

ExtremeSwitching™

Configuring VXLAN Gateway on VSP Operating System Software

Release 7.0 (VOSS)
9035358 Rev.01
April 2018

© 2018, Extreme Networks, Inc.
All Rights Reserved.

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Extreme Networks, Inc. assumes no liability for any errors. Extreme Networks, Inc. reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

"Documentation" means information published in varying mediums which may include product information, operating instructions and performance specifications that are generally made available to users of products. Documentation does not include marketing materials. Extreme Networks shall not be responsible for any modifications, additions, or deletions to the original published version of Documentation unless such modifications, additions, or deletions were performed by or on the express behalf of Extreme Networks. End User agrees to indemnify and hold harmless Extreme Networks, Extreme Networks' agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Extreme Networks is not responsible for the contents or reliability of any linked websites referenced within this site or Documentation provided by Extreme Networks. Extreme Networks is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Extreme Networks does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Extreme Networks provides a limited warranty on Extreme Networks hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Extreme Networks' standard warranty language, as well as information regarding support for this product while under warranty is available to Extreme Networks customers and other parties through the Extreme Networks Support website: <http://www.extremenetworks.com/support> under the link "Policies" or such successor site as designated by Extreme Networks. Please note that if You acquired the product(s) from an authorized Extreme Networks Channel Partner outside of the United States and Canada, the warranty is provided to You by said Extreme Networks Channel Partner and not by Extreme Networks.

"Hosted Service" means an Extreme Networks hosted service subscription that You acquire from either Extreme Networks or an authorized Extreme Networks Channel Partner (as applicable) and which is described further in Hosted SAS or other service description documentation regarding the applicable hosted service. If You purchase a Hosted Service subscription, the foregoing limited warranty may not apply but You may be entitled to support services in connection with the Hosted Service as described further in your service description documents for the applicable Hosted Service. Contact Extreme Networks or Extreme Networks Channel Partner (as applicable) for more information.

Hosted Service

THE FOLLOWING APPLIES ONLY IF YOU PURCHASE AN EXTREME NETWORKS HOSTED SERVICE SUBSCRIPTION FROM EXTREME NETWORKS OR AN EXTREME NETWORKS CHANNEL PARTNER (AS APPLICABLE), THE TERMS OF USE FOR HOSTED SERVICES ARE AVAILABLE ON THE EXTREME NETWORKS WEBSITE, <https://extremeportal.force.com> OR SUCH SUCCESSOR SITE AS DESIGNATED BY EXTREME NETWORKS, AND ARE APPLICABLE TO ANYONE WHO ACCESSES OR USES THE HOSTED SERVICE. BY ACCESSING OR USING THE HOSTED SERVICE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE DOING SO (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THE TERMS OF USE. IF YOU ARE ACCEPTING THE TERMS OF USE ON BEHALF A COMPANY OR OTHER LEGAL ENTITY, YOU

REPRESENT THAT YOU HAVE THE AUTHORITY TO BIND SUCH ENTITY TO THESE TERMS OF USE. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO ACCEPT THESE TERMS OF USE, YOU MUST NOT ACCESS OR USE THE HOSTED SERVICE OR AUTHORIZE ANYONE TO ACCESS OR USE THE HOSTED SERVICE.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE EXTREME NETWORKS WEBSITE, <https://extremeportal.force.com> OR SUCH SUCCESSOR SITE AS DESIGNATED BY EXTREME NETWORKS, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS EXTREME NETWORKS SOFTWARE, PURCHASED FROM EXTREME NETWORKS, INC., ANY EXTREME NETWORKS AFFILIATE, OR AN EXTREME NETWORKS CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH EXTREME NETWORKS OR AN EXTREME NETWORKS CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY EXTREME NETWORKS IN WRITING, EXTREME NETWORKS DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN EXTREME NETWORKS, AN EXTREME NETWORKS AFFILIATE OR AN EXTREME NETWORKS CHANNEL PARTNER; EXTREME NETWORKS RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND EXTREME NETWORKS, INC. OR THE APPLICABLE EXTREME NETWORKS AFFILIATE ("EXTREME NETWORKS").

Extreme Networks grants You a license within the scope of the license types described below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License as set forth below in the Designated System(s) License (DS) section as applicable. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. "Software" means computer programs in object code, provided by Extreme Networks or an Extreme Networks Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. "Designated Processor" means a single stand-alone computing device. "Server" means a set of Designated Processors that hosts (physically or virtually) a software application to be accessed by multiple users. "Instance" means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine ("VM") or similar deployment.

