so only that traffic is

paused and packet loss is avoided. All other traffic marked with different priorities continues to forward on the link until congestion occurs for those priorities as well.




Figure 1: PFC Example

To determine how traffic is treated when assigned to one of these eight priority values, the pre-defined OmniSwitch DCB profiles use the following general IEEE recommendations for mapping traffic types to the CoS priority values:

- Priority 1, BK—Background {Bulk transfers}
- Priority 0, BE—Best Effort {unprioritized user applications}
- Priority 2, EE—Excellent Effort {CEO's best effort}
- Priority 3, CA—Critical Applications {guaranteed minimum bandwidth}
- Priority 4, VI—Video {low latency < 100ms}
- Priority 5, VO—Voice {low latency < 10ms}
- Priority 6, IC—Internetwork Control {guaranteed delivery}
- Priority 7, NC—Network Control {guaranteed delivery}

Using PFC for OmniSwitch Flow Control

The main purpose of PFC is to provide a lossless transport for traffic sensitive to packet loss, such as storage area network (SAN), LAN backup, or management traffic, while at the same time allowing packet-drop congestion management for other traffic types on the same link. The OmniSwitch implementation of DCB accomplishes this through the use of embedded profiles that define the PFC configuration for each priority queue associated with each switch port.

There are 10 pre-defined profiles available, and the ability to create additional custom profiles is also supported. Each profile specifies which priority queue requires lossless (no packet drop) or lossy (packet drop) congestion management. When a switch port comes up, a single, pre-defined DCB profile is associated with that port by default. Currently DCB profile 8 is the default profile applied to each port. This profile defines lossy flow control for all eight queues (virtual lanes).

The following general steps are required to configure a lossless transport lane for a specific type of traffic:

- 1** Select a pre-defined DCB profile or create a custom profile that defines lossless flow control for a specific CoS priority value.
- 2** Determine the end-to-end lossless data path through the OmniSwitch network and apply the profile identified in Step 1 to each port in that path.
- 3** Make sure that the traffic requiring lossless transmission is marked with the same CoS priority value to which the DCB profile will apply lossless flow control. The switch will only pause frames sent with the same priority value as the lossless priority value specified in the DCB profile.
- 4** Make sure that all ingress ports are configured as QoS trusted ports with the default classification set to 802.1p.

For more information about DCB profiles, see [“Using DCB Profiles” on page 2-15](#).

For more information about CoS 802.1p priority bit classification and marking, see [“How Traffic is Classified and Marked” on page 24-6 in Chapter 24, “Configuring QoS,” of the *OmniSwitch AOS Release 7 Network Configuration Guide*](#).

Enhanced Transmission Selection Overview

Priority-based Flow Control (PFC) uses the Class of Service (COS) 802.1p priority values (0–7) to determine which traffic to pause when congestion occurs. Enhanced Transmission Selection (ETS) groups these priorities (traffic types) into Traffic Classes and then allocates a percentage of minimum bandwidth to each class.

Traffic Classes are groups of priorities that have similar traffic handling requirements (for example, LAN, SAN, and IPC). The amount of bandwidth assigned to each Traffic Class is a minimum guarantee of available bandwidth. Available bandwidth is defined as the amount of link bandwidth remaining after Traffic Classes not subject to ETS (for example, Strict Priority is used instead) or vendor-specific requirements are serviced. In addition,

- Available bandwidth for a Traffic Class is configurable in 1% increments.
- The maximum deviation of available bandwidth is 10% when all of the following is true:
 - > Only ETS traffic is flowing.
 - > No PFC pause frames have been received.
 - > All ETS Traffic Classes are offered enough load to consume their share of allocated bandwidth.
- Any unused portion of the minimum bandwidth is made available for use by other classes, but only until the original Traffic Class needs the bandwidth again. At such time, PFC (if enabled) will pause the required traffic to minimize packet loss until the bandwidth is regained.

Using ETS for OmniSwitch Bandwidth Allocation

The OmniSwitch implementation of DCB uses embedded profiles to define the ETS and PFC configuration for each priority queue associated with each switch port. ETS combines the CoS priorities used by PFC into Traffic Classes. Each DCB profile defines the ETS grouping of such priority values into Traffic Classes to which the specified bandwidth allocation, flow control (lossless or lossy), and scheduling are applied.

ETS requires a minimum of three Traffic Classes per port but supports a maximum of eight Traffic Classes per port. The use of DCB profiles to apply the ETS configuration facilitates the correct mapping of traffic types and priorities into Traffic Classes according to the IEEE 802.1Q-REV/D1-5 standard.

There are 10 pre-defined DCB profiles available, and the ability to create additional custom profiles is also supported. The following table shows the grouping of traffic types into a Traffic Class (TC) as applied by the pre-defined DCB profiles (see [“Using DCB Profiles” on page 2-15](#)).

Table 1: DCB Profile Traffic Classes

DCB Profile	Traffic Type Allocation	TC Priority Map
1	TC-0: Best Effort, Background, Excellent Effort, Critical Applications	0, 1, 2, 3
	TC-1: Voice, Video	5, 4
	TC-2: Network Control, Internetwork Control	7, 6
2	TC-0: Best Effort, Background	0, 1,
	TC-1: Critical Applications, Excellent Effort	3, 2
	TC-2: Voice, Video	5, 4
	TC-3: Network Control, Internetwork Control	7, 6
3	TC-0: Best Effort, Background	0, 1,
	TC-1: Critical Applications, Excellent Effort	3, 2
	TC-2: Voice, Video	5, 4
	TC-3: Internetwork Control	6
	TC-4: Network Control	7
4	TC-0: Background	1
	TC-1: Best Effort	0
	TC-2: Critical Applications, Excellent Effort	3, 2
	TC-3: Voice, Video	5, 4
	TC-4: Internetwork Control	6
5	TC-0: Background	1
	TC-1: Best Effort	0
	TC-2: Excellent Effort	2
	TC-3: Critical Applications	3
	TC-4: Voice, Video	5, 4
	TC-5: Internetwork Control	6
6	TC-6: Network Control	7
	TC-0: Background	1
	TC-1: Best Effort	0
	TC-2: Excellent Effort	2
	TC-3: Critical Applications	3
	TC-4: Video	4
	TC-5: Voice	5
TC-6: Internetwork Control	6	
	TC-7: Network Control	7

DCB Profile	Traffic Type Allocation	TC Priority Map
7	TC-0: Best Effort	0
	TC-1: Background	1
	TC-2: Excellent Effort	2
	TC-3: Critical Applications	3
	TC-4: Video	4
	TC-5: Voice	5
	TC-6: Internetwork Control	6
	TC-7: Network Control	7
8	TC-0: Best Effort	0
	TC-1: Background	1
	TC-2: Excellent Effort	2
	TC-3: Critical Applications	3
	TC-4: Video	4
	TC-5: Voice	5
	TC-6: Internetwork Control	6
	TC-7: Network Control	7
9	TC-0: Background	1
	TC-1: Best Effort	0
	TC-2: Excellent Effort	2
	TC-3: Critical Applications	3
	TC-4: Video	4
	TC-5: Voice	5
	TC-6: Internetwork Control	6
	TC-7: Network Control	7
10	TC-0: Background	1
	TC-1: Best Effort	0
	TC-2: Excellent Effort	2
	TC-3: Critical Applications	3
	TC-4: Video	4
	TC-5: Voice	5
	TC-6: Internetwork Control	6
	TC-7: Network Control	7

Figure 2 further illustrates ETS Traffic Class groupings by showing the Traffic Class minimum bandwidth guarantees as defined and applied by DCB profile 2:




Figure 2: ETS Traffic Classes (DCB Profile 2)

In the Figure 2 example,

- Traffic marked with priorities 0 and 1 is grouped into Traffic Class 0, which is assigned a minimum bandwidth guarantee of 27% of the available bandwidth.
- Traffic marked with priorities 2 and 3 is grouped into TC-1, which is assigned a minimum bandwidth guarantee of 40% of the available bandwidth.
- Traffic marked with priorities 4 and 5 is grouped into TC-2, which is assigned a minimum bandwidth guarantee of 33% of the available bandwidth.
- Traffic marked with priorities 6 and 7 is grouped into TC-3, but the TC-3 minimum bandwidth is set to zero. This is because DCB profile 2 applies Strict Priority scheduling to TC-3, not ETS. As a result, virtual lanes (priority queues) 6 and 7 are serviced first then ETS allocates the remaining bandwidth for TC-0, TC-1, and TC-2 priority queues.

For more information about DCB profiles, see [“Using DCB Profiles” on page 2-15](#).

For more information about 802.1p CoS priority bit classification and marking, see [“How Traffic is Classified and Marked” on page 24-6 in Chapter 24, “Configuring QoS,” of the *OmniSwitch AOS Release 7 Network Configuration Guide*](#).

Data Center Bridging Exchange Overview

The Data Center Bridging Exchange (DCBx) protocol communicates the DCB capabilities of each OmniSwitch and exchanges PFC and ETS configuration details with directly connected peer devices. In addition, DCBx functionality helps to ensure a consistent configuration across the network.

- DCBx discovers the DCB capabilities of directly connected peers (for example, does the port on the other end of the link support PFC and is it active).

- Misconfigurations are detected when DCB configuration parameter values are exchanged between two peers and the parameters are not the same on both peers.
- When a configuration mismatch is detected and both peers are “willing”, an automatic negotiation takes place to reconcile a common operational configuration for both peers.

DCBx uses the IEEE 802.1AB Link Layer Discover Protocol (LLDP) to discover and exchange information between two link peers. There are specific LLDP Type-Length-Values (TLVs) that DCBx defines for exchanging PFC and ETS attribute values, such as:

- Is PFC in use on the peer switch.
- What are the priority values that PFC manages.
- Which traffic type priority values are mapped to a specific Traffic Class.
- What are the minimum bandwidth guarantees defined for each Traffic Class.

An OmniSwitch is considered a DCB-capable switch if the Alcatel-Lucent OmniSwitch Data Center software license is installed. When this license is installed, DCBx is automatically enabled on each switch port and DCB profiles apply the PFC and ETS configuration to the QSet instance (QSI) for each port.

The following default DCBx configuration is applied to each DCBx-enabled port:

- PFC and ETS TLVs are enabled. The following TLVs are supported:
 - > PFC Configuration
 - > ETS Configuration
 - > ETS Recommendation
- The “willing” setting is enabled. This setting indicates that the switch is willing to negotiate changes to the operational DCBx configuration on the local port to match the DCBx configuration on the remote port.
- PFC defense mode is enabled. When PFC negotiations fail and the defense mode is enabled, no PFC functionality is applied to any priority but traffic flow is allowed to continue. However, if the defense mode is disabled when negotiations fail, the local PFC configuration is applied.

Note. Configuring DCBX operation for PFC and ETS is supported but dynamic configuration for Layer 2 and Layer 4 applications, such as FCoE and iSCSI via the DCBx Application TLV, is not supported.

DCBx in the Data Center Mesh

The diagram in Figure 3 shows where DCBx is enabled in a sample OmniSwitch Data Center Mesh topology. In this example,

- The OmniSwitch Data Center Software License is installed on each OmniSwitch to enable the DCB features (PFC, ETS, and DCBx) on each switch port.
- DCBx negotiates PFC and ETS parameters between the switches on every port. Manual configuration between an OmniSwitch and a non-OmniSwitch device (servers or switches from other vendors) may require manual configuration of DCB parameters on those devices.
- The servers could be sending SAN traffic at a priority negotiated for lossless and could use a different priority for iSCSI SAN traffic. Note that any application requiring lossless access to the network is supported as long as the traffic is stamped with the correct priority that is negotiated for lossless.

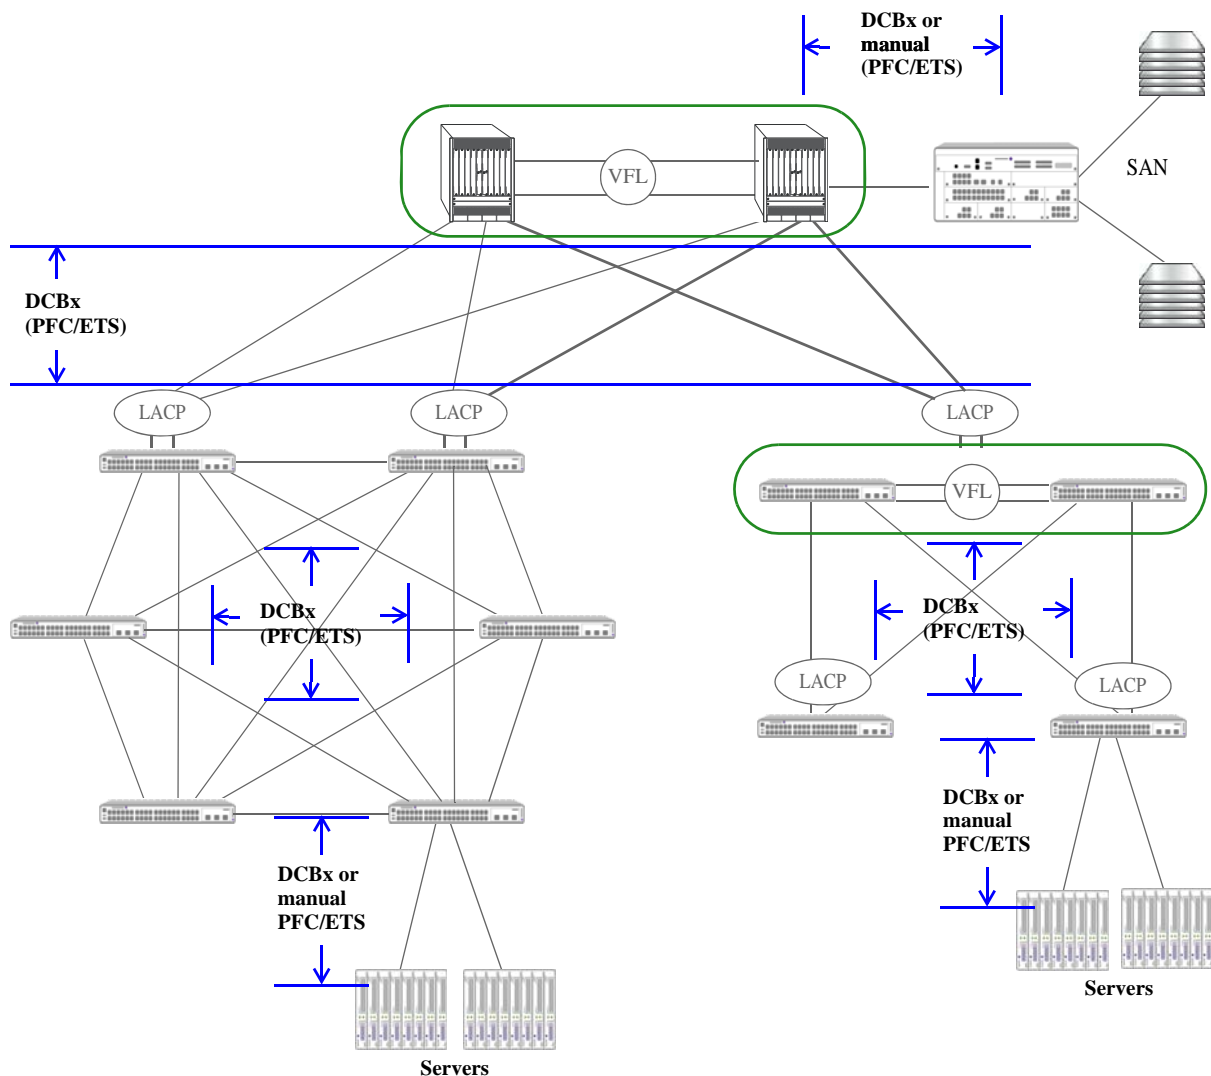


Figure 3: DCBx Example

For more information about DCBx, see [“Configuring DCBx Port Parameters”](#) on page 2-23.

For more information about DCB profiles, see [“Using DCB Profiles”](#) on page 2-15.

Interaction with Other Features

This section contains important information about how other OmniSwitch features interact with Data Center Bridging (DCB) features. Refer to the specific chapter for each feature to get more detailed information about how to configure and use the feature.

QoS

- This implementation of DCB provides enhanced QoS congestion and bandwidth allocation to support multiple traffic types on the same Ethernet link. However, DCB profiles are used to apply the DCB configuration instead of QoS profiles. These two profile types are mutually exclusive on the same port; only one or the other profile type is applied to a port at any given time.
- Even though a port is only associated with a single DCB profile or a single QoS profile, it is possible to have a mixture of both profile types on different ports within the same switch. For example, when a DCB-licensed switch boots up, all ports are associated with a DCB profile by default, but the profile association for any port can be changed to a QoS profile without disrupting the DCB profile associations of other ports.
- The DCB protocols, Priority-based Flow Control (PFC) and Enhanced Transmission Selection (ETS), use the Class of Service (CoS) 802.1p markings found in the frame header to define traffic groups. These markings are the result of QoS classification that occurs prior to and separately from the application of DCB functionality. DCB does not replace existing QoS classification and enqueueing mechanisms.

For more information about DCB and QoS profiles, see [“Using DCB Profiles” on page 2-15](#) and [“Multicast and Unicast Traffic Distribution” on page 2-24](#).

For more information about 802.1p CoS priority bit classification and marking, see [“How Traffic is Classified and Marked” on page 24-6](#) in [Chapter 24, “Configuring QoS,”](#) of the *OmniSwitch AOS Release 7 Network Configuration Guide*.

Shortest Path Bridging

In order to support DCBX on an SPB access port (for example, between a server and an OmniSwitch), create a Layer 2 profile to allow the access port to participate as an LLDP (802.1ab) protocol peer to the server. If this is not done, then LLDP packets will be treated as traffic and DCBX will not be negotiated between the access port and server.

The following is an example of an SPB configuration that creates Layer 2 profile “DCBX” and associates that profile with access port 1/1/6:

```
-> show configuration snapshot svcmgr
! SVCMgr:
service l2profile DCBX 802.1ab peer
service access port 1/1/6 l2profile DCBX
```

SNMP

If SNMP performs sets on the standard MIBs at the port level and the port is associated with:

- A custom DCB profile that is only associated with that one port, then the changes are made on that profile.

- A custom DCB profile that is associated with multiple ports, a new custom profile is created with the changes and the new profile is associated only with that port. The new profile is accessible through the OmniSwitch Command Line Interface (CLI).
- A default DCB profile, a new custom profile is created with the changes and the new profile is associated only with that port. The new profile is accessible through the OmniSwitch CLI.

For more information about DCB profiles, see [“Using DCB Profiles” on page 2-15](#).

Virtual Chassis

- DCBx is not supported on VF-Links.
- Only two profiles can be applied on the VF-links: DCB 8 (applied by default, all priorities are lossy) and DCB 9 (all priorities are lossless, PFC is enabled).

PFC Support Across the VFL

To support PFC across the VFL of Virtual Chassis, the links have to be configured in a certain manner. Only one VF Link is allowed from each port group based on the type of module being used to create the VFL. See Table 2 for the port group requirements:

Table 2: OS10K Module Port Groups

Module	Port Grouping
OS10K-XNI-U32E	Port Group 1 (Ports 1-16) Port Group 2 (Ports 17-32)
OS10K-XNI-U16E/U16L	Port Group 1 (Ports 1-16)
OS10K-QNI-U8E	Port Group 1 (Ports 1-4) Port Group 2 (Ports 5-8)
OS10K-QNI-U4E	Port Group 1 (Ports 1-4)
Note: The OS6900 platform and the OS10K-XNI-U32S module are not affected	

- If a chassis is operating in virtual chassis mode and the Data Center license is activated, then this port group requirement will be enforced. If there is no Data Center license installed this requirement is not needed.
- If the `vcsetup.cfg` file contains a configuration that does not adhere to this requirement, then on boot up only one VF-link from each group of ports listed above will be enabled. The remaining ports of the same port group will not be enabled. There is no guarantee as to which port will be selected to be enabled.
- During run time configuration if there is already a port from the port group that is a member of the VFL, then no additional ports from the port group can be added to the VFL and an error will be displayed if attempted.
- When configuring PFC over VFL on an OmniSwitch 6900, the VFL should not have more than eight ports comprising the VFL.

Using DCB Profiles

The DCB configuration for PFC, ETS, and DCBX is applied to switch ports through DCB profiles the same way that QSet profiles apply the QoS queue management configuration to switch ports. However, only DCB profiles or QSet profiles are allowed at any given time.

Note. QSet profiles and DCB profiles are mutually exclusive. If the OmniSwitch Data Center software license is installed, then DCB profiles are used. If this license is not installed, then QSet profiles are used.