License type(s)

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only: 1) on a number of Designated Processors up to the number indicated in the order; or 2) up to the number of Instances of the Software as indicated in the order, Documentation, or as authorized by Extreme Networks in writing. Extreme Networks may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Extreme Networks through electronic means established by Extreme Networks specifically for this purpose.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Extreme Networks. All content on this site, the documentation, Hosted Service, and the product provided by Extreme Networks including the selection, arrangement and design of the content is owned either by Extreme Networks or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part,

including any code and software unless expressly authorized by Extreme Networks. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Extreme Networks can be a criminal, as well as a civil offense under the applicable law.

Virtualization

The following applies if the product is deployed on a virtual machine. Each product has its own ordering code and license types. Note, unless otherwise stated, that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Extreme Networks Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

Third Party Components

"Third Party Components" mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements ("Third Party Components"), which contain terms regarding the rights to use certain portions of the Software ("Third Party Terms"). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Extreme Networks' website at: <http://www.extremenetworks.com/support/policies/software-licensing> or such successor site as designated by Extreme Networks. The open source software license terms provided as Third Party Terms are consistent with the license rights granted in these Software License Terms, and may contain additional rights benefiting You, such as modification and distribution of the open source software. The Third Party Terms shall take precedence over these Software License Terms, solely with respect to the applicable Third Party Components to the extent that these Software License Terms impose greater restrictions on You than the applicable Third Party Terms.

The following applies only if the H.264 (AVC) codec is distributed with the product. THIS PRODUCT IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE [HTTP://WWW.MPEGLA.COM](http://WWW.MPEGLA.COM).

Service Provider

THE FOLLOWING APPLIES TO EXTREME NETWORKS CHANNEL PARTNER'S HOSTING OF EXTREME NETWORKS PRODUCTS OR SERVICES. THE PRODUCT OR HOSTED SERVICE MAY USE THIRD PARTY COMPONENTS SUBJECT TO THIRD PARTY TERMS AND REQUIRE A SERVICE PROVIDER TO BE INDEPENDENTLY LICENSED DIRECTLY FROM THE THIRD PARTY SUPPLIER. AN EXTREME NETWORKS CHANNEL PARTNER'S HOSTING OF EXTREME NETWORKS PRODUCTS MUST BE AUTHORIZED IN WRITING BY EXTREME NETWORKS AND IF THOSE HOSTED PRODUCTS USE OR EMBED CERTAIN THIRD PARTY SOFTWARE, INCLUDING BUT NOT LIMITED TO MICROSOFT SOFTWARE OR CODECS, THE EXTREME NETWORKS CHANNEL PARTNER IS REQUIRED TO INDEPENDENTLY OBTAIN ANY APPLICABLE LICENSE AGREEMENTS, AT THE EXTREME NETWORKS CHANNEL PARTNER'S EXPENSE, DIRECTLY FROM THE APPLICABLE THIRD PARTY SUPPLIER.

WITH RESPECT TO CODECS, IF THE EXTREME NETWORKS CHANNEL PARTNER IS HOSTING ANY PRODUCTS THAT USE OR EMBED THE G.729 CODEC, H.264 CODEC, OR H.265 CODEC, THE EXTREME NETWORKS CHANNEL PARTNER ACKNOWLEDGES AND AGREES THE EXTREME NETWORKS CHANNEL PARTNER IS RESPONSIBLE FOR ANY AND ALL RELATED FEES AND/OR ROYALTIES. THE G.729 CODEC IS LICENSED BY SIPRO LAB TELECOM INC. SEE WWW.SIPRO.COM/CONTACT.HTML. THE H.264 (AVC) CODEC IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR

THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO: (I) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (II) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION FOR H.264 (AVC) AND H.265 (HEVC) CODECS MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE [HTTP://WWW.MPEGLA.COM](http://WWW.MPEGLA.COM).

Compliance with Laws

You acknowledge and agree that it is Your responsibility for complying with any applicable laws and regulations, including, but not limited to laws and regulations related to call recording, data privacy, intellectual property, trade secret, fraud, and music performance rights, in the country or territory where the Extreme Networks product is used.

Preventing Toll Fraud

"Toll Fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Security Vulnerabilities

Information about Extreme Networks' security support policies can be found in the Global Technical Assistance Center Knowledgebase at <https://gtacknowledge.extremenetworks.com/>.

Downloading Documentation

For the most current versions of Documentation, see the Extreme Networks Support website: <http://documentation.extremenetworks.com>, or such successor site as designated by Extreme Networks.