- If the OmniSwitch Data Center software license is not installed, the switch boots up as a “non-DCB” switch and QSet profiles are applied to switch ports.
- If the OmniSwitch Data Center software license is installed, the switch boots up as a DCB switch and profiles are applied as follows:
 - If there is no existing DCB configuration, the default DCB profile (DCP 8) is applied to all switch ports. This occurs even if the port was previously assigned to a non-default QSP (for example, QSP 2, 3, or 4).
 - If there is an existing DCB configuration, the profiles are applied based on that configuration.
- If a switch boots up in the DCB mode and no DCB configuration is present (only default DCP 8), the switch will start DCBX by default to make sure each port is auto-configurable via DCBX to match the peer configuration. This provides a “plug-and-play” installation process that allows a switch running the default DCB configuration to automatically adapt to the network.

The DCB profiles are based on the 802.1Q-REV/D1-5 standard to define how the switch classifies different traffic types and priority mappings and then groups those types into traffic classes. Profiles also specify the Traffic Class Flow (TCF), which is LL (lossless; PFC initiated upstream) or nL (lossy; PFC not initiated upstream).

There are 10 pre-defined DCB profiles (DCP 1–10) available that represent common applications of the DCB standards (see [“Pre-Defined DCB Profiles \(Unicast\)” on page 2-15](#)). Creating custom profiles is also allowed; a maximum of 128 (including the 10 pre-defined) profiles per switch is supported.

Pre-Defined DCB Profiles (Unicast)

A DCB profile defines a set of PFC, ETS, and DCBx parameters that are applied to different traffic classes. Using profiles simplifies the DCB configuration process across the network. Once the appropriate pre-defined profile is selected or a custom profile is created for a particular set of traffic classes, the profile is then applied to each OmniSwitch that lies in the transport path for the specified set of traffic classes.

This section contains the Traffic Class priority mappings and the related DCB configuration parameters for each of the ten pre-defined OmniSwitch DCB profiles. By default, each QSet port instance is associated with DCP 8. See [“Multicast and Unicast Traffic Distribution” on page 2-24](#) for more information.

DCB Profile 1							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 3							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%)		TC Flow Control (TF)	TC Scheduler (TS)
0	BE	Best Effort	0, 1, 2, 3	33	100	Lossy	ETS
1	VO	Voice (<10ms latency)	5, 4	67	100	Loss Less	ETS
2	NC	Network Control	7, 6	0	100	Lossy	Strict Priority

DCB Profile 2							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 4							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%)		TC Flow Control (TF)	TC Scheduler (TS)
0	BE	Best Effort	0, 1	27	100	Lossy	ETS
1	CA	Critical Applications	3, 2	40	100	Loss Less	ETS
2	VO	Voice (<10ms latency)	5, 4	33	100	Lossy	ETS
3	NC	Network Control	7, 6	0	100	Lossy	Strict Priority

DCB Profile 3							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 5							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%) Min	TC Bandwidth (%) Max	TC Flow Control (TF)	TC Scheduler (TS)
0	BE	Best Effort	0, 1	27	100	Lossy	ETS
1	CA	Critical Applications	3, 2	40	100	Lossy	ETS
2	VO	Voice (<10ms latency)	5, 4	33	100	Loss Less	ETS
3	IC	Internetwork Control	6	0	100	Lossy	Strict Priority
4	NC	Network Control	7	0	100	Lossy	Strict Priority

DCB Profile 4							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 6							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%) Min	TC Bandwidth (%) Max	TC Flow Control (TF)	TC Scheduler (TS)
0	BK	Background	1	7	100	Lossy	ETS
1	BE	Best Effort	0	20	100	Lossy	ETS
2	CA	Critical Applications	3, 2	40	100	Lossy	ETS
3	VO	Voice (<10ms latency)	5, 4	33	100	Loss Less	ETS
4	IC	Internetwork Control	6	0	100	Lossy	Strict Priority
5	NC	Network Control	7	0	100	Lossy	Strict Priority

DCB Profile 5							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 7							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%) Min	TC Bandwidth (%) Max	TC Flow Control (TF)	TC Scheduler (TS)
0	BK	Background	1	7	100	Lossy	ETS
1	BE	Best Effort	0	20	100	Lossy	ETS
2	EE	Excellent Effort	2	25	100	Lossy	ETS
3	CA	Critical Applications	3	15	100	Lossy	ETS
4	VO	Voice (<10ms latency)	5, 4	33	100	Loss Less	ETS
5	IC	Internetwork Control	6	0	100	Lossy	Strict Priority
6	NC	Network Control	7	0	100	Lossy	Strict Priority

DCB Profile 6							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 8							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%) Min	TC Bandwidth (%) Max	TC Flow Control (TF)	TC Scheduler (TS)
0	BK	Background	1	7	100	Lossy	ETS
1	BE	Best Effort	0	20	100	Lossy	ETS
2	EE	Excellent Effort	2	15	100	Lossy	ETS
3	CA	Critical Applications	3	25	100	Lossy	ETS
4	VI	Video (<100ms latency)	4	15	100	Loss Less	ETS
5	VO	Voice (<10ms latency)	5	18	100	Loss Less	ETS
6	IC	Internetwork Control	6	0	100	Lossy	Strict Priority
7	NC	Network Control	7	0	100	Lossy	Strict Priority

DCB Profile 7							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 8							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%)		TC Flow Control (TF)	TC Scheduler (TS)
0	BE	Best Effort	0	0	100	Loss Less	Strict Priority
1	BK	Background	1	0	100	Loss Less	Strict Priority
2	EE	Excellent Effort	2	0	100	Loss Less	Strict Priority
3	CA	Critical Applications	3	0	100	Loss Less	Strict Priority
4	VI	Video (<100ms latency)	4	0	100	Loss Less	Strict Priority
5	VO	Voice (<10ms latency)	5	0	100	Loss Less	Strict Priority
6	IC	Internetwork Control	6	0	100	Loss Less	Strict Priority
7	NC	Network Control	7	0	100	Loss Less	Strict Priority

DCB Profile 8 (Default Profile)							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 8							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%)		TC Flow Control (TF)	TC Scheduler (TS)
0	BE	Best Effort	0	0	100	Lossy	Strict Priority
1	BK	Background	1	0	100	Lossy	Strict Priority
2	EE	Excellent Effort	2	0	100	Lossy	Strict Priority
3	CA	Critical Applications	3	0	100	Lossy	Strict Priority
4	VI	Video (<100ms latency)	4	0	100	Lossy	Strict Priority
5	VO	Voice (<10ms latency)	5	0	100	Lossy	Strict Priority
6	IC	Internetwork Control	6	0	100	Lossy	Strict Priority
7	NC	Network Control	7	0	100	Lossy	Strict Priority

DCB Profile 9							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 8							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%)		TC Flow Control (TF)	TC Scheduler (TS)
				Min	Max		
0	BK	Background	1	12	100	Loss Less	ETS
1	BE	Best Effort	0	12	100	Loss Less	ETS
2	EE	Excellent Effort	2	12	100	Loss Less	ETS
3	CA	Critical Applications	3	12	100	Loss Less	ETS
4	VI	Video (<100ms latency)	4	13	100	Loss Less	ETS
5	VO	Voice (<10ms latency)	5	13	100	Loss Less	ETS
6	IC	Internetwork Control	6	13	100	Loss Less	ETS
7	NC	Network Control	7	13	100	Loss Less	ETS

DCB Profile 10							
Profile Control :: DCBX Configuration							
PFC :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Capability : 8				Recommendation :: Enable TLV : [*Yes/No] Recommended Profile : [DCB_Def_Profile_1/NULL]			
ETS :: Enable TLV : [*Yes/No] Willing : [*Yes/No] Max TC : 3				Application :: Enable TLV : [Yes/*No] Defense Mode :: [*Enable/Disable] Default MTU :: [9216]			
Profile Data :: Max Traffic Classes (TC) : 8							
Traffic Class (TC)	Traffic Type (TT)	TC Description (TD)	TC Priority Map (TP)	TC Bandwidth (%)		TC Flow Control (TF)	TC Scheduler (TS)
				Min	Max		
0	BK	Background	1	12	100	Lossy	ETS
1	BE	Best Effort	0	12	100	Lossy	ETS
2	EE	Excellent Effort	2	12	100	Lossy	ETS
3	CA	Critical Applications	3	12	100	Lossy	ETS
4	VI	Video (<100ms latency)	4	13	100	Lossy	ETS
5	VO	Voice (<10ms latency)	5	13	100	Lossy	ETS
6	IC	Internetwork Control	6	13	100	Lossy	ETS
7	NC	Network Control	7	13	100	Lossy	ETS

Configuring DCB Profiles

The default DCB profile (DCP 8) is automatically assigned to each QSet instance when a port goes active or a port joins a LAG. It is only necessary to assign a different profile if DCP 8 attributes are not sufficient.

Consider the following when configuring DCB profiles:

- DCB profiles 1–10 are predefined profiles that are not modifiable and cannot be deleted from the switch configuration.
- Creating a custom profile is allowed by importing one of the pre-defined profiles into a new profile ID between 11 and 128, then modifying the new profile attributes as necessary.
- There is only one DCP assigned to each QSet instance and only one QSet instance per port or link aggregate (LAG). However, a LAG may show multiple QSet instances, one for each port that is a member of the LAG.
- When a port leaves a LAG, the default DCP 8 profile is associated with the QSet instance for that port. In other words, if the QSet instance for a port was associated with DCP 4 when the port joined the LAG, the port is associated with DCP 8 when it leaves the LAG.

Creating a Custom Profile

The `qos qsp dcb import` command is used to create a custom profile using one of the 10 pre-defined DCB profiles as a template. For example, the following command creates profile 20 (DCP 20) using pre-defined DCP 5 as the template for the new profile:

```
-> qos qsp dcb 20 import qsp dcb 5
```

Once the custom profile is created, the following traffic class attributes for the profile are configurable using the `qos qsp dcb tc` command:

TC Attribute	Command Parameters	Description
PFC flow control	pfc flow-type nll ll	Changes the traffic class to lossy or loss less.
PFC link delay	pfc link-delay 10-100	Sets the actual headroom value for the traffic class. Setting an incorrect value for this command may result in traffic loss. In most cases, the profile default value of 0 for lossy and 52 for lossless is sufficient.
Bandwidth	min-bw max-bw	Configures the minimum and maximum bandwidth for the traffic class.
Recommended bandwidth	recommended bw	Configures the recommended minimum bandwidth for the traffic class.

Guidelines for Modifying Custom Profiles

- 1 Changing the priority-to-TC mapping is not allowed.
- 2 Changing the TC scheduler type (SP, ETS) is not allowed.
- 3 Enabling or disabling PFC is allowed on any TC in the custom profile. The PFC status applies to all priorities within the specific TC.

- 4 Changing Weights (Min Rate) on any of the ETS TC is allowed, with ?ETS_TC_Weights < = 100%Available_PR (after SP allocation).
- 5 Changing PIR (shaper)/Max Rate is allowed on all TC Classes/Types (essential for high Priority SP (runaway) classes).
- 6 When a custom profile is modified, the changes are applied to all ports that are associated with that custom profile. To apply specific changes to a single port (QSet instance), import a custom or default profile into a new custom profile, make the necessary changes, then apply the new custom profile to the port.

Changing the Profile Assignment

To assign a different profile to a specific QSet instance (QSI), use the `qos qsi qsp dcb` command. For example, the following commands assign DCP 2 to port 1/2 and DCP 3 to ports 2/1-10 and linkagg 5:

```
-> qos qsi port 1/2 qsp dcb 2
-> qos qsi port 2/1-10 qsp dcb 3
-> qos qsi linkagg 5 qsp dcb 3
```

To view the DCB profile configuration for the switch, use the `show qos qsp dcb` command.

```
-> show qos qsp dcb
```

Legends: Prio TC Map:

Represents the priority to traffic class mapping;
begins with priority 0 on the left and displays the
traffic class it belongs to.

#	Name	ETS			802.3x		
		Priority TC Map	PFC Cap	Max TC	Template-DCP #	Name	Pause-Ready
1	dcp-1	00001122	8	8	1	dcp-1	No
2	dcp-2	00112233	8	8	2	dcp-2	No
3	dcp-3	00112234	8	8	3	dcp-3	No
4	dcp-4	10223345	8	8	4	dcp-4	No
5	dcp-5	10234456	8	8	5	dcp-5	No
6	dcp-6	10234567	8	8	6	dcp-6	No
7	dcp-7	01234567	8	8	7	dcp-7	No
8	dcp-8	01234567	8	8	8	dcp-8	No
9	dcp-9	10234567	8	8	9	dcp-9	No
10	dcp-10	10234567	8	8	10	dcp-10	No
20	dcp-20	10234567	8	8	10	dcp-10	No

See the *OmniSwitch CLI Reference Guide* for more information about related `show` commands.

Configuring Support for Legacy Pause Frames

Configuring support for legacy PAUSE frames on a DCB OmniSwitch is useful when the switch needs to interoperate with a server that does not support PFC. Support for legacy PAUSE frames is configurable by creating a custom profile with the 802.3x pause attribute enabled.

- 1 Create a custom profile based on any pre-defined profile and enable PAUSE. For example:

```
-> qos qsp dcb 12 import qsp dcb dcp-9 802.3x-pause
```

2 Disable DCBX and PFC TLVs on the port that will support the PAUSE frames. For example:

```
-> qos qsi port 1/1/5 dcb dcbx admin-state disable
-> qos qsi port 1/1/5 dcb dcbx pfc tlv disable
```

3 Enable pause on the port. For example:

```
-> interfaces 1/1/5 pause tx-and-rx
```

4 Apply the custom profile, “dcp-9” to the port. For example:

```
-> qos qsi port 1/1/5 qsp dcb 11
```

When this type of custom profile is applied on an ingress port, a legacy PAUSE frame is sent on the ingress port back to the server during congestion on the egress.

Configuring DCBx Port Parameters

By default, the Data Center Bridging Exchange (DCBX) protocol is enabled on all switch ports when the switch boots up with the OmniSwitch Data Center software license. To disable DCBX for a port or link aggregate, use the **qos qsi dcb dcbx admin-state** command. For example:

```
-> qos qsi port 1/10 dcb dcbx admin-state disable
-> qos qsi linkagg 2 dcb dcbx admin-state disable
```

In addition to configuring the status of DCBX on a port, the following DCBX port attributes are also configurable using the **qos qsi dcb dcbx pfc** and **qos qsi dcb dcbx ets** commands:

- The “willing” bit setting (Yes or No) for PFC and ETS. By default, this parameter is set to Yes, which means that when a profile configuration mismatch occurs between two directly connected devices, each device will negotiate common values for PFC settings. The ETS settings from the other device are accepted by each device.
- The status of configuration TLV transmission for PFC and ETS (enabled by default).
- The status of the recommended TLV for ETS (enabled by default).

To verify the DCBX port configuration, use the **show qos qsi dcbx** command. For example:

```
-> show qos qsi dcb dcbx
```

Port	DCP	Name	DCBX	Stats	PFC	ETS					
						Admin	Admin	Defense	PFC TLV	PFC Will	Cfg TLV
1/1	8	dcp-8	Ena	Dis	Ena	Ena	Yes	Ena	Ena	Yes	Dis
1/2	8	dcp-8	Ena	Dis	Ena	Ena	Yes	Ena	Ena	Yes	Dis
1/3	8	dcp-8	Ena	Dis	Ena	Ena	Yes	Ena	Ena	Yes	Dis
1/4	8	dcp-8	Ena	Dis	Ena	Ena	Yes	Ena	Ena	Yes	Dis
1/5	8	dcp-8	Ena	Dis	Ena	Ena	Yes	Ena	Ena	Yes	Dis
1/6	8	dcp-8	Ena	Dis	Ena	Ena	Yes	Ena	Ena	Yes	Dis
1/7	8	dcp-8	Ena	Dis	Ena	Ena	Yes	Ena	Ena	Yes	Dis
1/8	8	dcp-8	Ena	Dis	Ena	Ena	Yes	Ena	Ena	Yes	Dis
10	8	dcp-8	Ena	Dis	Ena	Ena	Yes	Ena	Ena	Yes	Dis
vfl	8	dcp-8	Ena	Dis	Ena	Ena	Yes	Ena	Ena	Yes	Dis

See the *OmniSwitch CLI Reference Guide* for more information about the **qos qsi dcb dcbx** and related **show** commands.

Multicast and Unicast Traffic Distribution

The following Class of Service (CoS) model for unicast and multicast traffic is applied when either the default QSet profile (QSP 1) or the default DCB profile (DCP 8) is the active profile for the port.

Cos 0 - Lower Priority MC (0-3) = 10

Cos 1 - Higher Priority MC (4-7) = 52

Cos 3 - All Other Unicast UC(0-7) =108

Cos 7 - CPU Generated Packets = 127 (maximum weight)

For example:

- When sending two streams of 100% MC Lower Priority and 100% MC Higher Priority, the distribution should be 10 and 50 packets, which is approximately 17% of Lower Priority MC and 83% of Higher Priority.
- When sending Lower Priority MC 100% and UC 100%, the distribution is 9% of MC and 91% of UC.
- When sending Higher Priority MC 100% and UC 100%, the distribution is 32% of MC and 68% of UC.

Non-Default Profile

The CoS model implemented also applies for non-default QSet profiles (QSP 2, 3, and 4), except on the OmniSwitch 6900 and the following OmniSwitch 10K modules:

- OS10K-QNI-U8 (8 x 40G)
- OS10K-QNI-U4 (4 x 40G)
- OS10K-XNI-U32E (32 x 10G)
- OS10K-XNI-U16E (16 x 10G)
- OS10K-XNI-U16L (8 x 10G, 8 x 1G)

However, for non-default QSet profiles (QSP 2–4) and non-default DCB profiles (DCP 1–7, 9–128) on the OmniSwitch 6900 and the OmniSwitch 10K modules listed above, the multicast and unicast queue mapping is as follows:

Strict Priority Profiles (for example, DCP 7)

Queues	Priority	Precedence
UC7	7	Highest
MC3	7, 6	
UC6	6	
UC5	5	
MC2	5, 4	
UC4	4	
UC3	3	
MC1	3, 2	
UC2	2	
UC1	1	
MC0	1, 0	
UC0	0	Lowest

Weighted Round Robin (WRR) Profiles

Queues	Priority	Weight
UC7	7	W7
MC3	7, 6	Avg(W7,W6)
UC6	6	W6
UC5	5	W5
MC2	5, 4	Avg(W5,W4)
UC4	4	W4
UC3	3	W3
MC1	3, 2	Avg(W3,W2)
UC2	2	W2
UC1	1	W1
MC0	1, 0	Avg(W1,W0)
UC0	0	W0

Note: W_n = Weight of UC_n

Avg(W_n,W_m) = Average of Weights of UC_n & UC_m

Profile with a Mix of Strict Priority and WRR

Unicast queues configured as Strict Priority will inherit behavior from the Strict Priority model, and unicast queues configured as WRR will inherit behavior from the WRR model. Multicast queues will always follow the behavior that the corresponding unicast queues are following. For example:

- If UC7 and UC6 are Strict Priority, then the MC3 (priority 6 and 7) will also use Strict Priority.
- If UC7 and UC6 are Weighted Round Robin, then MC3 (priority 6 and 7) will also use Weighted Round Robin. The weight of MC3 will be the average of the weights for UC6 and UC7.

For DCB profile ETS behavior, where a Traffic Class (TC) can have more than one priority, multicast queues will follow the corresponding unicast queue behavior. For example:

DCB Profile 1:

TC 0 (Priority 0 -3)

TC 1 (Priority 4 -5)

TC 2 (Priority 6 -7)

TC 0 has UC0 through UC3 in Round Robin, so MC0 (priority 0 and 1) and MC1 (priority 2 and 3) will also participate in the Round Robin behavior of TC 0.

Multicast Source PFC on OmniSwitch 6900

Ingress admission control on the OmniSwitch 6900 does not distinguish between unicast and multicast traffic. Therefore, a multicast source connected to a port which is PFC aware will react to congestion thereby pausing transmission. This will affect multicast hosts not in the congestion path.

When a multicast source is attached to a port on a OmniSwitch 6900, make sure that PFC is not enabled for that particular priority on the ingress. This can be done by configuring the port to use DCP 8 (all priorities are lossless) or for instance, DCB-1 (priority 4 and 5 are lossless, so multicast may be sent at any other priority other than priority 4 or 5).

If multicast sources are configured to react to PFC, it will affect subscribers not in the congestion path.

Verifying the DCB Configuration

Displaying the Data Center Bridging (DCB) configuration is helpful to verify the actual configuration on each switch in the data center mesh topology. To display information about the DCB configuration, use the **show** commands listed in this section.

show qos qsp dcb	Displays the configured DCB profiles and the traffic classes associated with the DCB profile.
show qos qsi dcbx	Displays the Data Center Bridging Exchange (DCBX) port configuration and status.
show qos qsi dcb ets	Displays the DCB Enhanced Transmission Selection (ETS) port configuration and status.
show qos qsi dcbx pfc	Displays the DCB Priority-based Flow Control (PFC) port configuration and status.
show qos pfc-lossless-usage	Displays the PFC lossless priority usage for the switch.
show qos qsi dcb pfc stats	Displays PFC port statistics.
show configuration snapshot vfc	Displays the current running configuration for DCB, QSP, PFC, and ETS.

3 Configuring Shortest Path Bridging

The Alcatel-Lucent OmniSwitch supports Shortest Path Bridging MAC (SPBM), as defined in the IEEE 802.1aq standard. SPBM uses the Provider Backbone Bridge (PBB) network model to encapsulate (using IEEE 802.1ah headers) and tunnel customer traffic through the network backbone. The shortest path trees upon which the PBB network infrastructure operates are determined using a version of the Intermediate System-to-Intermediate System (IS-IS) link state protocol that supports TLV extensions for SPB (ISIS-SPB).

Incorporating SPBM into the data center infrastructure provides the following benefits:

- Transparently extends Layer 2 connections (VLAN segments) across a large virtual service Layer 2 backbone network.
- Maintains a loop-free network while providing efficient use of available bandwidth, especially in a mesh topology. All connections between all switches in the topology remain active (no blocking of redundant links).
- A shortest path is automatically calculated between each bridge and every other bridge in the data center mesh, resulting in low latency and sub-second convergence times needed to support critical data center bridging requirements.
- Can process a large number of customer MAC addresses without overrunning provider network resources. Customer MAC addresses are only learned on Backbone Edge Bridges (BEB), where customer traffic is then encapsulated and tunneled through the network core infrastructure. Backbone Core Bridges (BCB) do not have to learn any customer MAC addresses.
- Provides a clear separation of customer traffic (between different customers and between the provider network domain). Entry points for customer traffic are clearly defined on the participating BEBs. Customer traffic is identified and associated with a specific service instance bound to the PBB infrastructure.
- Integration with Virtual Machine Network Profiles (vNPs) to support virtual machine (VM) discovery and mobility.

In This Chapter

This chapter provides an overview about how Shortest Path Bridging MAC (SPB-M) works and how to configure SPB-M through the Command Line Interface (CLI). CLI commands are used in the configuration examples; for more details about the syntax of commands, see the *OmniSwitch CLI Reference Guide*.

This chapter includes the following topics:

- “Shortest Path Bridging Specifications” on page 3-3.
- “SPBM Parameter Defaults” on page 3-4.
- “SPBM Interface Defaults” on page 3-4.
- “SPBM Service Defaults” on page 3-5.
- “Shortest Path Bridging Overview” on page 3-6.
- “Interaction With Other Features” on page 3-15.
- “Quick Steps for Configuring SPB” on page 3-18.
- “Configuring SPBM” on page 3-21.
- “Verifying the SPB Backbone and Services” on page 3-42.

Shortest Path Bridging Specifications

Platforms Supported	OmniSwitch 10K, 6900
OmniSwitch Software License	Advanced
IEEE Standards Supported	802.1aq/D3.6: Draft February 10, 2011— <i>Virtual Bridged Local Area Networks-Amendment 9: Shortest Path Bridging</i> 802.1ah/D4.2: DRAFT March 26, 2008— <i>Virtual Bridged Local Area Networks–Amendment 6: Provider Backbone Bridging</i>
IETF Internet-Drafts Supported	draft-ietf-isis-ieee-aq-05.txt— <i>ISIS Extensions Supporting IEEE 802.1aq Shortest Path Bridging</i>
SPB mode supported	SPB-M (MAC-in-MAC)
Maximum number of ISIS-SPB instances per switch.	1
Maximum number of BVLANs per switch	4
Number of equal cost tree (ECT) algorithms supported.	16
Maximum number of service instance identifiers (I-SIDs) per switch	1K
Maximum number of VLANs or SVLANs per I-SID	4K
Maximum Transmission Unit (MTU) size for SPB services.	9K (not configurable at this time)

SPBM Parameter Defaults

Parameter Description	Command	Default
ISIS-SPB status for the switch.	spb isis admin-state	Disabled
Equal Cost Tree (ECT) ID number for the backbone VLAN (BVLAN).	spb isis bvlan ect-id	1 or next available ECT ID number on the local switch.
Control BVLAN for the switch.	spb isis control-bvlan	None
The BVLAN tandem multicast mode (only applies to associated SPB services running in tandem mode).	spb isis bvlan tandem-multicast-mode	Source and Group (S, G)
Priority value for the ISIS-SPB instance.	spb isis bridge-priority	32768
Wait time intervals, in milliseconds, for shortest path first (SPF) calculations.	spb isis spf-wait	maximum wait: 1000 initial wait : 100 second wait : 300
Wait time intervals, in milliseconds, for link state PDU (LSP) transmissions.	spb isis lsp-wait	maximum wait: 1000 initial wait : 0 second wait : 300
Graceful restart status for the switch.	spb isis graceful-restart	Disabled
Graceful restart helper status for the switch.	spb isis graceful-restart helper	Disabled

SPBM Interface Defaults

Parameter Description	Command	Default
SPB interface status	spb isis interface	Disabled
SPB interface time interval between each hello packet transmission.	spb isis interface hello-interval	9 seconds
SPB interface hello multiplier used to determine hello packet hold time.	spb isis interface hello-multiplier	3
SPB interface link cost to reach the peer bridge.	spb isis interface metric	10

SPBM Service Defaults

By default, there are no SPBM service components configured for the switch. However, when a service is created, the following default values apply:

Parameter Description	Command	Default
SPB service administrative status.	service spb admin-state	Disabled
SPB service multicast replication mode.	service spb multicast-mode	Head-end
SPB service VLAN translation.	service spb vlan-xlation	Disabled
SPB service maximum transmission unit (MTU) value.	Not configurable at this time	9194
SPB service statistics collection.	service spb stats	Disabled
SPB service description.	service spb description	None.
Default profile automatically applied to access ports.	service access l2profile	def-access-profile
Layer 2 profile that specifies how control packets are processed on service access ports.	service l2profile	def-access-profile: STP, GVRP, MVRP = tunnel 802.3ad = peer 802.1x, 802.1ab, AMAP = drop CSCO PDU, VLAN, uplink = drop
VLAN translation for the service access port.	service access vlan-xlation	Disabled
Service access point (SAP) administrative status.	service spb sap admin-state	Disabled
SAP encapsulation.	service spb sap	0 (untagged traffic).
SAP trust mode.	service spb sap trusted	Trusted
SAP statistics collection.	service spb sap stats	Disabled
SAP description.	service spb sap description	None

Shortest Path Bridging Overview

The Alcatel-Lucent OmniSwitch implementation of Shortest Path Bridging (SPB) supports SPB MAC (SPBM) as defined in the IEEE 802.1aq standard. SPBM is defined for use in Provider Backbone Bridge (PBB) networks as specified in the IEEE 802.1ah standard.

SPBM provides a mechanism to automatically define a shortest path tree (SPT) bridging configuration through a Layer 2 Ethernet network. SPBM Ethernet services use this configuration to encapsulate and tunnel data through the PBB network. The following main components of the OmniSwitch implementation of SPBM provide this type of functionality:

- **ISIS-SPB**—A version of the Intermediate to Intermediate System (IS-IS) link state protocol that supports SPB TLV extensions. SPBM uses ISIS-SPB to build sets of symmetric shortest path trees (SPTs) between any SPB switch.
- **Provider Backbone Bridge (PBB) IEEE 802.1ah**— Defines a MAC-in-MAC data encapsulation path for PBB networks that is supported by SPBM.
- **Provider Backbone Bridge Network (PBBN)**—A network comprised of Backbone Edge Bridges (BEBs) and Backbone Core Bridges (BCB) that is used to interconnect Provider Bridge Networks (PBN) with other networks.
- **Backbone Edge Bridge (BEB)**—A SPB switch positioned at the edge of the PBB network that learns and encapsulates (adds an 802.1ah backbone header to) customer frames for transport across the backbone network. The BEB interconnects the customer network space with PBB network space.
- **Backbone Core Bridge (BCB)**—A SPB node that resides inside the PBB network core. The BCB employs the same BVLAN on two or more network ports. This BVLAN does not terminate on the switch itself; traffic ingressing on an SPB network port is switched to other SPB network ports. As a result, the BCB does not have to learn any of the customer MAC addresses. It mainly serves as a transit bridge for the PBB network.
- **SPBM Service**—An OmniSwitch Service Manager service configured on the BEBs. Each service maps to a service instance identifier (I-SID) which is bound to a backbone VLAN. One backbone VLAN can accommodate multiple I-SIDs.
- **Backbone VLAN (BVLAN)**—A VLAN that serves as a transport VLAN for the SPBM service instances and to connect SPB bridges together through SPT sets. Unlike standard VLANs, BVLANs do not learn source MAC addresses or flood unknown destination or multicast frames. Instead, BVLANs only forward on the basis of the forwarding database (FDB) as populated by the ISIS-SPB protocol.

The following illustration shows how SPBM uses the above components to tunnel customer traffic through a Provider Backbone Bridge Network:




Figure 1: SPBM Network Components

In this network,

- The BEBs are SPBM capable (ISIS-SPB configured and enabled) and form a shortest path bridging network that also includes the SPBM capable Backbone Core Bridges (BCBs).
- Each bridge calculates a shortest path tree (SPT) for each BVLAN with itself as the root of each tree.
- SPB Ethernet service instances identified by I-SIDs are created on each BEB. Each I-SID is associated with a BVLAN ID. The BVLAN is configured on each bridge (BEB and BCB) in the backbone network. However, the I-SID itself and the I-SID association with the BVLAN is only configured on each BEB that will service customer traffic.
- A Service Access Point (SAP) is configured on each BEB to identify the access port on which customer traffic will enter the PBBN, the SPB service instance that will tunnel the traffic through the network, and the type of customer traffic to forward (for example, only specific CVLAN IDs, untagged traffic only, or all tagged traffic). Basically, the SAP binds access ports and the specified customer traffic received on those ports to the service.
- Layer 2 traffic from the connected edge networks enters the BEBs through access ports. The SAP configuration on the receiving access port is applied to classify which frames are mapped to which services, if any.
- Classified traffic is then encapsulated into 802.1ah frames by the BEB before the frames are transmitted through the backbone network.
- The 802.1ah encapsulated frames are forwarded on the shortest path through the entire PBBN to reach the intended destination BEB. The BCBs switch traffic based on the destination backbone MAC address (BMAC)—bridge MAC address of the BEB—provided in the 802.1ah header and do not process any I-SID information in the frame.

SPBM Shortest Path Trees

The shortest path between two points is a straight line. Shortest Path Bridging (SPB) implements frame forwarding on the shortest path between any two bridges in an Ethernet network. The shortest path trees (SPTs) calculated by SPB provide the shortest and most efficient path to and from the intended destination. SPTs are formed along the direct, straight-line links between switches to make up an overall path through the topology that provides a robust, efficient direction for network traffic to travel.

The SPBM network topology consists of two layers:

- **The backbone infrastructure (control plane) layer.** ISIS-SPB builds the backbone layer by defining loop-free, shortest path trees (SPTs) through the backbone network.
- **The services (data plane) layer.** The service layer is based on the Provider Backbone Bridging (PBB) framework as defined in the IEEE 802.1ah standard. SPBM supports the 802.1ah MAC-in-MAC method for data encapsulation. SPBM services transport the encapsulated traffic over the ISIS-SPB infrastructure. (See “[SPB Services](#)” on page 3-12 for more information).

This section contains an example of ISIS-SPB operations in a small SPBM network. In addition to describing how shortest path trees are created in the BVLAN domain, the flow of unicast and multicast traffic through the network, this example also shows the benefits of using SPB over Spanning Tree for VLAN traffic distribution.

Spanning Tree

The following diagram shows an example Provider Backbone Bridge (PBB) network with a single backbone VLAN using the Spanning Tree protocol for network loop protection with the same path cost on all links:

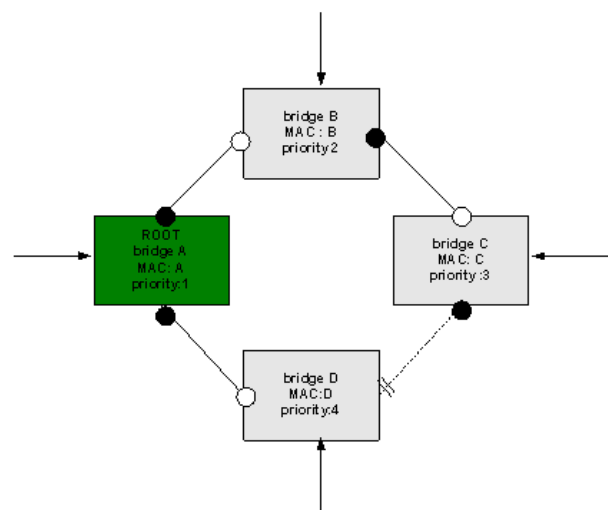


Figure 2: Spanning Tree Topology

In this example, Bridge A is the Root bridge. As a result, customer traffic entering Bridge A would always use the shortest path to reach every other bridge in the network. However, traffic entering Bridge D that is destined for Bridge C must traverse the path through Bridge A to reach Bridge C, even though Bridge D is directly connected to Bridge C. Clearly the path from Bridge D to Bridge C is not the shortest path in this case.

ISIS-SPB

The IEEE 802.1aq standard for SPB specifies the use of the IS-IS link state protocol instead of Spanning Tree to form sets of shortest path trees through the network. When SPB is used, each bridge is the Root for all traffic entering that bridge. As a result, each bridge can provide the shortest path to every other bridge in the network.

The bridging methodology needed to allow each bridge to serve as its own root bridge is enforced through the use of SPB BVLANS. This type of VLAN does not learn customer MAC addresses or flood unknown unicast and multicast traffic. In addition, network loops are mitigated through strict ingress checks based on the source MAC address of frames received on the BVLAN (frames received from an unexpected source are discarded).

SPBM uses an extended version of the IS-IS protocol that supports SPB (ISIS-SPB) to calculate the SPBM network topology. In addition, the learning and propagation of source MAC addresses is handled through the ISIS-SPB control plane, instead of through the data plane.

When calculating the SPBM network topology, ISIS-SPB must meet Layer 2 requirements to create congruent and symmetric paths. To do this, SPBM supports 16 predefined Equal Cost Tree (ECT) algorithms to break ties when two or more equal cost paths to the same destination are discovered. The same ECT algorithm is configured for the same BVLAN ID on each SPB switch in the network to ensure congruent, symmetric paths for the service traffic bound to that BVLAN.

Basically, to create a unicast tree, SPBM simply computes the shortest path from every bridge with each bridge serving as the Root (as shown below) and populates the Layer 2 forwarding database (FDB) on the SPB bridges with MAC addresses.

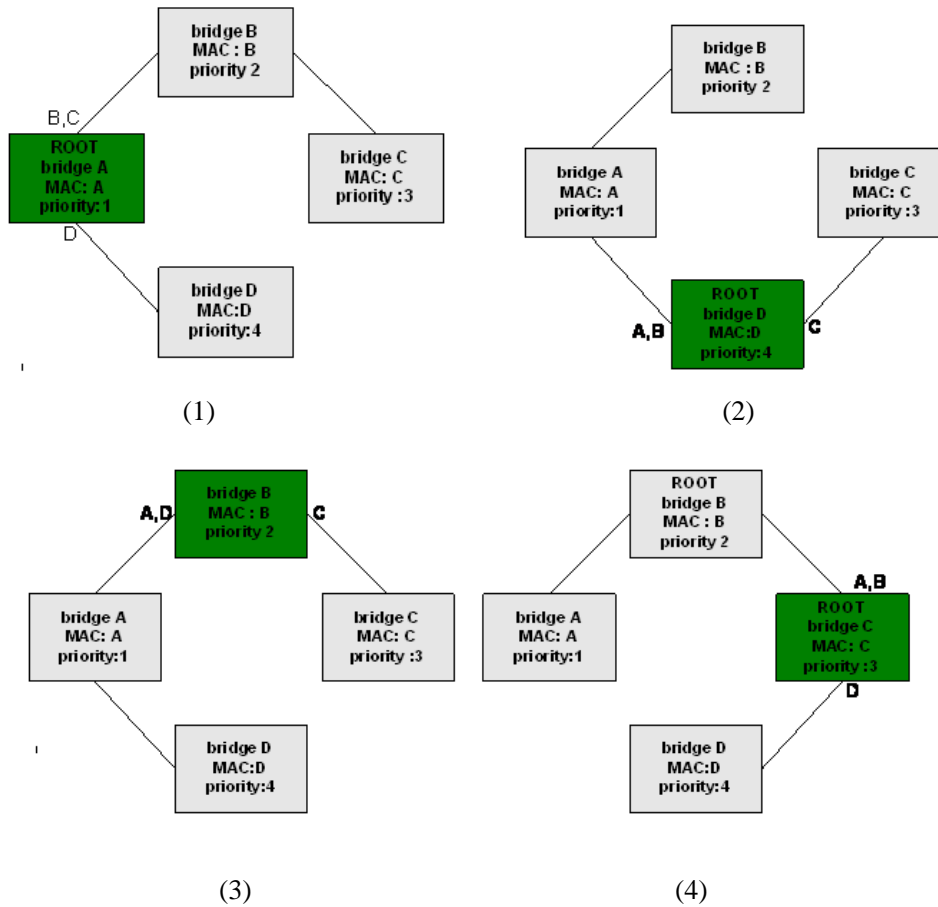


Figure 3: ISIS-SPB Shortest Path Calculations

The ISIS-SPB unicast trees shown in Figure 3 were built as follows:

- 1** Bridge A calculates the shortest path tree to Bridge B and then programs its FDB with MAC address B on the link, as shown in 4(1).
- 2** Bridge A will then calculate shortest paths to Bridge C and Bridge D and programs the MAC addresses according to the path computed.
- 3** All other bridges follow the same procedure (note that the actual computation is much more optimized and the description here is only for illustration purposes).
- 4** The following traffic pattern for this example network is the result of the ISIS-SPB SPT calculations:




Figure 4: ISIS-SPB Topology

As shown in Figure 4, all the backbone MAC (BMAC) addresses are learned by the switches when ISIS-SPB converges. The path taken by each unicast flow (for example ABC, CBA) are reverse path congruent and travel the shortest path through the network.

In the ISIS-SPB topology (Figure 4), the link between Bridge D and Bridge C carries traffic, whereas in the Spanning Tree topology (Figure 3), this link is blocked. Although these examples are based on traffic distribution for a single BVLAN, the ability to make all links in the topology available at all times is especially advantageous in highly redundant, meshed networks.

Although the link between Bridge D and Bridge C is used in the ISIS-SPB topology, traffic flow is relatively low in comparison to the other links. To make better use of this link, a second BVLAN could be created and assigned a different ECT algorithm to trigger ISIS-SPB calculations of a separate set of SPTs for the second BVLAN. This is similar to creating a new Multiple Spanning Tree (MST) instance in a Spanning Tree topology to create a different tree and assigning a new VLAN to that instance.

Each ECT algorithm uses a different calculation to break ties when paths between SPB bridges are equal cost. Another method to influence the SPT calculation is to modify the bridge priority for the switch or change the link cost metric for the SPB interface connection between two switches.

Multicast Traffic

SPBM supports two methods for replicating and forwarding multicast traffic (or unknown destination traffic) received from customer equipment: head-end replication and tandem replication.

- **Head-end replication.** Multicast traffic is replicated once for each receiver, encapsulated with the BMAC address, and then sent as a unicast packet to each destination. This method is more suited for networks where there is a low demand for multicast traffic.
- **Tandem replication.** Multicast traffic is replicated only where there is a fork in the SPT and each branch has at least one receiver. Each multicast source bridge in the SPBM network is the root for a multicast distribution tree (MDT). A MDT is created per-source per-BVLAN and it is pruned according to whether the SPB node is on the shortest path of a multicast transmitter and receiver. For those

MDTs that cross a given Backbone Edge Bridge (BCB), that BCB needs to generate a multicast forwarding table for each such MDT.