Contact Extreme Networks Support

See the Extreme Networks Support website: <http://www.extremenetworks.com/support> for product or Hosted Service notices and articles, or to report a problem with your Extreme Networks product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Extreme Networks Support website: <http://www.extremenetworks.com/support/contact/> (or such successor site as designated by Extreme Networks), scroll to the bottom of the page, and select Contact Extreme Networks Support.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, Hosted Service(s), and product(s) provided by Extreme Networks are the registered or unregistered Marks of Extreme Networks, Inc., its affiliates, its licensors, its suppliers, or other third parties. Users are not permitted to use such Marks without prior written consent from Extreme Networks or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Extreme Networks or the applicable third party.

Extreme Networks is a registered trademark of Extreme Networks, Inc.

All non-Extreme Networks trademarks are the property of their respective owners. Linux[®] is the registered trademark of Linus Torvalds in the U.S. and other countries.

For additional information on Extreme Networks trademarks, please see: <http://www.extremenetworks.com/company/legal/>

Contents

Chapter 1: Preface	6
Purpose.....	6
Training.....	6
Providing Feedback to Us.....	6
Getting Help.....	6
Extreme Networks Documentation.....	7
Subscribing to Service Notifications.....	8
Chapter 2: New in this document	9
Introducing OVSDB protocol support for VXLAN Gateway.....	9
Notice about feature support.....	9
Chapter 3: SPBM and IS-IS configuration workflow	11
Chapter 4: VXLAN Gateway fundamentals	13
VXLAN Gateway licensing.....	16
VXLAN modes	16
Types of VXLAN Gateway deployments	17
VXLAN-to-VLAN Layer 2 gateway deployment.....	17
VXLAN-to-VXLAN Layer 3 gateway deployment.....	18
VXLAN-to-SPB gateway deployment.....	19
VXLAN Gateway considerations and limitations.....	20
OVSDB protocol support for VXLAN Gateway considerations and limitations.....	22
Chapter 5: VXLAN Gateway configuration using the CLI	25
Configuring the VXLAN Gateway boot flag.....	25
Configuring VXLAN Gateway.....	26
Flushing the MAC forwarding table.....	28
VXLAN Gateway show commands.....	29
Displaying VTEP source information.....	29
Displaying remote VTEP information.....	29
Displaying the name of the remote VTEP.....	30
Displaying VNID information.....	30
Displaying the MAC addresses in the VNID FDB.....	31
Displaying the VXLAN running configuration.....	31
Chapter 6: OVSDB protocol support for VXLAN Gateway configuration using CLI	33
Configuring OVSDB managed interfaces.....	33
Configuring OVSDB protocol support for VXLAN Gateway.....	34
Configuring OVSDB replication.....	36
Displaying the OVSDB configuration.....	37
Displaying the OVSDB controller status.....	38
Displaying the OVSDB replication state.....	38
Displaying the OVSDB managed interfaces.....	39

Chapter 7: VXLAN Gateway configuration using EDM	40
Configuring the VXLAN Gateway boot flag.....	40
Configuring a VTEP source IP address.....	40
Configuring a remote VTEP.....	41
Configuring a VNID.....	42
Configuring VNID endpoints.....	43
Configuring ELAN endpoints.....	44
Displaying the VTEP next hop.....	45
Displaying the VNID forwarding database.....	45
Chapter 8: OVSDB protocol support for VXLAN Gateway configuration using EDM	47
Configuring OVSDB globally.....	47
Configuring an OVSDB controller.....	48
Displaying the OVSDB controller status.....	49
Configuring OVSDB managed interface.....	50
Chapter 9: VXLAN Gateway configuration example	51

Chapter 1: Preface

Purpose

This document provides information on the VXLAN Gateway feature for the following products:

- Extreme Networks VSP 7200 Series
- Extreme Networks VSP 8000 Series

This document contains procedural and conceptual information to help you configure and manage VXLAN Gateway on support platforms. Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Training

Ongoing product training is available. For more information or to register, you can access the Web site at www.extremenetworks.com/education/.

Providing Feedback to Us

We are always striving to improve our documentation and help you work better, so we want to hear from you! We welcome all feedback but especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team about this document, please contact us using our short [online feedback form](#). You can also email us directly at internalinfodev@extremenetworks.com

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

- [GTAC \(Global Technical Assistance Center\) for Immediate Support](#)
 - Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact
 - Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- [GTAC Knowledge](#) – Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- [The Hub](#) – A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- [Support Portal](#) – Manage cases, downloads, service contracts, product licensing, and training and certifications.