Multicast traffic originating from a bridge is encapsulated with a special multicast BMAC DA that identifies the source of the traffic and then forwarded on the tree. Participating bridges that receive the packet will then know the source of the traffic and will use the multicast forwarding information for that source to switch the packet to the appropriate destination.

SPB Services

The SPBM network topology consists of two layers:

- **The backbone infrastructure (control plane) layer.** ISIS-SPB builds the backbone layer by defining loop-free, shortest path trees (SPTs) through the backbone network (see [“SPBM Shortest Path Trees” on page 3-8](#) for more information).
- **The services (data plane) layer.** The service layer is based on the Provider Backbone Bridging (PBB) framework as defined in the IEEE 802.1ah standard. SPBM supports the 802.1ah MAC-in-MAC method for data encapsulation. SPBM services transport the encapsulated traffic over the ISIS-SPB infrastructure.

The SPB service layer framework is comprised of the following components:

- **Backbone Edge Bridge (BEB).** An OmniSwitch is considered a BEB if the switch is SPB capable and at least one service access point (SAP) and one SPB interface is configured on the switch. The BEB marks the boundary between the customer network and the PBB network (PBBN).
- **Backbone Core Bridge (BCB).** An OmniSwitch is considered a BCB if the switch is SPB capable and no SAPs are configured but at least one SPB interface is configured on the switch to forward encapsulated SPBM network traffic. Note that the requirement for configuring a BCB is based on whether or not the network topology includes a transit bridge.
- **Service Instance Identifier (I-SID).** Configured only on a BEB, this component identifies a backbone service instance that will tunnel the encapsulated data traffic through the PBBN between BEBs. The I-SID is bound to a BVLAN ID and a Service Manager SPB service ID when the service is created.
- **Access Port.** A port or link aggregate configured as an SPB access port. This type of port is configured on the BEBs and defines the point at which traffic from other provider networks or directly from customer networks enters the PBBN. The access port is also associated with a Layer 2 profile that specifies how to process protocol control frames received on the port
- **Service Access Point (SAP)**—A SAP is a logical service entity (also referred to as a virtual port) that is configured on a BEB to bind an access port to an SPB service ID and specify the type of customer traffic ((untagged, single-tagged, double-tagged, or all) to encapsulate and tunnel through the PBBN.
- **SPB Interface (Network Port)**—A port or link aggregate configured as an SPB interface that resides on either a BEB or a BCB and connects to the backbone network. Network ports carry customer traffic encapsulated in 802.1ah frames and are associated with all BVLANs on the switch. Customer traffic ingressing on a network port is switched to another network port (on BCBs) or to an access port (on BEBs).

Once the ISIS-SPB infrastructure and the SPB service-based architecture is defined, the following service components are dynamically created by the OmniSwitch. No user-configuration is required.

- **Service Distribution Point (SDP)**—A SDP provides a logical point at which customer traffic is directed from one BEB to another BEB. SDPs are used to set up distributed services, which consist of

at least one SAP on a local node, one SAP on a remote node, and an SDP binding the service on both nodes.

- **Mesh SDP**—A mesh SDP represents the binding of a SPB service instance to an SDP. The SDP then distributes the service connectivity to other BEBs through the ISIS-SPB shortest path trees.

Sample SPBM Network Topology

The following diagram provides a sample SPBM network topology that shows how the SPBM service and ISIS-SPB backbone layers work together to basically extend (or virtualize) customer traffic across a Provider Backbone Bridge Network (PBBN):

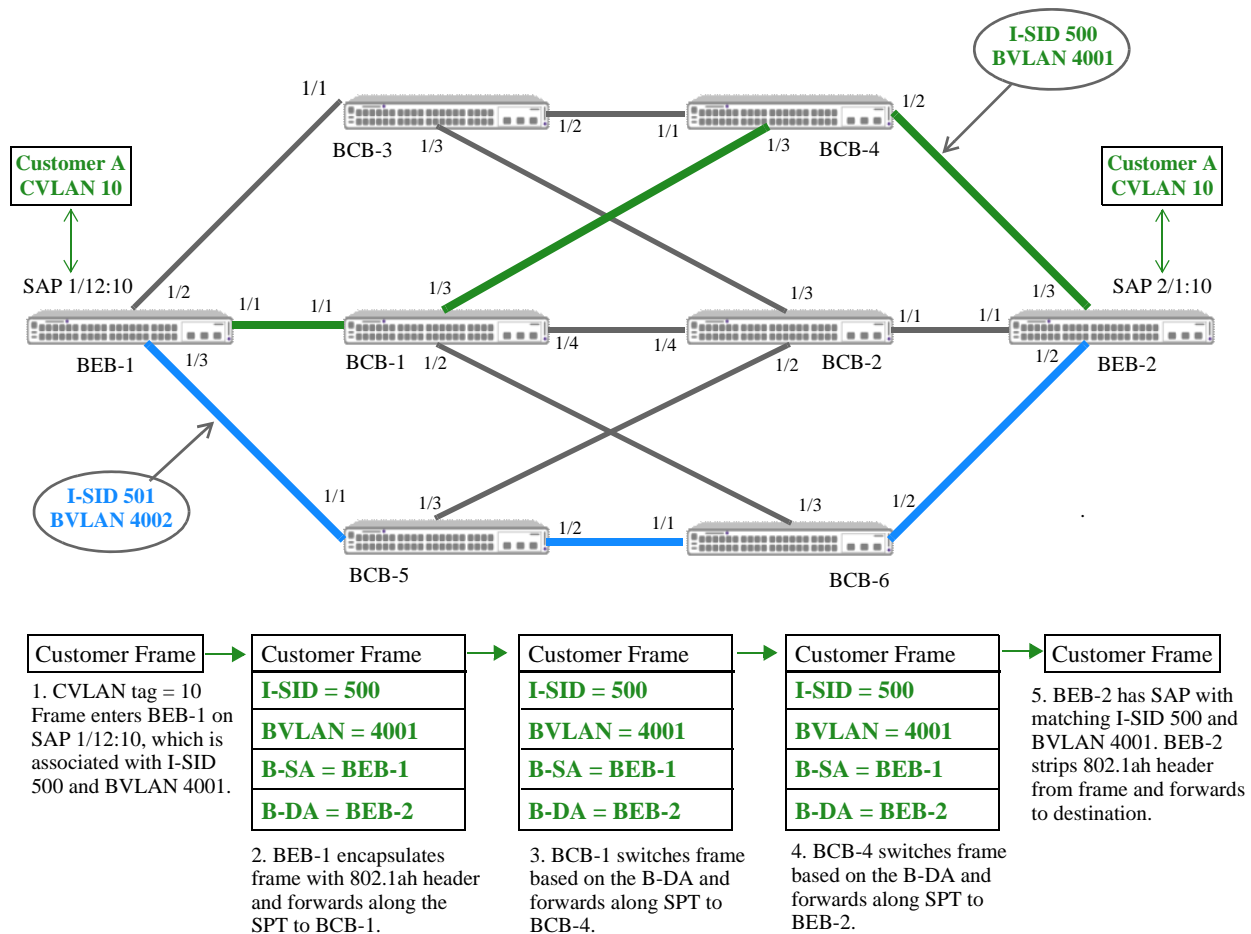


Figure 5: Sample SPBM Network

In this sample SPBM topology:

- The packet flow for Customer A frames tagged with VLAN 10 is shown as a typical example. These frames are mapped to an SPB service that represents a binding of I-SID 500 to BVLAN 4001. The path for this binding is shown in green.
- An additional path, shown in blue, is for another SPB service that represents a binding of I-SID 501 to BVLAN 4002. This provides an example of how adding an additional BVLAN and service configuration to the topology can provide an alternate service path for other traffic from the same customer or traffic from a completely different customer.

- SPB BVLAN 4001 and 4002 are created and assigned to ECT ID 1 and 2, respectively, on every switch (BEBs and BCBs) in the topology. These BVLANs serve as the transport entity on which ISIS-SPB builds the shortest path trees and SPB services tunnel data.
- The switch ports connecting each SPB switch with the next-hop SPB switch are configured as SPB interface ports. This type of port is used to forward ISIS-SPB control packets and serves as a network port for tunneling encapsulated traffic through SPB services.
- The service access points (SAPs) created on BEB-1 and BEB-2 determine which frames from Customer A are accepted on the SAP port, where they are then encapsulated and mapped to the associated service. Other SAPs exist on these switches for the other service path.
- When a frame tagged with VLAN 10 ingresses on port 1/12, the frame is encapsulated in an 802.1ah header. The header specifies the BMAC for BEB-1 as the B-SA, the BMAC for BEB-2 as the B-DA, the SAP I-SID (500), and the SAP BVLAN (4001).
- All other frames ingressing on SAP 1/12:10 that are not tagged with VLAN 10 are dropped, unless there are other SAPs configured for that port that will classify those frames.
- The encapsulated frame is then forwarded along the BVLAN 4001 shortest path tree (SPT) to BEB-2, where the 802.1ah header is stripped off and the frame is forwarded to the appropriate destination port.
- The entire process for encapsulating and tunneling customer frames is the same for frames ingressing on port 2/1 of BEB-2 destined for BEB-1.

How it Works

- There is one instance of ISIS-SPB supported in the backbone topology. This instance is activated once the BVLANs and SPB interfaces are created and the administrative status of ISIS-SPB is enabled for each switch.
- When ISIS-SPB is administratively enabled on each switch, all the configured SPB interfaces start to advertise Hello packets to discover and establish adjacencies with other SPB switches.
- Once adjacencies are established, link state packets (LSPs) are generated with SPB-specific TLVs and shortest path trees from each switch to all other switches are calculated.
- Each SPB switch learns the backbone MAC (BMAC) address and associated BVLAN IDs of every SPB switch in the network and stores that information in a local forwarding database. The BMAC address is the bridge MAC address of the switch and is advertised by ISIS-SPB as the System ID.
- ISIS-SPB then informs Service Manager of the reachability of the BMAC/BVLAN combinations. This information is used to automatically create a service distribution point (SDP) between the same BVLAN on each BEB.
- When ISIS-SPB receives advertisement of a service instance identifier (I-SID) from a remote BEB that matches an I-SID created on the local switch, the SDP (BMAC/BVLAN) of the remote BEB is bound to the I-SID. The binding of a service to an SDP is referred to as a mesh SDP.
- Basically, a SDP is a dynamically created logical entity that distributes service connectivity to other BEBs through the ISIS-SPB shortest path trees. When customer frames are then classified into a specific SAP, the frames are encapsulated and tunneled through the mesh SDP (service/SDP bind) associated with that SAP.

Interaction With Other Features

This section contains important information about Shortest Path Bridging MAC (SPBM) interaction with other OmniSwitch features. Refer to the specific chapter for each feature to get more detailed information about how to configure and use the feature.

Backbone VLANs (VLAN Manager)

VLAN Manager CLI commands are used to create a SPB backbone VLAN (BVLAN). Although a BVLAN is created in a similar manner as a standard VLAN, BVLANs differ from standard VLANs as follows:

- No Spanning Tree control—the Spanning Tree protocol is automatically disabled on each BVLAN and all ports associated with each BVLAN will remain in a forwarding state. However, Spanning Tree can remain operational on other types of VLANs.
- No source MAC address learning—normal hardware learning is disabled on BVLANs. Instead, the forwarding database (FDB) is populated by the ISIS-SPB protocol.
- There is no flooding of unknown destination or multicast frames.
- Ingress filtering based on the source MAC address—frames received on ports that do not have an incoming source MAC address pre-programmed by ISIS-SPB are discarded.
- IP interfaces are not supported on BVLANs.

Dynamic Service Access Points (UNP)

The Universal Network Profile (UNP) feature supports two types of profiles: VLAN and service. The service profile type triggers the dynamic creation of a service access point (SAP) when traffic received on a UNP access port is classified and assigned to that profile. If the service that the SAP is associated with does not exist, the service is also dynamically created.

The main advantage of UNP service profiles is to allow incoming traffic to trigger dynamic SAP creation, thus reducing the amount of manual configuration required. In addition, no other protocols are required on the switch or host device to support this functionality.

Hardware

All SPB functionality is supported on the OmniSwitch 6900 and OmniSwitch 10K switches. However there is one exception: the OS10K_XNI_U32 module does not support the configuration of access ports and service access point (SAPs). As a result, the level of SPB support on an OmniSwitch 10K with an OS10K_XNI_U32 module is as follows:

- The switch can serve as a Backbone Core Bridge (BCB) if ports on the OS10K_XNI_U32 are only configured as SPB interfaces (network ports).
- The switch can serve as a Backbone Edge Bridge (BEB) in a mixed chassis only if access ports and SAPs are configured on other module types and no SPB at all is configured on the OS10K_XNI_U32.

Link Aggregation

- Both static and dynamic link aggregates are configurable as SPBM service access ports and as SPBM network interfaces.
- Note that a link aggregate must consist of all access ports or all network ports. SPBM functionality is not supported on link aggregates that consist of a mixture of SPBM ports and standard switch ports.

OAM

- OAM support per the IEEE 802.1ah standard for Provider Backbone Bridging (PBB) is applied at the customer VLAN (CVLAN) and backbone VLAN (BVLAN) level. Support at the service instance (I-SID) level is not available.
- The OmniSwitch proprietary Layer 2 ping and traceroute features are available to troubleshoot CVLAN and BVLAN domains, including an I-SID check.

Quality of Service (QoS)

- The priority assignment of a user frame is determined at an access point. A Service Access Point (SAP) on an SPB access port can be configured to be trusted or un-trusted. If a SAP is configured to be trusted, then internal priority for ingress traffic on that SAP is derived from tagged or NULL tagged ingress packet priority or from default port priority if ingress packet is untagged. If a SAP is untrusted then internal priority can be user configured.
- QoS performs the following actions on ports configured as access ports:
 - > Access ports are automatically trusted and the default classification is set to 802.1p.
 - > The trust status and classification are not user-configurable on access ports.
 - > All QoS CLI configuration is blocked on access ports. This includes physical ports and ports that are members of a link aggregate.
 - > Untagged L2 control packets (such as BPDU, GVRP, AMAP) are always tunneled (if enabled) through the SPB domain with the default EXP bits set at 7, so that they can arrive at the destination at the highest priority of 7. Trusted/untrusted SAPs configured on access ports will not affect the priority assignment for Layer 2 control packets.
- QoS priority (802.1p) is applied as follows to trusted and untrusted SAPs:

SAP Configuration	Allowed Configuration	
Tagged (VLAN 1–4094)	Trusted	Tagged traffic priority derived from tags.
	Untrusted	Tagged traffic priority configured by user.
QinQ (outer VLAN 1–4094)	Trusted	Tagged traffic priority derived from outer tags.
	Untrusted	Tagged traffic priority configured by user.
Wild Card	Trusted	Tagged traffic priority derived from tags. Untagged traffic Port default (PRI 0).
	Untrusted	Tagged/ traffic priority configured by user
Untagged	Trusted	Untagged Traffic Port default (PRI 0)
	Untrusted	Priority configured by user.

- > By default, a SAP is trusted with best effort priority (0)

- A SAP can be dynamically changed to trusted/untrusted without admin down the sap
- A SAP priority may only be set when a SAP is untrusted.
- When a SAP is changed from untrusted to trusted, any previously assigned priority is reset with best effort priority (0).
- A trusted SAP that defines a double-tagged encapsulation (QinQ) will use the outer VLAN tag to determine the priority of the frame.
- Priority handling at the edge and core components of an SPBM topology:
 - On a ingress Backbone Edge Bridge (BEB), a frame is classified to a SAP. The internal priority is determined based on the QoS settings of the SAP (for example, trusted vs. untrusted, default priority). This internal priority is mapped to the backbone VLAN (BVLAN) tag of the tunnel encapsulation.
 - The Backbone Core Bridge (BCB) acts as a Layer 2 device that switches the frame across ingress to egress ports in the BVLAN domain. The BVLAN tag is used to determine the internal priority queue on the egress port where the frame is enqueued.
 - On a egress BEB, the internal priority is determined from the BVLAN tag. The frame is decapsulated and enqueued to the egress queue(s) of the access port(s) based on this internal priority.

Quick Steps for Configuring SPB

This section provides a quick tutorial for configuring the SPBM network backbone (control plane) and the service encapsulation path (data plane). The Command Line Interface (CLI) commands provided in this section are used to configure the “[Sample SPBM Network Topology](#)” on page 3-13.

Quick Steps for Configuring the SPBM Backbone

The following quick steps are used on each switch in the SPBM backbone that will participate in the “[Sample SPBM Network Topology](#)” on page 3-13. This includes both edge and transit (core) switches.

- 1 Use the **system name** command to assign a unique system name to each SPB switch in the domain.

```
-> system name BEB-1
-> system name BEB-2
-> system name BCB-1
-> system name BCB-2
-> system name BCB-3
-> system name BCB-4
-> system name BCB-5
-> system name BCB-6
```

- 2 Use the **spb bvlan** command to create BVLANs 4001 and 4002 on each switch (edge and core switches) that will participate in the SPBM topology.

```
-> spb bvlan 4001
-> spb bvlan 4002
```

- 3 Use the **spb isis bvlan ect-id** command to change the equal cost tree (ECT) algorithm ID for the specified BVLAN, as necessary, to make sure that the same ECT ID is assigned to the same BVLAN ID on each switch in the SPBM topology.

```
-> spb isis bvlan 4001 ect-id 1
-> spb isis bvlan 4002 ect-id 2
```

- 4 Use the **spb isis control-bvlan** command to designate one of the BVLANs on each SPB switch as the control BVLAN for the SPB instance. The control BVLAN is used to exchange ISIS-SPB control packets with neighboring SPB switches.

```
-> spb isis control-bvlan 4001
```

- 5 Use the **spb isis interface** command to configure a port or link aggregate as a SPB interface. This type of interface sends PDUs to detect neighboring SPB switches and form adjacencies and also serves as a network port that is used to carry encapsulated service traffic through the SPBM backbone network.

```
-> spb isis interface port 1/1-3
-> spb isis interface port 1/1-4
```

- 6 Use the **spb isis admin-state** command to enable the SPB instance for the switch. Enabling ISIS-SPB on the switch triggers the transmission of hello packets from the SPB interfaces, which starts the process of defining the SPB infrastructure and calculating the shortest path trees (SPTs) through the topology.

```
-> spb isis admin-state enable
```

Quick Steps for Configuring SPB Services

The following quick steps use the OmniSwitch Service Manager commands to configure the logical entities that comprise the SPB services in the [“Sample SPBM Network Topology”](#) on page 3-13.

1 Use the **service access** command to configure a port or link aggregate on which customer traffic is received as a SPB service access port.

```
-> service access port 1/1
-> service access port 2/1
```

2 Use the **service spb** command to create a SPB service and associate that service with a backbone service instance identifier (I-SID) and BVLAN.

```
-> service spb 1 isid 500 bvlan 4001 admin-state enable
-> service spb 2 isid 501 bvlan 4002 admin-state enable
```

3 Use the **service spb sap** command to create a service access point (SAP) by associating a SPB service with SAP ID. A SAP ID is comprised of a port or link aggregate and an encapsulation value that identifies the customer traffic to associate with the service.

```
-> service spb 1 sap port 1/12:10 admin-state enable
-> service spb 1 sap port 2/1:10 admin-state enable
-> service spb 2 sap port 1/12:0 admin-state enable
-> service spb 2 sap port 1/12:all admin-state enable
```

In this example, SAPs 1/12:10 and 2/1:10 map traffic ingressing on these SAPs that has an outer tag (customer VLAN tag) equal to 10 to the service associated with the SAP. In this case, SPB service 1 (ISID=500, BVLAN=4001). SAPs 1/12:all and 2/1:all map all tagged and untagged traffic ingressing on these SAPs to the service associated with the SAP.

Sample Command Configuration

This section provides the sequence of commands used on each switch to configure the [“Sample SPBM Network Topology”](#) on page 3-13. Note that the SPBM backbone is configured on every switch first, then the SPBM service architecture is configured second. Following this order of configuration is highly recommended to ensure proper switch participation in ISIS-SPB adjacencies and shortest path tree calculations.

SPBM Backbone Commands

The **system name** and Shortest Path Bridging (**spb**) commands are used to configure the SPBM backbone infrastructure for the sample topology, as shown:

BEB-1	BEB-2	BCB-1
-> system name BEB-1	-> system name BEB-2	-> system name BCB-1
-> spb bvlan 4001	-> spb bvlan 4001	-> spb bvlan 4001
-> spb bvlan 4002	-> spb bvlan 4002	-> spb bvlan 4002
-> spb isis bvlan 4001 ect-id 1	-> spb isis bvlan 4001 ect-id 1	-> spb isis bvlan 4001 ect-id 1
-> spb isis bvlan 4002 ect-id 2	-> spb isis bvlan 4002 ect-id 2	-> spb isis bvlan 4002 ect-id 2
-> spb isis control-bvlan 4001	-> spb isis control-bvlan 4001	-> spb isis control-bvlan 4001
-> spb interface port 1/1-3	-> spb interface port 1/1-3	-> spb interface port 1/1-4
-> spb isis admin-state enable	-> spb isis admin-state enable	-> spb isis admin-state enable

BCB-2

```
-> system name BCB-2
-> spb bvlan 4001
-> spb bvlan 4002
-> spb isis bvlan 4001 ect-id 1
-> spb isis bvlan 4002 ect-id 2
-> spb isis control-bvlan 4001
-> spb interface port 1/1-4
-> spb isis admin-state enable
```

BCB-3

```
-> system name BCB-3
-> spb bvlan 4001
-> spb bvlan 4002
-> spb isis bvlan 4001 ect-id 1
-> spb isis bvlan 4002 ect-id 2
-> spb isis control-bvlan 4001
-> spb interface port 1/1-3
-> spb isis admin-state enable
```

BCB-4

```
-> system name BCB-4
-> spb bvlan 4001
-> spb bvlan 4002
-> spb isis bvlan 4001 ect-id 1
-> spb isis bvlan 4002 ect-id 2
-> spb isis control-bvlan 4001
-> spb interface port 1/1-3
-> spb isis admin-state enable
```

BCB-5

```
-> system name BCB-5
-> spb bvlan 4001
-> spb bvlan 4002
-> spb isis bvlan 4001 ect-id 1
-> spb isis bvlan 4002 ect-id 2
-> spb isis control-bvlan 4001
-> spb interface port 1/1-3
-> spb isis admin-state enable
```

BCB-6

```
-> system name BCB-6
-> spb bvlan 4001
-> spb bvlan 4002
-> spb isis bvlan 4001 ect-id 1
-> spb isis bvlan 4002 ect-id 2
-> spb isis control-bvlan 4001
-> spb interface port 1/1-3
-> spb isis admin-state enable
```

SPBM Service Commands

The Service Manager (**service**) commands are used to build the SPBM services architecture for the sample topology, as shown. Note that services are only configured on designated BEB switches.

BEB-1

```
-> service access port 1/12
-> service spb 1 isid 500 bvlan 4001
-> service spb 2 isid 501 bvlan 4002
-> service spb 1 sap port 1/12:10 admin-state enable
-> service spb 2 sap port 1/12:0 admin-state enable
-> service spb 2 sap port 1/12:all admin-state enable
```

BEB-2

```
-> service access port 1/12
-> service spb 1 isid 500 bvlan 4001
-> service spb 2 isid 501 bvlan 4002
-> service spb 1 sap 1/12:10 admin-state enable
-> service spb 2 sap 1/12:0 admin-state enable
-> service spb 2 sap 1/12:all admin-state enable
```

Configuring SPBM

Configuring the SPBM backbone and service layers requires several steps. These steps are outlined here and further described throughout this section. For a brief tutorial on configuring SPBM, see [“Quick Steps for Configuring SPB” on page 3-18](#).

Configure the SPBM Backbone (ISIS-SPB)

Only switches that are SPB capable can participate in the SPBM network topology. The following configuration steps are required to make an OmniSwitch a SPB-capable node:

- 1 Create a BVLAN.** The BVLAN provides the foundation of the SPBM infrastructure. A BVLAN is associated with an equal cost tree (ECT) algorithm ID and a SPB service instance ID that is used to carry customer traffic through the backbone network. See [“Backbone VLANs” on page 3-22](#).
- 2 Configure SPB interfaces.** A SPB interface is associated with each BVLAN that is configured on the switch. At the ISIS-SPB level, this type of interface sends and received ISIS Hello packets and link state PDU (LSP) to discover adjacent SPB switches and calculate the shortest path trees through the SPBM network topology. At the services level, the SPB interfaces serve as network ports that are used to carry encapsulated customer traffic through the network. See [“Configuring SPB Interfaces” on page 3-25](#).
- 3 Configure global ISIS-SPB parameters.** In addition to enabling/disabling the ISIS-SPB instance for the switch, global configuration includes settings such as a system name for the switch, global bridge parameters, and various wait time intervals. When ISIS-SPB is enabled for the switch, default settings for these global bridge parameters and wait time intervals are active. It is only necessary to change these values if the default settings are not sufficient. See [“Configuring Global ISIS-SPB Parameters” on page 3-26](#).

For more information about SPBM commands, see [Chapter 7, “Shortest Path Bridging Commands,”](#) in the *OmniSwitch CLI Reference Guide*.

Configure SPBM Services

The OmniSwitch Service Manager application is used to configure the services layer of the SPBM network topology. A service is defined by a specific set of logical entities that are configured only on the backbone edge bridges (BEBs) of the network. the following configuration steps are required to define the service-based architecture for a SPBM network:

- 1 Create a SPBM service.** A Service Manager service ID is associated with a BVLAN, a backbone service instance identifier (I-SID), and a service access point (SAP) to identify the customer traffic that the service will tunnel through the provider network. See [“Creating a SPB Service” on page 3-32](#).
- 2 Configure access (customer-facing) ports.** One or more access ports are associated with a service access point (SAP) to identify to the service which ports will receive customer traffic that the service will process for tunneling through the provider network. When an access port is associated with a SAP, the SAP parameter attributes are applied to traffic received on the access port. See [“Configuring Service Access Ports” on page 3-35](#).
- 3 Define access port profile attributes.** A default Layer 2 profile is automatically assigned to an access port at the time the port is configured as an access port. This profile determines how control frames received on the port are processed. It is only necessary to configure a Layer 2 profile if the default attribute values are not sufficient. See [“Configuring Layer 2 Profiles for Access Ports” on page 3-36](#).

4 Configure a SPB service access point (SAP). A SAP binds a SPB service to an access (customer-facing) port and defines which customer traffic to tunnel through the service. Each SAP is associated to one service name, but a single service can have multiple SAPs to which it is associated. See [“Verifying the SPB Service Configuration” on page 3-34.](#)

For more information about Service Manager commands, see [Chapter 45, “Service Manager Commands,”](#) in the *OmniSwitch CLI Reference Guide*.

SPB Configuration Guidelines

Configuring a SPBM network topology involves setting up two layers of functionality: the ISIS-SPB backbone infrastructure and the Provider Backbone Bridge (802.1ah) services layer for MAC-in-MAC encapsulation. Review the guidelines in this section before attempting to configure the various components of the SPBM infrastructure and services.

ISIS-SPB

This implementation of the ISIS-SPB protocol supports only a single topology with a multi-topology identifier (MT-ID) of zero.

- The ISIS-SPB protocol instance is independent of IP IS-IS, or other network layer protocol identifiers (NLPIDs) riding in the same IS-IS implementation. However, ISIS-SPB and IP IS-IS can coexist on the same switch.
- ISIS-SPB interfaces, link state packet databases (LSPDB), and forwarding information are all created and maintained within the single ISIS-SPB instance.
- IS-IS Level 1 point-to-point adjacencies are supported; Level 2 is not supported at this time.
- SPB interfaces are associated with a link metric cost that is configurable, thus providing the ability to change the logical topology created by the ISIS-SPB instance. However, if different metric values are configured on each side of a link, ISIS-SPB will choose the higher-valued one as the metric to use for both sides. This is necessary to enforce the symmetry of SPT calculations in both directions across the link.
- Enabling SPB for the switch automatically triggers the transmission of Hello packets from the SPB interfaces, thus starting the process of discovery and forming adjacencies to build shortest path trees.

Backbone VLANs

- The BVLAN configuration must be the same on each SPB switch within the PBB network. For example, if BVLAN 10 with an ECT ID of 1 is configured on one switch, then BVLAN 10 with an ECT ID of 1 must exist on all other SPB bridges in the network to ensure proper calculation of the ISIS-SPB shortest path trees through the backbone.
- If more than one BVLAN is needed, configure each BVLAN with a different ECT algorithm ID. For example, if two BVLANS (BVLAN 4001 and BVLAN 4002) are needed for a specific SPBM topology, then create BVLAN 4001 with ECT ID 1 and BVLAN 4002 with ECT ID 2 on each switch that is going to participate in the topology.

Note. When adding another BVLAN to an existing SPBM topology instance, create the new BVLAN and its associated ECT ID on every switch first, then configure the SPB service association for the BVLAN. Creating SPB services before the BVLAN configuration is complete on all switches can cause problems with forming adjacencies or may even cause an SPB switch to drop existing adjacencies.

- In most cases one BVLAN is sufficient for virtualizing traffic through the network backbone. However, configuring more than one BVLAN provides alternate routes for tunneling customer traffic. This can also provide a form of load balancing by distributing traffic over different BVLAN segments.
- All encapsulated traffic within the BVLAN domain is unicast with a resolved source and destination BMAC addresses. Frames received on BVLAN ports that do not have an incoming source MAC address pre-programmed by ISIS-SPB are discarded.

Configuring BVLANS

The SPBM backbone VLAN (BVLAN) provides the foundation on which ISIS-SPB shortest path trees are built and SPBM services tunnel encapsulated customer data through the Provider Backbone Bridge network (PBBN). Configuring a BVLAN on a switch is also the first step in setting up the ISIS-SPB infrastructure and in making an OmniSwitch an SPB-capable node.

Note. The BVLAN configuration must be the same on each OmniSwitch that is going to participate in the SPBM network topology. So if BVLAN 4001 is created on one switch, then BVLAN 4001 must be created on all other switches in the SPBM network.

To create a BVLAN, use the **spb bvlan** command with the optional **name** parameter. For example:

```
-> spb bvlan 4001 name spb-4001
```

If the **name** parameter is not specified with this command, the VLAN ID is used for the name by default. For example, the following command creates BVLAN 4001 with “VLAN 4001” as the name:

```
-> spb bvlan 4001
```

To remove a BVLAN, use the **no** form of the **spb bvlan** command. For example:

```
-> no spb bvlan 4001
```

Assigning the Equal Cost Tree ID

ISIS-SPB calculations may result in multiple paths of equal costs. The Equal Cost Tree (ECT) ID specifies a tie-breaking algorithm that is used when ISIS-SPB is calculating a set of shortest path trees from one switch to all other switches in the SPB domain. When a BVLAN is created, an ECT ID is automatically assigned to the BVLAN. If it is the first BVLAN created on the switch, ECT ID 1 is assigned, otherwise the next available ID number is used.

Each BVLAN created must be duplicated on all other participating switches in the SPBM network and must use the same ECT ID number for that BVLAN on each switch. A BVLAN created on one switch may not be automatically assigned the same ECT ID on another switch. As a result, it may be necessary to modify the ECT ID number using the **spb isis bvlan ect-id** command. For example:

```
-> spb isis bvlan 4002 ect-id 2
```

Note. When adding another BVLAN to an existing SPBM topology instance, create the new BVLAN and its associated ECT ID on every switch first, then configure the SPB service association for the BVLAN. Creating SPB services before the BVLAN configuration is complete on all switches can cause problems with forming adjacencies or may even cause an SPB switch to drop existing adjacencies.

Configuring the Control BVLAN

One of the BVLANs configured on each switch serves as the control BVLAN for the ISIS-SPB instance. The control BVLAN exchanges ISIS-SPB control packets with neighboring SPB switches on behalf of all BVLANs configured on the local switch. The control packets are tagged with the control BVLAN ID.

By default, the first BVLAN created (or if there is only one), is the control BVLAN. To designate a different BVLAN as the control BVLAN, use the `spb isis control-bvlan` command. For example:

```
-> spb isis control-bvlan 4002
```

A control BVLAN also carries regular encapsulated SPB domain traffic in addition to ISIS-SPB control packets. In other words, a VLAN can serve as both a regular BVLAN and a control BVLAN at the same time.

Configuring the Tandem Multicast Mode

The tandem multicast mode (*,G) or (S,G) of a BVLAN is applied only to SPB services associated with the BVLAN that are using tandem replication for multicast traffic. When a BVLAN is created, the (S,G) tandem multicast mode is applied by default.

To change the tandem multicast mode for a BVLAN, use the `spb isis bvlan tandem-multicast-mode` command and specify either `gmode` (*,G) or `sgmode` (S,G). For example:

```
-> spb isis bvlan 4001 tandem-multicast-mode sgmode
-> spb isis bvlan 4002 tandem-multicast-mode gmode
```

Verifying the BVLAN Configuration

To view the BVLAN configuration for the switch, use the `show spb isis bvlans` command. For example:

```
-> show spb isis bvlans
```

```
SPB ISIS BVLANS:
```

BVLAN	ECT-algorithm	In Use	Services mapped	Num ISIDS	Tandem Multicast	Root Bridge (Name : MAC Address)
4001	00-80-c2-01	YES	YES	52	SGMODE	
4002	00-80-c2-02	YES	YES	51	SGMODE	
4003	00-80-c2-03	YES	YES	52	SGMODE	
4004	00-80-c2-04	YES	YES	51	SGMODE	

```
BVLANS :      4
```


The BVLAN is a special type of VLAN that is created and maintained by VLAN Manager. As a result, it also appears in the VLAN Manager **show** command displays. For example, in the following **show vlan** output display, VLANs 4001 through 4004 are included and “spb” appears in the “type” column:

```
-> show vlan
vlan  type  admin  oper  ip    mtu   name
-----+-----+-----+-----+-----+-----+-----
   1   std    Dis    Dis   Dis   1500  VLAN 1
1000  std    Ena    Ena   Ena   1500  VLAN 1000
4001  spb    Ena    Ena   Dis   1524  VLAN 4001
4002  spb    Ena    Ena   Dis   1524  VLAN 4002
4003  spb    Ena    Ena   Dis   1524  VLAN 4003
4004  spb    Ena    Ena   Dis   1524  VLAN 4004
4094  mcm    Ena    Dis   Dis   9198  MCM IPC
```

To view configuration information for an individual BVLAN, use the **show vlan** command and specify the BVLAN ID. For example:

```
-> show vlan 4001
Name                : VLAN 4001,
Type                : Backbone vlan,
Administrative State : enabled,
Operational State   : disabled,
IP Router Port      : disabled,
IP MTU              : 1524
```

Configuring SPB Interfaces

A port or link aggregate is configurable as a SPB interface. Each switch in the SPBM topology should have at least one SPB interface configured. The SPB interface serves more than one purpose:

- Advertises IS-IS Hello packets to discover SPB neighbors and establish adjacencies.
- After adjacencies are established, exchanges link state packets (LSPs) with SPB neighbors to build a local LSP database (LSPDB). A switch’s adjacencies are reflected in the contents of its link state packets. This relationship between adjacencies and link state allows the protocol to detect downed routers in a timely fashion.
- Serves as a network port by forwarding encapsulated SPB service traffic on backbone VLANs (BVLANS) through the SPBM Provider Backbone Bridge (PBB) network.

To configure a port or link aggregate as SPB interface, use the **spb isis interface** command. For example:

```
-> spb isis interface port 1/10
-> spb isis interface linkagg 5
```

When a port is converted to a SPB interface, the interface is automatically assigned to all existing BVLANS. There is one ISIS-SPB instance per switch, and each BVLAN and SPB interface are associated with that instance. However, it is also possible to tag SPB interfaces to carry traffic for standard VLANs.

The **spb isis interface** command is also used to optionally configure the following parameter values:

- **admin-state**—Administratively enables or disables the SPB interface. By default, the interface is enabled when the SPB interface is created.
- **hello-interval**—Specifies the amount of time, in seconds, to wait between each transmission of a hello packet from this interface. By default, the hello time interval is set to nine seconds.

- **hello-multiplier**—Specifies an integer value that is multiplied by the hello interval time to determine the amount of time, in seconds, a receiving bridge holds onto the hello packets transmitted from this interface. By default, the hello multiplier is set to three.
- **metric**—An integer value that specifies the link cost to reach the destination backbone MAC (BMAC). By default, the link cost is set to ten. Changing the link metric value provides a method for changing the logical topology as calculated by ISIS-SPB.

The following command examples change the default hello and metric values for the SPB interface:

```
-> spb isis interface port 4/7 hello-interval 60
-> spb isis interface linkagg 3 hello-multiplier 10
-> spb isis interface port 2/1 metric 100
-> spb isis interface linkagg 5 hello-interval 20 hello-multiplier 5 metric 200
```

Verifying the SPB Interface Configuration

To view the SPB interface configuration for the switch, use the **show spb isis interface** command. For example:

```
-> show spb isis interface
SPB ISIS Interfaces:
```

Interface	Level	CircID	Oper state	Admin state	Link Metric	Hello Intvl	Hello Mult
1/1	L1	1	UP	UP	10	9	3
1/2	L1	2	UP	UP	10	9	3
1/3	L1	3	UP	UP	10	9	3
1/4	L1	4	DOWN	UP	10	9	3
1/5	L1	5	DOWN	UP	10	9	3
1/6	L1	6	DOWN	UP	22	9	3
1/7	L1	7	DOWN	UP	10	9	3

Interfaces : 7

Configuring Global ISIS-SPB Parameters

This section describes the global configuration for the ISIS-SPB instance, which includes the following:

- [“Configuring the System Name” on page 3-27.](#)
- [“Configuring the SPB Bridge Priority” on page 3-27.](#)
- [“Configuring the ISIS-SPB Area Address” on page 3-27.](#)
- [“Configuring the Shortest Path Source ID” on page 3-28.](#)
- [“Configuring the Shortest Path First Wait Time” on page 3-28.](#)
- [“Configuring the Link State Packet Wait Time” on page 3-29.](#)
- [“Configuring the Overload State” on page 3-30.](#)
- [“Configuring Redundant Switches for Graceful Restart” on page 3-31.](#)
- [“Enabling/Disabling ISIS-SPB” on page 3-31.](#)

To verify the global configuration parameter values for the switch, use the **show spb isis info** command. For example:

```
-> show spb isis info
SPB ISIS Bridge Info:
  System Id           = e8e7.3233.1831,
  System Hostname     = BEB-1,
  SPSourceID         = 03-18-31,
  SPBM System Mode    = auto,
  BridgePriority      = 32768 (0x8000),
  MT ID              = 0,
  Control BVLAN      = 4001,
  Area Address        = 0.0.0,
  Level Capability    = L1,
  Admin State         = UP,
  LSDB Overload      = Disabled,
  Last Enabled        = Thu Aug  2 22:43:19 2012,
  Last SPF            = Fri Aug  3 18:15:51 2012,
  SPF Wait            = Max: 1000 ms, Initial: 100 ms, Second: 300 ms,
  LSP Lifetime        = 1200,
  LSP Wait            = Max: 1000 ms, Initial: 0 ms, Second: 300 ms,
  Graceful Restart    = Disabled,
  GR helper-mode      = Disabled,
  # of L1 LSPs       = 8
```

Configuring the System Name

Configuring a system name is required on each switch that is going to participate in the SPBM topology. To configure a system name for the switch, use the **system name** command. For example:

```
-> system name BEB-1
```

ISIS-SPB advertises the system name to identify the switch to other SPB peer switches.

Configuring the SPB Bridge Priority

A bridge is ranked within the SPB topology by its bridge ID (an eight byte hex number). The bridge priority value makes up the upper two bytes of the eight-byte SPB bridge ID. The lower six bytes of the Bridge ID contain the system ID, which is the dedicated bridge MAC address of the SPB bridge.

The bridge priority is used in shortest path tree calculations. The lower the priority value, the higher the priority. Setting a different bridge priority value on different SPB bridges will override the system ID significance during the shortest path tree (SPT) calculation.

By default, all SPB switches are assigned a priority value of 32768. To change the bridge priority value for a switch, use the **spb isis bridge-priority** command. For example:

```
-> spb isis bridge-priority 25590
```

Configuring the ISIS-SPB Area Address

By default, the IS-IS area address for the ISIS-SPB instance is set to 0.0.0, which is typically sufficient for this implementation of SPBM. Both ISIS-SPB and ISIS-IP instances may coexist on the same switch as long as they don't use the same area address.

If changing the area address is necessary, use the **spb isis area-address** command. For example:

```
-> spb isis area-address 1.1.1
```

Note. Each switch that is going to participate in the SPB topology must use the same area address and must use an address that is different from the ISIS-IP area address.