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Extreme Networks Documentation

To find Extreme Networks product guides, visit our documentation pages at:

Current Product Documentation	www.extremenetworks.com/documentation/
Archived Documentation (for previous versions and legacy products)	www.extremenetworks.com/support/documentation-archives/
Release Notes	www.extremenetworks.com/support/release-notes

Open Source Declarations

Some software files have been licensed under certain open source licenses. More information is available at: www.extremenetworks.com/support/policies/software-licensing.

Subscribing to Service Notifications

Subscribe to receive an email notification for product and software release announcements, Vulnerability Notices, and Service Notifications.

About this task

You can modify your product selections at any time.

Procedure

1. In an Internet browser, go to <http://www.extremenetworks.com/support/service-notification-form/>.
2. Type your first and last name.
3. Type the name of your company.
4. Type your email address.
5. Type your job title.
6. Select the industry in which your company operates.
7. Confirm your geographic information is correct.
8. Select the products for which you would like to receive notifications.
9. Click **Submit**.

Chapter 2: New in this document

The following sections describe what is new in this document.

Introducing OVSDB protocol support for VXLAN Gateway

DEMO FEATURE — OVSDB protocol support for VXLAN Gateway

You can now configure and manage the VXLAN Gateway hardware-based VTEP functions with Open vSwitch Database Management Protocol (OVSDB). You must have at least one Network Virtualization Controller (NVC) configured as an Open vSwitch (OVS) to use the OVSDB features. The VXLAN Gateway must have OVSDB enabled and the Network Virtualization Controller (NVC) must communicate on the OVSDB managed interface. The NVC can manage the VNID, VNID to I-SID bindings, Remote-VTEP, and VNI to Remote-VTEP associations. You only need to manually configure the Source-VTEP and NVC IP addresses. The NVC distributes the hosts MAC and IP addresses learned by the VXLAN Gateway.

Important:

OVSDB protocol support for VXLAN Gateway is a demo feature. Do not use the OVSDB features in production environments.

The VXLAN Gateway on the switch terminates VXLAN tunnels and is known as a VXLAN tunnel endpoint (VTEP). The switch VXLAN Gateway is a hardware-based VTEP that allows VXLAN to communicate with VLANs, other VXLANs, as well as Fabric Connect I-SIDs.

Note:

VXLAN Gateway and associated features require a Premier License.

Notice about feature support

This document includes content for multiple hardware platforms across different software releases. As a result, the content can include features not supported by your hardware in the current software release.

If a documented command, parameter, tab, or field does not appear on your hardware, it is not supported.

For information about feature support, see *Release Notes*.

New in this document

For information about physical hardware restrictions, see your hardware documentation.

Chapter 3: SPBM and IS-IS configuration workflow

The following section describes the generic work flow to configure SPBM and IS-IS infrastructure and services on your network.

*** Note:**

This section is an overview. For further details on the SPBM and IS-IS infrastructure and configuration, see the documents described in the Documentation sources section below.

1. Infrastructure configuration:

As a first step, you must configure your basic infrastructure for Shortest Path Bridging MAC (SPBM).

2. Services configuration:

After you complete the infrastructure configuration, you configure the appropriate services for your network to run on top of your base architecture. This includes:

- Layer 2 and Layer 3 VSNs
- IP Shortcuts
- Inter-VSN routing

3. Operations and Management:

To debug connectivity issues and isolate network faults in the SPBM network, you can use Connectivity Fault Management (CFM).

Documentation sources:

- For information on basic SPBM infrastructure and IS-IS configuration and Layer 2 services, see *Configuring Fabric Basics and Layer 2 Services*.

This document also contains information on configuring Fabric Extend, which enables your enterprise to extend Fabric Connect technology over Layer 2 or Layer 3 core networks.

- For information on Fabric Layer 3 services configuration, see *Configuring Fabric Layer 3 Services*.
- For information on IP Multicast over Fabric Connect configuration and services, see *Configuring Fabric Multicast Services*. This document also contains information about configuring the SPB-PIM Gateway (SPB-PIM GW), which provides multicast inter-domain communication between an SPB network and a PIM network. The SPB-PIM GW can also connect two independent SPB domains.
- For information on CFM, see *Troubleshooting*.

- For information on VXLAN Gateway configuration, see *Configuring VXLAN Gateway*.

Chapter 4: VXLAN Gateway fundamentals

VXLAN Gateway terminates virtual extensible LAN (VXLAN) tunnels that stretch emulated Layer 2 segments over an IP network.