Configuring the Shortest Path Source ID

The shortest path (SP) source ID, identifies the source of multicast frames and is relevant only in multi-cast tandem replication mode. By default, the last three least significant bytes of the system ID (local bridge MAC address) is used for the source ID value.

To change the source ID value, use the **spb isis source-id** command. For example:

```
-> spb isis source-id 07-0b-d3
```

To set the source ID back to the default value, use the **spb isis source-id** command with the **auto** parameter. For example:

```
-> spb isis source-id auto
```

Configuring the Shortest Path First Wait Time

The **spb isis spf-wait** command is used to configure the time intervals between the first, second, and subsequent ISIS-SPB shortest path first (SPF) calculations.

Subsequent SPF calculations, if required, are generated at exponentially increasing intervals of the SPF second wait time interval until the maximum wait time interval value is reached. For example, if the second-wait interval value is set to 1000 milliseconds, then the next SPF calculation is triggered after 2000 milliseconds and the next SPF calculation after that is triggered at 4000 milliseconds, and so on, until the maximum-wait interval value is reached.

When the maximum interval value is reached, the SPF wait interval will stay at the maximum value until there are no more SPF calculations scheduled during that interval. After a full interval without any SPF calculations, the SPF wait interval will reset back to the initial wait time interval value.

The following **spb isis spf-wait** command parameters are used to configure the SPF timers:

- **max-wait**—The maximum number of milliseconds to wait between two consecutive SPF calculations. The default maximum wait time value is set to 1000 milliseconds. Specify a maximum value that is the same or greater than the second wait time value.
- **initial-wait**—The number of milliseconds to wait before triggering an initial SPF calculation after a topology change. The default initial wait time value is set to 100 milliseconds. Specify a value that is the same or less than the maximum wait time value.
- **second-wait**—The number of milliseconds to wait between the first and second SPF calculation. The default second wait time value is set to 300 milliseconds. Specify a value that is the same or less than the maximum wait time value.

For example, the following command changes the SPF wait time values for the local SPB instance:

```
-> spb isis spf-wait max-wait 2500 initial-wait 1000 second-wait 1500
```

To change one or more of the wait time values, it is only necessary to specify the parameter for the desired change. For example:

```
-> spb isis spf-wait max-wait 5000
-> spb isis spf-wait initial-wait 1000
-> spb isis spf-wait second-wait 2000
```

To set the wait time values back to the default settings, use the **spb isis spf-wait** command without specifying any of the parameters. For example:

```
-> spb isis spf-wait
```

Configuring the Link State Packet Wait Time

The **spb isis lsp-wait** command is used to configure the time intervals between the first, second, and subsequently generated link state packets (LSPs).

Subsequent LSP, if required, are generated at exponentially increasing intervals of the LSP second wait time interval until the maximum value is reached. For example, if the second-wait interval value is set to 10 seconds, then the next LSP is generation is triggered after 20 seconds and the next LSP generated after that is triggered at 40 seconds, and so on, until the maximum wait time interval value is reached.

When the maximum interval value is reached, the LSP wait interval will stay at the maximum value until there are no more LSP generations during that interval. After a full interval without any LSP generations, the LSP wait interval will reset back to the initial wait time interval value.

The following **spb isis lsp-wait** command parameters are used to configure the SPF timers:

- **max-wait**—The maximum number of seconds to wait between two consecutively generated LSPs. The default maximum wait time value is set to 1000 milliseconds. Specify a maximum value that is the same or greater than the second wait time value.
- **initial-wait**—The number of seconds to wait before triggering an initial LSP generation after a topology change. The default initial wait time value is set to 0 milliseconds. Specify a value that is the same or less than the maximum wait time value.
- **second-wait**—The minimum number of seconds to wait between the first and second generated LSPs. The default second wait time value is set to 300 milliseconds. Specify a value that is the same or less than the maximum wait time value.

For example, the following command changes the LSP wait time values for the local SPB instance:

```
-> spb isis lsp-wait max-wait 2000 initial-wait 1000 second-wait 1500
```

To change one or more of the wait time values, it is only necessary to specify the parameter for the desired change. For example:

```
-> spb isis lsp-wait max-wait 5000
-> spb isis lsp-wait initial-wait 2500
-> spb isis lsp-wait second-wait 3000
```

To set the wait time values back to the default settings, use the **spb isis lsp-wait** command without specifying any of the parameters. For example:

```
-> spb isis lsp-wait
```

Configuring the Overload State

This implementation of ISIS-SPB supports the overload state mechanism, which allows an instance of ISIS-SPB to inform its neighbors that the instance is nearing or exceeding its capabilities. When peers see that a switch is advertising in this state, they will select an alternate path around the overloaded switch.

The ISIS-SPB instance for a switch may dynamically trigger the overload state condition when the instance detects that it is nearing or has reached resource limits. However, it is also possible to manually trigger the overload state condition using the **spb isis overload** command. For example:

```
-> spb isis overload
```

Some advantages of manually triggering the overload state condition, even if the instance is nowhere near its resource limits, are as follows:

- The switch is designated as “leaf node” that should never carry transit traffic. Configuring the link metric value for the SPB interfaces on the switch and attached peers is another method for preventing the switch from receiving transit traffic, but enabling the overload state is a much quicker way to achieve the same results and requires less configuration.
- When there is a need to remove the switch from service (temporarily or permanently). In this scenario, network availability is increased because peer switches will detect the overload state of the switch and gracefully transition to alternate paths, while the “manually overloaded” switch continues to forward packets. Just simply shutting the switch down would cause more disruption to network traffic.

When the overload state is either dynamically or manually enabled for the switch, the overload bit is set in LSP 0 to indicate that this ISIS-SPB instance is not available to accept transit traffic. However, an ISIS-SPB switch operating in the overload state is still used only if there is no alternate path to reach the intended destination.

When the overload state is enabled, the switch will operate in this state for an infinite amount of time. To configure the switch to remain in the overload state for only a specific amount of time (in seconds), use the **spb isis overload** command with the optional **timeout** parameter. For example:

```
-> spb isis overload timeout 500
```

To disable the overload state, use the **no** form of the **spb isis overload** command. For example:

```
-> no spb isis overload
```

It is also possible to specify that the overload state is enabled for the switch after every system bootup. This is done using the **spb isis overload-on-boot** command, which also has an optional **timeout** parameter. For example:

```
-> spb isis overload-on-boot timeout 500
```

To disable the overload-on-boot option, use the **no** form of the **spb isis overload-on-boot** command. For example:

```
-> no spb isis overload-on-boot timeout 500
```

Note that the **no spb isis overload** command does not disable the overload-on-bootup option.

Configuring Redundant Switches for Graceful Restart

By default, ISIS-SPB graceful restart is disabled. When graceful restart is enabled, the switch can either be a helper (which helps a neighbor router to restart) or a restarting router, or both. When graceful restart is enabled on the switch, the helper mode is automatically enabled by default.

To configure ISIS-SPB graceful restart support on an OmniSwitch, use the **spb isis graceful-restart** command. For example, to configure graceful restart on the router, enter:

```
-> spb isis graceful-restart
```

The helper mode can be disabled on the switch with the **spb isis graceful-restart helper** command. For example, to disable the helper support for neighboring switches, enter the following:

```
-> ip isis graceful-restart helper disable
```

To disable support for graceful restart, use the **no** form of the **spb isis graceful-restart** command. For example:

```
-> no spb isis graceful-restart
```

Enabling/Disabling ISIS-SPB

By default ISIS-SPB is disabled on the switch. To enable ISIS-SPB, use the **spb isis admin-state** command with the **enable** option. For example:

```
-> spb isis admin-state enable
```

To disable the ISIS-SPB instance on the switch, enter the **spb isis admin-state** command with the **disable** option. When the ISIS-SPB status is disabled for the switch, the related configuration settings and statistics are retained.

```
-> spb isis admin-state disable
```

Note. Enabling ISIS-SPB on a switch starts the process of ISIS-SPB discovery, adjacency building, and shortest path tree calculations. Make sure that the SPBM configuration is set up first, then enable ISIS-SPB on each switch that will participate in the SPBM network.

Creating a SPB Service

A SPB service is identified by a service ID number, which represents an association between a backbone service instance identifier (I-SID) and an existing BVLAN. Basically, creating a SPB service binds the backbone I-SID to a BVLAN ID. All traffic mapped to the specific I-SID is then encapsulated and forwarded on the associated BVLAN to the intended destination.

The **service spb** command is used to create a SPB service. For example, the following command creates SPB service 1 and binds I-SID 100 to BVLAN 4001:

```
-> service spb 1 isid 500 bvlan 4001
```

The BVLAN ID specified with the **service spb** command must already exist in the switch configuration. However, the I-SID number specified creates a new I-SID that is bound to the BVLAN for this service.

Note. When adding another BVLAN to an existing SPBM topology instance, create the new BVLAN and its associated ECT ID on every switch first, then configure the SPB service association for the BVLAN. Creating SPB services before the BVLAN configuration is complete on all switches can cause problems with forming adjacencies or may even cause an SPB switch to drop existing adjacencies.

Modifying Default SPB Service Parameters

The following SPB service parameter values are set by default at the time the service is created. If necessary, use the specified commands to change the default values.

Parameter Description	Command	Default
Service description.	service spb description	None
Administrative status for statistics collection.	service spb stats	Disabled
Multicast replication mode	service spb multicast-mode	head-end
VLAN translation	service spb vlan-xlation	Disabled
Administrative status of the service	service spb admin-state	Disabled

Refer to the *OmniSwitch CLI Reference Guide* for more information about the above parameters and related commands.

Using VLAN Translation

VLAN translation refers to the egress translation of VLAN tags on service access points (SAPs). When enabled for a service, the VLAN tags for outgoing customer frames on SAPs associated with that service are processed according to the local SAP configuration (the SAP on which the frames will egress) and not according to the configuration of the SAP on which the frames were received.

- If the local SAP is configured for untagged traffic (*slot/port:0*), the egress traffic is always sent out as untagged.
- If the local SAP is configured for 802.1q-tagged traffic (*slot/port:ctag*), the egress traffic is single-tagged with the tag value specified by the *ctag* (customer VLAN tag) value.
- If the local SAP is configured for double-tagged traffic (*slot/port:outer_tag.: inner_tag*), the egress traffic is double-tagged with the tag values specified by the *outer_tag* and *inner_tag* values.

When VLAN translation is disabled, frames simply egress without any modification of the VLAN tags. In other words, the frames are transparently bridged without tag modification.

The following table shows the required translation (tag is added or replaced) that takes place when the egress SAP configuration is applied to the possible frame types (untagged, tagged, double-tagged). Note that in this table the terms “ITAG” and “OTAG” refer to inner tag and outer tag, respectively.

Egress SAP (action required based on SAP type)			
	Untagged SAP	Single Tagged SAP	Double-Tagged SAP
Incoming Frame	Remove OTAG	Replace OTAG	Replace OTAG
	Remove ITAG	(Note: Replace = implicit add) Remove ITAG	(Note: Replace = implicit add) Add/Replace ITAG
Untagged	No tags, so no action taken	Add the SAP OTAG	Add the SAP OTAG Add the SAP ITAG.
Single-tagged	Remove the OTAG	Replace the OTAG	Add ITAG Replace OTAG
Double-tagged	Remove the ITAG Remove the ITAG	Remove the ITAG Replace the OTAG	Replace ITAG Replace OTAG

VLAN translation is enabled at two different levels: the service level and access port level. To enable translation at the service level, use the **service spb vlan-xlation** command. For example:

```
-> service spb 1 vlan-xlation enable
```

To enable VLAN translation for all services, use the **all** parameter with the same command. For example:

```
-> service spb all vlan-xlation enable
```

To disable VLAN translation, use the **service spb vlan-xlation** command with the disable parameter. For example:

```
-> service spb 1 vlan-xlation disable
-> service spb all vlan-xlation disable
```

To enable VLAN translation at the port level, use the **service access vlan-xlation** command. For example:

```
-> service access port 1/11 vlan-xlation enable
```

See “[Configuring Service Access Ports](#)” on page 3-35 for more information.

Enable the Service

By default, the SPB service is disabled when the service is created. Once the service is created and any optional service parameters are configured, use the **service spb admin-state** command with the **enable** option to enable the service. For example:

```
-> service spb 1 admin-state enable
```

To disable the service, enter the following command:

```
-> service spb 1 admin-state disable
```

Deleting a SPB Service

Before deleting a service from the switch configuration, disable the administrative status of the service. Once this is done, use the **no** form of the **service spb** command to delete the service. For example:

```
-> no service spb 1
```

Verifying the SPB Service Configuration

To view the SPB service configuration for the switch, use the **show spb isis services** command. For example:

```
-> show service spb
```

Legend: * denotes a dynamic object

SPB Service Info

```
SystemId : 00e0.ble7.0188, SrcId : 0x70188, SystemName : BEB-1
```

ServiceId	Adm	Oper	Stats	SAP Count	Bind Count	Isid	BVlan	MCast Mode	(T/R)
1	Up	Up	N	4	1	1000	4001	Headend	(0/0)
2	Up	Up	N	4	1	1001	4001	Headend	(0/0)
3	Up	Up	N	4	1	1002	4001	Headend	(0/0)
4	Up	Up	N	4	1	1003	4001	Headend	(0/0)
5	Up	Up	N	4	1	1004	4001	Headend	(0/0)
6	Up	Up	N	4	1	1005	4001	Headend	(0/0)
7	Up	Up	N	4	1	1006	4001	Headend	(0/0)
8	Up	Up	N	4	1	1007	4001	Headend	(0/0)
9	Up	Up	N	4	1	1008	4001	Headend	(0/0)
10	Up	Up	N	4	1	1009	4001	Headend	(0/0)

To view the configuration for an individual service, use the **show spb isis services** command and specify the SPB service ID. For example:

```
-> show service spb 1
```

SPB Service Detailed Info

```
Service Id       : 1, Description      : ,
ISID             : 1000, BVlan           : 4001,
Multicast-Mode  : Headend, Tx/Rx Bits     : 0/0,
Admin Status    : Up, Oper Status     : Up,
Stats Status    : No, Vlan Translation : No,
Service Type    : SPB, Allocation Type  : Static,
MTU             : 9194, Def Mesh VC Id  : 1,
SAP Count       : 4, SDP Bind Count  : 1,
Ingress Pkts    : 0, Ingress Bytes   : 0,
Egress Pkts     : 0, Egress Bytes    : 0,
Mgmt Change     : 08/10/2012 13:14:43, Status Change  : 08/10/2012 13:14:00
```

Configuring Service Access Points (SAPs)

A SAP identifies the location where customer traffic enters the Provider Backbone Bridge Network (PBBN) edge, the type of customer traffic to service, parameters to apply to the traffic, and the service that will process the traffic for tunneling through the provider network.

Configuring a SAP requires several steps. These steps are outlined here and further described throughout this section:

- Configure customer-facing ports or link aggregates as service access ports.
- Configure Layer 2 profiles to determine how control packets are processed on access ports.
- Create a SAP by associating a SAP ID with a SPB service ID. A SAP ID is comprised of an access port and an encapsulation value, which is used to identify the type of customer traffic (untagged, single-tagged, or double-tagged) to map to the associated service.

SAP Configuration Guidelines

Consider the following when configuring a SAP:

- A SAP is a unique local entity for any given device. The same SAP ID value can be used on other BEB switches.
- There are no SAPs configured by default; explicit configuration of a SAP is required.
- A SAP is administratively disabled at the time the SAP is created.
- When a SAP is deleted, all configuration parameters for the SAP are also deleted.
- A SAP is owned by and associated with the service that was specified at the time the SAP was created.
- If a port is administratively shutdown, all SAPs on that port become operationally out of service.
- Both fixed ports and link aggregates are configurable as access ports. Only access ports are associated with SAPs.
- Bridging functionality is not supported on access ports or link aggregates.

Configuring Service Access Ports

Each SAP is comprised of an access port or link aggregate and an encapsulation type value. Access ports are customer-facing ports that reside on a provider edge router. Traffic received on these ports is classified for one or more SAPs and forwarded onto the intended destination by the associated SPB service.

To configure a port or link aggregate as an access port, use the `service access` command. For example, the following command configures port 1/2 and link aggregate 5 as access ports:

```
-> service access port 1/2
-> service access linkagg 5
```

To revert an access port back to a regular switch port or link aggregate, use the no form of the service access command. For example:

```
-> no service access port 1/2
-> no service access linkagg 5
```

VLAN Translation on Access Ports

VLAN translation refers to the egress translation of VLAN tags on service access points (SAPs). For more information about VLAN translation is applied, see [“Using VLAN Translation” on page 3-32](#).

By default, VLAN translation is disabled on access ports. Enabling VLAN translation on an access port implicitly enables translation for all SAPs associated with that port. However, translation must also be enabled for the services associated with these SAPs. This ensures that all SAPs associated with a service will apply VLAN translation.

To enable VLAN translation on an access port, use the **service access vlan-xlation** command with the **enable** option. For example:

```
-> service access port 1/3 vlan-xlation enable
-> service access linkagg 10 vlan-xlation enable
```

To disable VLAN translation on an access port, use the **service access vlan-xlation** command with the **disable** option. For example:

```
-> service access port 1/3 vlan-xlation disable
-> service access linkagg 10 vlan-xlation disable
```

Configuring Layer 2 Profiles for Access Ports

A Layer 2 profile determines how control frames ingressing on an access port are processed. When a port is configured as an access port, a default Layer 2 profile (**def-access-profile**) is applied to the port with the following default values for processing control frames:

Protocol	Default
STP	tunnel
802.1x	drop
802.1ab	drop
802.3ad	peer
GVRP	tunnel
MVRP	tunnel
AMAP	discard
CISCO PDU	discard
CISCO VLAN	discard
CISCO uplink	discard

If the default profile values are not sufficient, use the **service l2profile** command with the **tunnel**, **discard**, and **peer** options to create a new profile. For example, the following command creates a profile named “DropL2”:

```
-> service l2profile DropL2 stp discard gvrp discard 802.1ab discard
```

Consider the following when configuring Layer 2 profiles:

- Not all of the control protocols are currently supported with the **peer**, **tunnel**, and **discard** parameters. Use the following table to determine the parameter combinations that are supported:

Protocol	Reserved MAC	peer	discard	tunnel
STP	01-80-C2-00-00-00	no	yes	yes
802.1x	01-80-C2-00-00-03	no	yes	yes
802.1ab	01-80-C2-00-00-0E	yes	yes	yes
802.3ad	01-80-C2-00-00-02	yes	no	no
GVRP	01-80-C2-00-00-21	no	yes	yes
MVRP	—	no	yes	yes
AMAP	00-20-DA-00-70-04	yes	yes	no
CISCO PDU	01-00-0C-CC-CC-CD	no	yes	yes
CISCO VLAN	01-00-0C-CC-CD-CE	no	yes	yes
CISCO uplink	01-00-0C-CC-CD-CF	no	yes	yes

- When a profile is created, the new profile inherits the default profile settings for processing control frames. The default settings are applied with the new profile unless they are explicitly changed. For example, the profile “DropL2” was configured to discard STP, GVRP, and 802.1ab frames. No other protocol settings were changed, so the default settings still apply for the other protocols.
- Remove any profile associations with access ports before attempting to modify or delete the profile.

To delete a Layer 2 profile, use the **no** form of the **service l2profile** command. For example, the following command deletes the “DropL2” profile:

```
-> no service l2profile DropL2
```

Use the **show service l2profile** command to view a list of profiles that are already configured for the switch. This command also displays the attribute values for each profile.

Assigning Layer 2 Profiles to Access Ports

After a Layer 2 profile is created, it is then necessary to assign the profile to an access port or link aggregate. When this is done, the current profile associated with the port is replaced with the new profile.

The **service access l2profile** command is used to assign a new profile to an access port. For example, the following command assigns the “DropL2” profile to access port 1/4:

```
-> service access port 1/4 l2profile DropL2
-> service access linkagg 5 l2profile DropL2
```

To change the profile associated with the access port back to the default profile (**def-access-profile**), use the **default** option with the **service access l2profile** command. For example:

```
-> service port 1/4 l2profile default
-> service access linkagg 5 l2profile DropL2
```

Use the **show service access** command to display profile associations for access ports.

Verifying the Access Port Configuration

To view the access port configuration for the switch, use the `show service access` command. For example:

```
-> show service access
Port      Link  SAP      SAP      Vlan
Id        Status Type     Count    Xlation L2Profile
-----+-----+-----+-----+-----+-----
1/11      Up    Manual   100      N        def-access-profile
1/12      Down  Manual   100      N        def-access-profile
1/13      Down  Manual   100      N        def-access-profile
1/14      Down  Manual   100      N        def-access-profile
1/15      Down  Dynamic  0         Y        def-access-profile
1/16      Up    Dynamic  1         Y        def-access-profile
1/17      Down  Dynamic  0         Y        def-access-profile
1/18      Down  Dynamic  0         Y        def-access-profile
```

Total Access Ports: 8

Creating the Service Access Point

Each SPB service is bound to at least one Service Access Point (SAP). A SAP identifies the point at which customer traffic enters the Provider Backbone Bridge Network (PBBN). Creating a SAP on a SPB switch designates that switch as a Backbone Edge Bridge (BEB) in the PBBN. An SPB switch that does not have a SAP but does have at least one BVLAN and an SPB interface is designated as Backbone Core Bridge (BCB) in the PBBN.

Once the SPB topology is determined and switches that will serve as BEBs are identified, a SAP is created on each BEB. A SAP is created by associating a SAP ID with a SPB service. A SAP ID is comprised of a customer-facing port (referred to as an access port) and an encapsulation value that is used to identify the type of customer traffic (untagged, single-tagged, or double-tagged) to map to the associated service.

The `service spb sap` command is used to configure a SAP. This command specifies the SPB service ID number and the SAP ID (slot/port:encapsulation). The following parameter values are used with this command to specify the encapsulation value:

SAP Encapsulation Value	Customer Traffic Serviced
0 (null)	All untagged packets; tagged packets are dropped.
all	All tagged and untagged packets not already classified into another SAP*
<i>qtag</i>	Only traffic 802.1q-tagged with the specified VLAN ID.
<i>outer_qtag.innnger_qtag</i>	Only traffic double-tagged (QinQ) with the specified outer and inner VLAN IDs.

*Note that the **:all** (wildcard) parameter is also configurable as the inner tag value for double-tagged frames (for example, "10:all" specifies double-tagged packets with an outer tag equal to 10 and an inner tag with any value).

The following `service spb sap` command example creates a SAP that will direct customer traffic ingress-ing on access port 1/4 that is tagged with VLAN ID 50 to service 100:

```
-> service spb 100 sap 1/4:50 description "BEB1 to SPB100 CVLAN 50"
```

In the above example, the 1/4:50 designation is referred to as the SAP ID or the encapsulation ID. This means that if no other SAPs are configured for port 1/4, then any traffic ingressing on that port is dropped if the traffic is not tagged with VLAN 50.

It is possible to configure more than one SAP for the same access port, which provides a method for segregating incoming traffic into multiple services. For example, the following SAP configuration for port 2/3 sends incoming traffic to three different services based on the VLAN tags of the frames received:

```
-> service spb 2000 sap port 2/3:all
-> service spb 200 sap port 2/3:100
-> service spb 1000 sap port 2/3:100.200
```

In this example,

- Frames double-tagged with 100 (outer tag) and 200 (inner tag) are sent on service 1000.
- Frames single-tagged with VLAN 100 are sent on service 200.
- All other frames (those that are not single-tagged with 100 or double-tagged with 100 and 200) are sent on service 2000.

The following SAP ID classification precedence is applied when there are multiple SAPs for one access port:

- 1 Double-tagged (Outer VLAN + Inner VLAN) - Highest
- 2 Double-tagged (Outer VLAN + all)
- 3 Single-tagged (VLAN)
- 4 Single-tagged (wildcard)
- 5 Untagged - Lowest.

Modifying Default SAP Parameters

The following parameter values are set by default at the time the SAP is created. If necessary, use the specified commands to change the default values.

Parameter Description	Command	Default
SAP description.	service spb sap description	None
SAP trust mode	service spb sap trusted	Trusted
Administrative status for the SAP	service spb sap admin-state	Disabled
Administrative status for statistics collection.	service spb sap stats	Disabled

Refer to the *OmniSwitch CLI Reference Guide* for more information about the above parameters and commands.

Configuring the SAP Trust Mode

The [service spb sap trusted](#) command is used to configure the trust mode for a SAP. A trusted SAP can accept 802.1p values in incoming packets; an untrusted SAP will set any 802.1p values to zero in incoming packets, unless an 802.1p value is configured with this command.

Note that untagged Layer 2 control packets (for example, BPDU, GVRP, and AMAP) are always tunneled (if enabled) through the Provider Backbone Bridge (PBB) network with the default EXP bits set to 7, so that they can arrive at the destination bridge at the highest COS queue of 7. As a result, trusted and untrusted SAPs configured on the access ports will not affect the Layer 2 control packets ingressing on the access ports.

By default, a SAP is trusted with the priority set to best effort (zero). Use the **no** form of the **service spb sap trusted** command with the **priority** option to change the SAP mode to untrusted. For example:

```
-> service spb 100 sap 1/4:50 no trusted priority 7
```

When a SAP is trusted, the priority value contained in tagged customer packets is used; untagged packets are assigned the default priority value (zero). When a SAP is untrusted, the priority value configured for the SAP is assigned to both tagged and untagged customer packets.

Enabling/Disabling the SAP

By default, a SAP is disabled at the time the SAP is created. To enable the SAP administrative status, use the **service spb sap admin-state** command. For example:

```
-> service spb 100 sap port 1/4:50 admin-state enable
-> service spb 200 sap linkagg 5:all admin-state enable
```

To disable the SAP, enter the following command:

```
-> service spb 100 sap port 1/4:50 admin-state disable
-> service spb 200 sap linkagg 5:all admin-state disable
```

Deleting the SAP

When a SAP is administratively disabled, the SAP configuration is not removed from the switch. To delete a SAP from the switch configuration, use the **no** form of the **service spb sap** command. For example:

```
-> service spb 100 no sap port 1/4:50
-> service spb 200 no sap linkagg 5:all
```

Verifying the SAP Configuration

A SAP is a type of virtual port that is associated with a SPB service. To determine the SAP configuration for a specific service, use the **show service spb ports** command to view the virtual ports associated with a specific service. For example:

```
-> show service spb 1 ports
```

Legend: * denotes a dynamic object

SPB Service Info

```
Admin : Up, Oper : Up, Stats : N, Mtu : 9194, VlanXlation : N,
ISID : 1000, Bvlan : 4001, MCast-Mode : Headend, Tx/Rx : 0/0
```

Identifier	Sap Trusted:Priority/			Sap Description /	
	Adm	Oper	Stats	Sdp SystemId:Bvlan	Intf Sdp SystemName
sap:1/11:1000	Up	Up	N	Y:x	1/11 -
sap:1/12:1000	Up	Down	N	Y:x	1/12 -
sap:1/13:1000	Up	Down	N	Y:x	1/13 -
sap:1/14:1000	Up	Down	N	Y:x	1/14 -
sdp:32776:1*	Up	Up	Y	e8e7.3233.1831:4001	1/1 BEB-1

Total Ports: 5

To then view configuration information for a specific SAP, use the **show service spb sap** command. For example:

```
-> show service spb 1 sap port 1/11:1000
```

```
SAP Detailed Info
```

SAP Id	: 1/11:1000,	Description	: ,
Admin Status	: Up,	Oper Status	: Up,
Stats Status	: No,	Vlan Translation	: No,
Service Type	: SPB,	Allocation Type	: Static,
Trusted	: Yes,	Priority	: 0,
Ingress Pkts	: 0,	Ingress Bytes	: 0,
Egress Pkts	: 0,	Egress Bytes	: 0,
Mgmt Change	: 08/07/2012 23:39:29,	Status Change	: 08/10/2012 15:13:08

Verifying the SPB Backbone and Services

Displaying the SPBM configuration is helpful to verify the actual configuration on each SPB switch in the topology and to troubleshoot ISIS-SPB backbone and SPB service connectivity.

Verifying the ISIS-SPB Backbone Configuration

To display information about the ISIS-SPB infrastructure (backbone), use the **show** commands listed in this section.

show spb isis info	Displays the global status and configuration for the ISIS-SPB instance on the switch.
show spb isis bvlan	Displays the backbone VLAN (BVLAN) configuration for the switch.
show spb isis interface	Displays the SPB interface (network port) configuration for the switch.
show spb isis adjacency	Displays information about the ISIS-SPB adjacencies created for the SPB switch.
show spb isis database	Displays ISIS-SPB topology information maintained in the link state database (LSDB).
show spb isis nodes	Displays the discovered node-level parameter values for all of the ISIS-SPB switches participating in the topology.
show spb isis unicast-table	Displays the unicast forwarding information for the BVLAN topology.
show spb isis services	Displays a network-wide view of existing services to help verify that SPB services are correctly advertised and learned by ISIS-SPB
show spb isis spf	Displays the shortest path first (SPF) information to all known SPB switches for a specific BVLAN.
show spb isis multicast-table	Displays the multicast forwarding entries for services.
show spb isis multicast-sources	Displays all the known multicast sources across the SPB domain and BVLANs.
show spb isis multicast-sources-spf	Displays the shortest path first (SPF) readability for a known multicast source bridge for a specific BVLAN.
show spb isis ingress-mac-filter	Displays the ingress MAC filter for multicast traffic for a given BVLAN operating in the (*,G) mode.

Verifying the SPB Service Configuration

To display information about the Service Manager configuration for SPB service connectivity, use the **show** commands listed in this section

show service access	Displays the service access (customer-facing) port configuration.
show service l2profile	Displays the Layer 2 profile definitions. These profiles are applied to service access ports to determine how Layer 2 control protocol frames are processed on these ports.
show service	Displays the service configuration.
show service spb ports	Displays all the virtual ports (SAPs, SDPs) that are associated with an SPB service.

show service spb sap	Displays the configuration information for the specified SAP ID associated with the specified service.
show service sdp	Displays the dynamic Service Distribution Point (SDP) configuration.
show service mesh-sdp	Displays the dynamic SDP-to-service binding configuration.
show service spb debug-info	Displays debug information for the virtual ports associated with the SPB service.
show service spb counters	Displays the traffic statistics for the specified SPB service and associated virtual ports.
clear service spb counters	Clears the traffic statistics for the specified SPB service and associated virtual ports.

For more information about the resulting displays from these commands, see the *OmniSwitch CLI Reference Guide*.

4 Virtual Machine Classification

Server virtualization in the data center is happening now. The OmniSwitch implements two features to help automate the discovery and movement of virtual machines (VMs):

- **Virtual Network Profile (vNP)**—A vNP is basically a Universal Network Profile (UNP) that will classify virtual machines in the same manner as any other device connected to a UNP port. Once a virtual machine (or any other such device) is assigned to a vNP, the virtual machine traffic is bound to a VLAN or Shortest Path Bridging (SPB) service as defined in the profile. In addition, any QoS policies associated with the profile are also applied to the VM traffic.
- **Edge Virtual Bridging (EVB)**—based on the IEEE 802.1Qbg standard, EVB defines architecture to standardized connections between hosts with virtual machines and switches. The OmniSwitch implementation provides the ability for the switch to interact with EVB servers.
 - EVB runs between an EVB bridge (OmniSwitch running EVB) and an EVB station (server running EVB) to discover VMs running on the station or detect VM movement and associate each VM with an appropriate VLAN bridging instance.
 - The Link Layer Delivery Protocol (LLDP) is used by the EVB bridge to discover the EVB station and establish links between the station and bridge.

In This Chapter

This chapter describes the functionality provided with the vNP and EVB features and how these features are used to classify virtual machines into the OmniSwitch bridging or SPB service domain. It provides information about configuring vNP and EVB through the Command Line Interface (CLI). CLI commands are used in the configuration examples; for more details about the syntax of commands, see the *OmniSwitch CLI Reference Guide*.

The following topics and procedures are included in this chapter:

- [“UNP \(vNP\) Specifications” on page 4-2.](#)
- [“EVB Specifications” on page 4-2.](#)
- [“Server Virtualization Overview” on page 4-3.](#)
- [“Classifying Virtual Machines” on page 4-4.](#)

For more information about UNP, see [Chapter 26, “Configuring Universal Network Profiles,”](#) in the *OmniSwitch AOS Release 7 Network Configuration Guide*.

UNP (vNP) Specifications

Platforms supported	OmniSwitch 10K, 6900
Number of UNPs per switch	4K (includes static and dynamic profiles).
Number of UNP users per switch	2K
Authentication type	MAC-based authentication
Profile type	VLAN or Shortest Path Bridging (SPB)
UNP port type	Bridge (VLAN-based classification) or access (service-based classification)
UNP classification rules	MAC address, MAC-range, IP address, and VLAN tag
Number of QoS policy lists per switch	32 (includes the default list)
Number of QoS policy lists per UNP	1

EVB Specifications

Platforms Supported	OmniSwitch 10K, 6900
OmniSwitch Software License	Data Center
IEEE Standards Supported	P802.1Qbg Standard Draft, Revision D2.2. February 18, 2012— <i>Virtual Bridged Local Area Networks—Amendment 21: Edge Virtual Bridging</i>
EVB mode	Bridging (virtual machines request the required CVLAN ID tag)
Edge Relay (ER) support	Single ER per switch port. The ER can operate as a Virtual Ethernet Port Aggregator (VEPA) or as a Virtual Ethernet Bridge (VEB).

Server Virtualization Overview

Data center virtualization, as well as the virtualization of many other types of networks, facilitates the sharing and management of network resources. The virtual network is in essence a single, logical entity that is not restricted by physical boundaries. This is similar in concept to the use of virtual local area networks (VLANs) and virtual private networks (VPNs).

A key factor in the transition of the traditional data center network is the widespread adoption of server virtualization. A virtualized server is comprised of the following main components (see Figure 1):

- Guest virtual machine—the software representation of a self-contained operating system and application software instance that runs on a host server.
- A host server—the underlying physical hardware that provides computing resources for guest virtual machines.
- Hypervisor—a software program that provisions and manages the guest virtual machines for the host server. This program enables multiple guest virtual machines to run in isolation, side-by-side on the same physical server.




Figure 1: Virtualized Server

Server virtualization provides the following benefits:

- **Consolidation of server resources.** The ability to run multiple virtual machines on one server reduces the number of servers required to support data center applications. Using virtual machines maximizes physical server resources; the server is no longer tied to only one application or a single task.
- **Power savings.** The need for fewer servers means less power usage. In addition, existing servers can remain powered down until they are needed for virtual machine support. It is not necessary to have an underutilized server up and running all the time.
- **Hardware independence.** Virtual machines operate independently from the physical hardware on which they reside. This means that the configuration of virtual machine components (for example, RAM, CPU, network interface) can differ from the underlying components of the physical server. In addition, each virtual machine on the same physical server can run a different operating system.
- **Virtual machine mobility.** A virtual machine (VM) is basically an encapsulation of virtual hardware resources with its own operating system and applications. These resources operate independently from the underlying physical server hardware. As a result, virtual machines can easily move from one location to another without having to make underlying resource changes or disrupt live applications.

Classifying Virtual Machines

To facilitate the discovery and movement of virtual machines, the OmniSwitch provides the Universal Network Profile (UNP) feature and supports the IEEE P802.1Qbg Edge Virtual Bridging (EVB) protocol.

UNP Overview

This section provides an overview of UNP functionality, which is not restricted to supporting only data center solutions. For more information about UNP, see [Chapter 26, “Configuring Universal Network Profiles,”](#) in the *OmniSwitch AOS Release 7 Network Configuration Guide*.

The UNP feature provides the ability to define and apply network access control to specific types of devices by grouping such devices according to specific matching profile criteria. This allows network administrators to create virtual network profiles (vNPs) and user network profiles from a unified framework of operation and administration.

Basically, UNP functionality is used to define profile-based VLANs or SPB service access points (SAPs) to which network devices are assigned. The profile can allow, deny, or require actions by users or machines on the network. Because membership to a VLAN or service is based on UNP profile criteria, devices assigned to the VLAN or service are not tied to a specific port or switch. This flexibility allows device mobility within the network while maintaining network security.

Implementing UNP functionality to create virtual network profiles is particularly useful in an Alcatel-Lucent data center switching architecture. A vNP can define information such as access control rights, a VLAN or SPB service assignment, along with expected quality of service (QoS) levels for data center machines. This type of information can help virtual machines to securely bind to a VLAN bridged domain or an SPB service domain.

Profile Types

The OmniSwitch supports two separate traffic domains: VLAN and service. Traffic is classified into each domain based on the UNP profile type.

- **VLAN profile.** This type of profile creates a VLAN-port association (VPA) for device traffic that is classified into the profile. The VPA represents an association between the UNP port on which the device traffic was received and the VLAN ID specified by the profile. In other words, once classified into a specific profile, device traffic is tagged to forward on the UNP VLAN.
- **Service profile.** The OmniSwitch supports Shortest Path Bridging (SPB) services, which are based on the Provider Backbone Bridge Network (PBBN) architecture. This type of profile creates an association between device traffic that is classified into the profile and a SPB service access point (SAP).

Edge Port Types

Virtual machines bound to either a VLAN domain or a service domain can run on hosts with and without Edge Virtual Bridging (EVB) enabled. This results in three types of data center edge ports, all of which are provided through the OmniSwitch UNP and EVB features:

- UNP-enabled bridge port
- UNP-enabled service access port
- EVB-enabled bridge port

EVB service ports are not supported. See the [“Using EVB” on page 4-8](#) for more information.

UNP also provides the ability to group UNP ports into a single logical customer domain. Once a UNP port is assigned to a specific customer domain ID, only profile classification rules associated with the same domain are applied to that port.

Using customer domains to segment network traffic ties a profile, such as a vNP, to specific UNP ports. For example, UNP ports carrying traffic for Customer A are grouped into domain 2, and profiles tailored for Customer A are also assigned to domain 2. As a result, domain 2 profiles are applied only to domain 2 UNP ports.

Classification Rules

A vNP can also define specific classification rules that are used to assign virtual machines to a profile. The following classification rules (shown in the order of precedence) are used to classify traffic received on UNP bridge and access ports:

- 1 MAC address + VLAN tag
- 2 MAC address
- 3 MAC address range + VLAN tag
- 4 MAC address range
- 5 IP address + VLAN tag
- 6 IP address
- 7 VLAN tag

In addition to classification rules, UNP provides the ability to trust the VLAN tag of incoming traffic. This means that virtual machines can be provisioned into the VLAN tag pre-assigned by the virtualized server hypervisor. Each UNP port can then be assigned a default UNP to classify untagged traffic when all classification mechanisms either fail or are not available.

How it Works

Dynamic assignment of devices using UNPs is achieved through port-based functionality that provides the ability to authenticate and classify device traffic. Authentication verifies the device identity and provides a UNP name. In the event authentication is not available or is unsuccessful, classification rules associated with the UNPs, as well as the UNP port configuration attributes, are applied to the traffic to determine the UNP assignment.

There is no global switch setting to invoke UNP (vNP) functionality. Instead, switch ports are configured as a UNP bridge port or a UNP access port and profiles are defined to determine the dynamic VLAN or SPB service assignment for devices connected through the UNP ports.

When UNP is enabled on a switch port or link aggregate, the following device classification process is triggered when the port receives traffic:

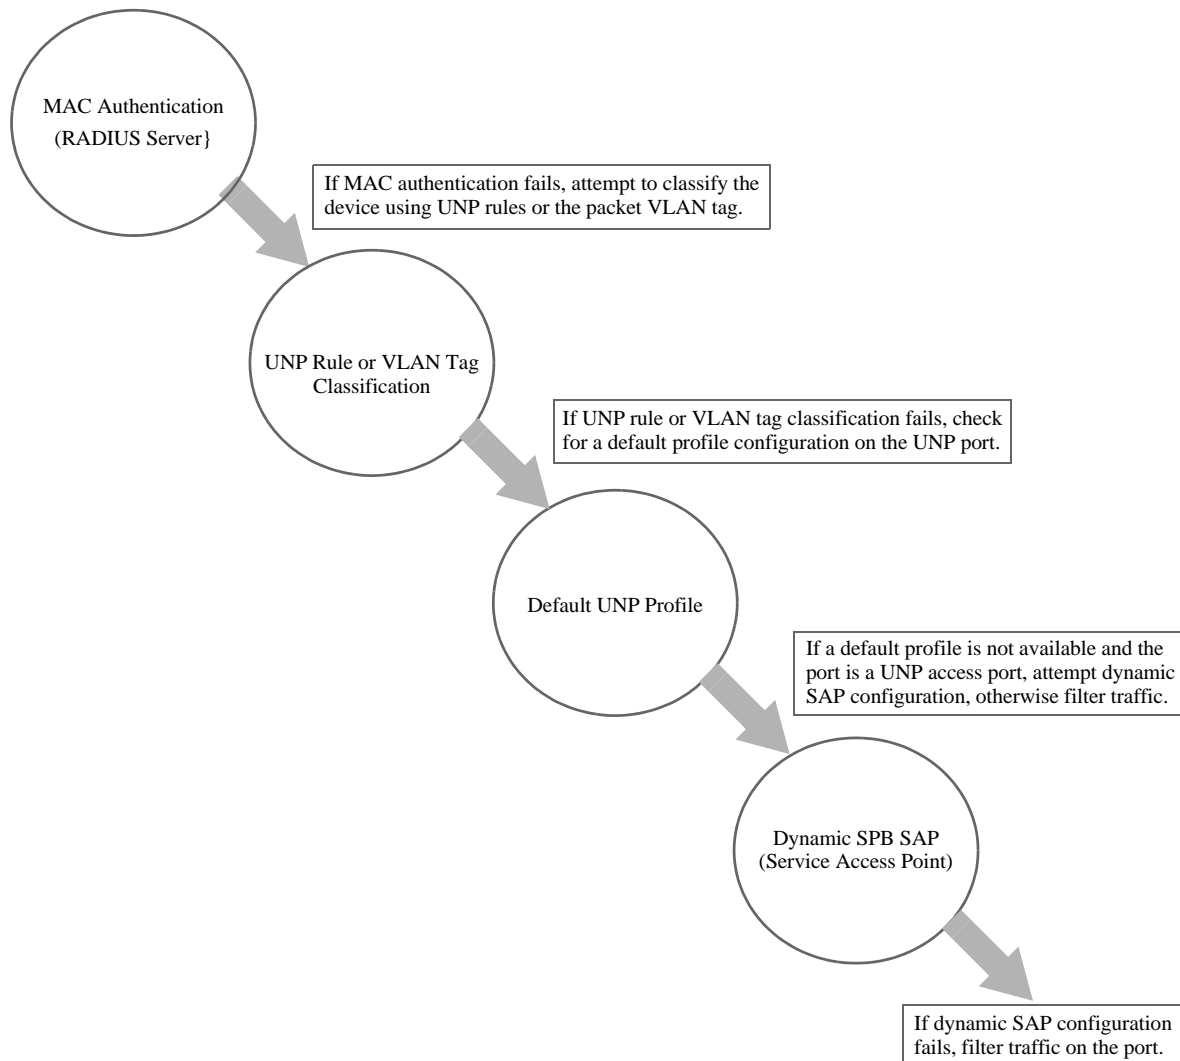


Figure 1: UNP Classification Overview

Virtual Network Profile Example

The following simplified illustration shows an example of vNP functionality in a virtualized network:

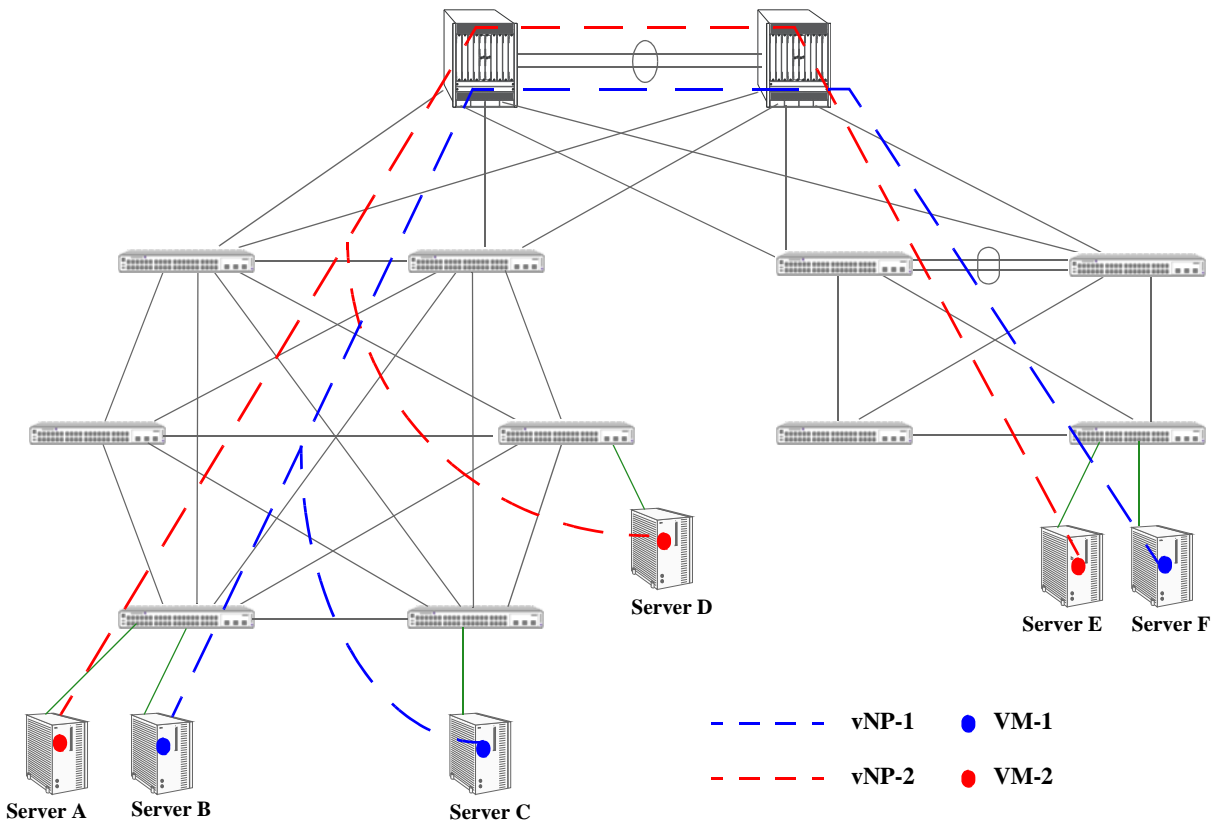


Figure 2: Virtual Network Profile–VM Mobility

In this example,

- On each switch that is connected to a server, virtual network profiles vNP-1 and vNP-2 are configured to classify and forward traffic for virtual machines VM-1 and VM-2.
- The virtual network profiles specify either a VLAN ID or service ID that will carry virtual machine traffic through the data center mesh to other servers (either within the same mesh or through the cloud to another mesh) and if any QoS policies are applied to virtual machine traffic.
- Each port connected to a server is enabled with UNP functionality.
- When the MAC address for VM-1 and for VM-2 is discovered on a UNP port, the switch applies the profile configuration to each MAC address to classify each virtual machine into a VLAN or SPB service. As a result,
 - VM-1 is assigned to the vNP-1 profile. The profile attributes then determine the VLAN or service assignment and if any QoS policies are applied to VM-1 traffic.
 - VM-2 is assigned to the vNP-2 profile. The profile attributes then determine the VLAN or service assignment and if any QoS policies are applied to VM-2 traffic.
- If VM-1 or VM-2 is moved to another server, as shown in this sample network, VM-1 will get assigned to vNP-1 and VM-2 will get assigned to vNP-2, thus the network access control configuration for each virtual machine is preserved whenever it moves to another server.

Using EVB

In the data center environment, data center servers run specialized hypervisor software to provision and manage multiple virtual machines (VMs) within the host server. These VMs can be dynamically created, deleted, or even moved to other host servers in the network.

Each VM may require different network connections and services. Typically, similar types of VMs are grouped together to form a VM network to control the unicast and multicast domains and to make a connection to storage area networks (SAN) or wide area networks (WAN).

The OmniSwitch implementation of Edge Virtual Bridging (EVB) helps automate the discovery of virtual machines and connect them to the appropriate network domain. EVB runs between a station (data center server) and a bridge (an OmniSwitch).

EVB defines logical components of the station (server): Virtual Station Interface (VSI), Edge Relay (ER), and Uplink Relay Port. The OmniSwitch supports a single ER per switch port. The ER can operate as a Virtual Ethernet Port Aggregator (VEPA) or Virtual Ethernet Bridge (VEB).

- VEPA (Virtual Ethernet Port Aggregator)—the hypervisor treats the NIC as a single interface connected to each VM and all outgoing traffic is sent through the NIC to an external switch (the Edge Relay function per IEEE 802.1Qbg in the host server).
- VEB (Virtual Ethernet Bridge)—the hypervisor creates a virtual Ethernet switch inside the host that can bridge data between VMs and send data out to the network as needed.




Figure 3: EVB Architecture

Only the EVB VLAN bridging mode is supported at this time. In this mode,

- The VM requests specific C-VLAN(s). Note that the VM must provide the requested C-VLANs when it sends EVB protocol packets to the switch.
- The switch dynamically creates the corresponding VLANs and propagates those VLANs to the other switches in the network using the MVRP protocol.

- After the VM is associated to the bridge using the EVB protocol, then the VM can send data traffic tagged with the requested C-VLAN.

UNP and EVB are mutually exclusive on the same port; only UNP or EVB functionality is allowed on a specific port or link aggregate at any given time.

- If a port is configured as an EVB bridging port, then the EVB protocol is used between the station and bridge to discover VMs.
- If a port is configured as a UNP bridging or access port, then the regular UNP classification methods are used to assign the VM traffic to a dynamic VLAN or dynamic SAP.
- The UNP method is mainly used for stations that are not using EVB.
- The EVB method is only for interaction with EVB stations.

Tracking Virtual Machines

Virtual hypervisor managers can perform virtual machine (VM) movement without intervention. When this type of movement occurs, the hypervisor manager indicates to which hypervisor (virtual server) the virtual machine was moved. However, it is not easy to tell which OmniSwitch port in the data center mesh the virtual machine is now using to access the network.

Alcatel-Lucent network management tools provide management and visibility for virtual machines in the mesh. The suite of management tools includes:

- OmniVista 2500 Virtual Machine Management (VMM) for management of VM mobility and integration with standard hypervisors.
- OmniVista 2500 NMS for mesh provisioning and management.

OmniVista 2500 provides a single pane of management for VMs across the network, which includes VM visibility and their point of association to the network. This gives network administrators a dashboard to determine and track the following information:

- Virtual machine locations.
- Virtual machine types.
- The switch ports to which the virtual machines are connected.
- The duration of the connections.
- Which virtual network profile (vNP) the virtual machine is using.

VMM is an embedded application within OmniVista that consists of the following components:

- **vCenter integration:** VM discovery, VM polling and event listener service will listen to preference changes and event notifications.
- **vNP configuration:** Profiles of UNP configuration that specify the configuration associated with each VM VLAN. The profiles can be assigned to switches, as and when needed.
- **VLAN Notification:** Displays a list of VM instances where the network is missing some configuration to effectively handle VMs attached. This allows the administrator to be proactive in identifying and correcting any necessary network configuration.
- **OmniVista Locator:** Provides the ability to search for specific VMs using various search criteria.

A Software License and Copyright Statements

This appendix contains Alcatel-Lucent and third-party software vendor license and copyright statements.

Alcatel-Lucent License Agreement

ALCATEL-LUCENT SOFTWARE LICENSE AGREEMENT

IMPORTANT. Please read the terms and conditions of this license agreement carefully before opening this package.

By opening this package, you accept and agree to the terms of this license agreement. If you are not willing to be bound by the terms of this license agreement, do not open this package. Please promptly return the product and any materials in unopened form to the place where you obtained it for a full refund.

1. **License Grant.** This is a license, not a sales agreement, between you (the “Licensee”) and Alcatel-Lucent. Alcatel-Lucent hereby grants to Licensee, and Licensee accepts, a non-exclusive license to use program media and computer software contained therein (the “Licensed Files”) and the accompanying user documentation (collectively the “Licensed Materials”), only as authorized in this License Agreement. Licensee, subject to the terms of this License Agreement, may use one copy of the Licensed Files on the Licensee’s system. Licensee agrees not to assign, sublicense, transfer, pledge, lease, rent, or share their rights under this License Agreement. Licensee may retain the program media for backup purposes with retention of the copyright and other proprietary notices. Except as authorized under this paragraph, no copies of the Licensed Materials or any portions thereof may be made by Licensee and Licensee shall not modify, decompile, disassemble, reverse engineer, or otherwise attempt to derive the Source Code. Licensee is also advised that Alcatel-Lucent products contain embedded software known as firmware which resides in silicon. Licensee may not copy the firmware or transfer the firmware to another medium.

2. **Alcatel-Lucent’s Rights.** Licensee acknowledges and agrees that the Licensed Materials are the sole property of Alcatel-Lucent and its licensors (herein “its licensors”), protected by U.S. copyright law, trademark law, and are licensed on a right to use basis. Licensee further acknowledges and agrees that all rights, title, and interest in and to the Licensed Materials are and shall remain with Alcatel-Lucent and its licensors and that no such right, license, or interest shall be asserted with respect to such copyrights and trademarks. This License Agreement does not convey to Licensee an interest in or to the Licensed Materials, but only a limited right to use revocable in accordance with the terms of this License Agreement.

3. Confidentiality. Alcatel-Lucent considers the Licensed Files to contain valuable trade secrets of Alcatel-Lucent, the unauthorized disclosure of which could cause irreparable harm to Alcatel-Lucent. Except as expressly set forth herein, Licensee agrees to use reasonable efforts not to disclose the Licensed Files to any third party and not to use the Licensed Files other than for the purpose authorized by this License Agreement. This confidentiality obligation shall continue after any termination of this License Agreement.

4. Indemnity. Licensee agrees to indemnify, defend and hold Alcatel-Lucent harmless from any claim, lawsuit, legal proceeding, settlement or judgment (including without limitation Alcatel-Lucent's reasonable United States and local attorneys' and expert witnesses' fees and costs) arising out of or in connection with the unauthorized copying, marketing, performance or distribution of the Licensed Files.

5. Limited Warranty. Alcatel-Lucent warrants, for Licensee's benefit alone, that the program media shall, for a period of ninety (90) days from the date of commencement of this License Agreement (referred to as the Warranty Period), be free from defects in material and workmanship. Alcatel-Lucent further warrants, for Licensee benefit alone, that during the Warranty Period the Licensed Files shall operate substantially in accordance with the functional specifications in the User Guide. If during the Warranty Period, a defect in the Licensed Files appears, Licensee may return the Licensed Files to Alcatel-Lucent for either replacement or, if so elected by Alcatel-Lucent, refund of amounts paid by Licensee under this License Agreement. EXCEPT FOR THE WARRANTIES SET FORTH ABOVE, THE LICENSED MATERIALS ARE LICENSED "AS IS" AND ALCATEL-LUCENT AND ITS LICENSORS DISCLAIM ANY AND ALL OTHER WARRANTIES, WHETHER EXPRESS OR IMPLIED, INCLUDING (WITHOUT LIMITATION) ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. SOME STATES DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES SO THE ABOVE EXCLUSIONS MAY NOT APPLY TO LICENSEE. THIS WARRANTY GIVES THE LICENSEE SPECIFIC LEGAL RIGHTS. LICENSEE MAY ALSO HAVE OTHER RIGHTS WHICH VARY FROM STATE TO STATE.

6. Limitation of Liability. Alcatel-Lucent's cumulative liability to Licensee or any other party for any loss or damages resulting from any claims, demands, or actions arising out of or relating to this License Agreement shall not exceed the license fee paid to Alcatel-Lucent for the Licensed Materials. IN NO EVENT SHALL ALCATEL-LUCENT BE LIABLE FOR ANY INDIRECT, INCIDENTAL, CONSEQUENTIAL, SPECIAL, OR EXEMPLARY DAMAGES OR LOST PROFITS, EVEN IF ALCATEL-LUCENT HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. SOME STATES DO NOT ALLOW THE LIMITATION OR EXCLUSION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE LIMITATION OR EXCLUSION TO INCIDENTAL OR CONSEQUENTIAL DAMAGES MAY NOT APPLY TO LICENSEE.

7. Export Control. This product is subject to the jurisdiction of the United States. Licensee may not export or reexport the Licensed Files, without complying with all United States export laws and regulations, including but not limited to (i) obtaining prior authorization from the U.S. Department of Commerce if a validated export license is required, and (ii) obtaining "written assurances" from licensees, if required.

8. Support and Maintenance. Except as may be provided in a separate agreement between Alcatel-Lucent and Licensee, if any, Alcatel-Lucent is under no obligation to maintain or support the copies of the Licensed Files made and distributed hereunder and Alcatel-Lucent has no obligation to furnish Licensee with any further assistance, documentation or information of any nature or kind.

9. Term. This License Agreement is effective upon Licensee opening this package and shall continue until terminated. Licensee may terminate this License Agreement at any time by returning the Licensed Materials and all copies thereof and extracts therefrom to Alcatel-Lucent and certifying to Alcatel-Lucent in writing that all Licensed Materials and all copies thereof and extracts therefrom have been returned or erased by the memory of Licensee's computer or made non-readable. Alcatel-Lucent may terminate this License Agreement upon the breach by Licensee of any term hereof. Upon such termination by

Alcatel-Lucent, Licensee agrees to return to Alcatel-Lucent or destroy the Licensed Materials and all copies and portions thereof.

10. Governing Law. This License Agreement shall be construed and governed in accordance with the laws of the State of California.

11. Severability. Should any term of this License Agreement be declared void or unenforceable by any court of competent jurisdiction, such declaration shall have no effect on the remaining terms herein.

12. No Waiver. The failure of either party to enforce any rights granted hereunder or to take action against the other party in the event of any breach hereunder shall not be deemed a waiver by that party as to subsequent enforcement of rights or subsequent actions in the event of future breaches.

13. Notes to United States Government Users. Software and documentation are provided with restricted rights. Use, duplication or disclosure by the government is subject to (i) restrictions set forth in GSA ADP Schedule Contract with Alcatel-Lucent's reseller(s), or (ii) restrictions set forth in subparagraph (c) (1) and (2) of 48 CFR 52.227-19, as applicable.

14. Third Party Materials. Licensee is notified that the Licensed Files contain third party software and materials licensed to Alcatel-Lucent by certain third party licensors. Some third party licensors are third party beneficiaries to this License Agreement with full rights of enforcement. Please refer to the section entitled "[Third Party Licenses and Notices](#)" on page -4 for the third party license and notice terms.

Third Party Licenses and Notices

Legal Notices applicable to any software distributed alone or in connection with the product to which this document pertains, are contained in files within the software itself located at: **/flash/foss**.

Also, if needed, we provide all FOSS (Free and Open Source Software) source code used into this release at the following URL: <https://service.esd.alcatel-lucent.com/portal/page/portal/EService/release>

Index

B

backbone VLAN 4-15, 4-23
BVLAN *see* backbone VLAN

D

Data Center Bridging 3-1
 Interaction with Other Features 3-13
 Profiles 3-15
Data Center Bridging Exchange 3-1, 3-10
DCB *see* Data Center Bridging
DCBx *see* Data Center Bridging Exchange

E

Edge Virtual Bridging 5-1
Enhanced Transmission Selection 3-1, 3-7
ETS *see* Enhanced Transmission Selection
EVB *see* Edge Virtual Bridging

I

ISIS-SPB 4-9

P

PBB *see* Provider Backbone Bridge
PFC *see* Priority-based Flow Control
Priority-based Flow Control 3-1, 3-5
Provider Backbone Bridge
 802.1AH 4-6
 network 4-6
 SPB services 4-12

Q

qos qsi qsp dcb command 3-22
qos qsi qsp dcb dcbx admin-state command 3-23
qos qsi qsp dcb dcbx ets command 3-23
qos qsi qsp dcb dcbx pfc command 3-23
qos qsp dcb import command 3-21
qos qsp dcb tc command 3-21

S

service access l2profile command 4-37
service access vlan-xlation command 4-33, 4-36
service l2profile command 4-36
service spb admin-state command 4-33
service spb command 4-32
service spb sap admin-state command 4-40
service spb sap command 4-38
service spb sap trusted command 4-39

service spb vlan-xlation command 4-33

Shortest Path Bridging 4-1, 4-6

 benefits 4-1
 bridge ID 4-27
 bridge priority 4-27
 BVLAN 4-23
 ISIS-SPB 4-9
 services 4-12
 shortest path trees 4-8
 SPT 4-8
 system ID 4-27
 topology example 4-13

show qos qsi dcbx command 3-23

show qos qsp dcb command 3-22

spb isis admin-state command 4-31

spb isis area-address command 4-27

spb isis bridge-priority command 4-27

spb isis lsp-wait command 4-29

spb isis overload command 4-30

spb isis overload-on-boot command 4-30

spb isis source-id command 4-28

spb isis spf-wait command 4-28

SPB *see* Shortest Path Bridging

U

Universal Network Profile
 Profile Types 5-4
 Virtual Network Profile 5-1

V

Virtual Machines
 Classifying 5-4
 Tracking 5-9