VXLAN is an encapsulation protocol for running an overlay network over an existing Layer 3 infrastructure. This tunneling scheme uses a VLAN-like encapsulation mechanism to encapsulate MAC-based traffic. The VXLAN overlays separate workloads from physical networks, which facilitates the movement of VMs in a multi-tenant environment.

For more information about VXLAN, see the Internet Engineering Task Force (IETF) standard RFC 7348.

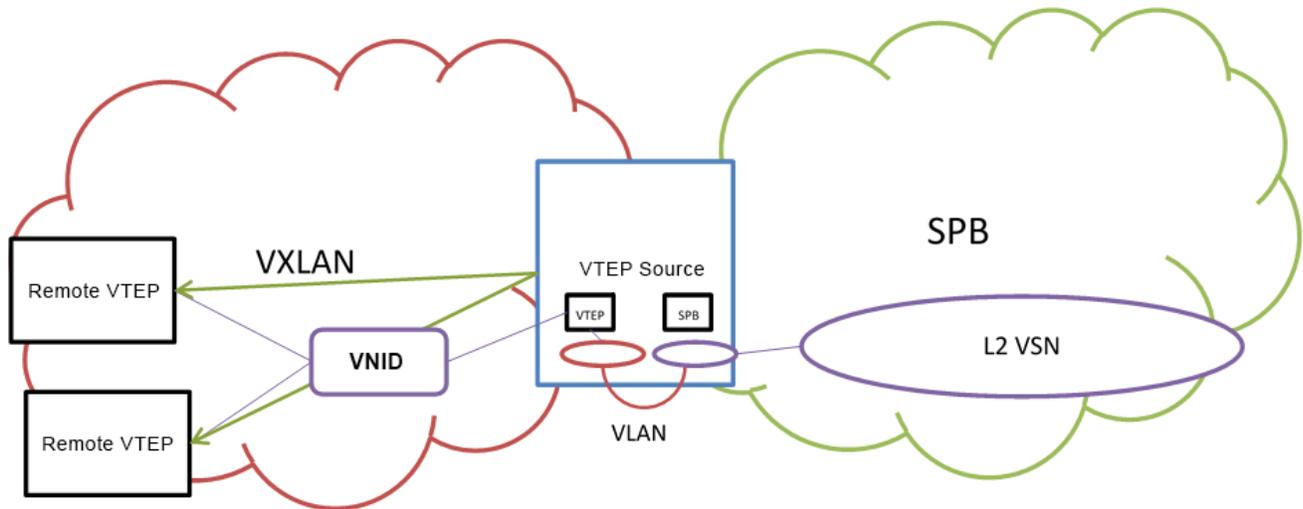
 **Note:**

VXLAN Gateway is supported on the VSP 7200 and VSP 8000 Series switches.

VXLAN Gateway as a VTEP

The VXLAN Gateway feature terminates VXLAN tunnels and operates as a VXLAN tunnel endpoint (VTEP). The VXLAN Gateway is a hardware-based VTEP that allows a VXLAN to communicate with VLANs, other VXLANs, as well as Fabric Connect I-SIDs.

The following figure shows a VXLAN Gateway (labeled as VTEP Source) connecting VXLAN segments to an SPB network. This type of deployment provides a solution for VXLAN in data centers to interoperate with Fabric Connect networks (labeled as SPB).



VXLAN segments

VXLAN tunnels are also called VXLAN segments. VXLAN segments provide the ability to separate, abstract, and decouple the physical topology from a *logical* or *virtual* topology by using encapsulated tunneling. VXLAN segments allows only VMs within the same VXLAN tunnel to communicate with each other.

VNID

Each VXLAN segment has a 24-bit segment ID called a VXLAN Network Identifier (VNID). VNID allows up to *16 million* VXLAN segments to coexist within the same administrative domain, which enables VMs to migrate between servers in the same data center or between distributed data centers. Each VTEP can support multiple VNIDs. For information about maximum scaling numbers, see *Release Notes*.

The VNID uses the inner MAC frame originated by the individual VM. Overlapping MAC addresses can cross VXLAN segments, but never have traffic crossover because the traffic is isolated using the VNID qualifier. Due to this encapsulation, only the VTEP knows the VNID and its associated VXLAN tunnel. VMs are unaware of the VXLAN.

VM traffic

Only VMs within the same VXLAN segment can communicate with each other. The following steps show how a VM communicates with a VM on a different host within the same VXLAN segment:

1. The VM sends a MAC frame destined to the target.
2. The VTEP looks up the VNID that the VM is associated with to verify that the destination MAC is on the same segment. If the destination MAC is on the same segment, the process continues.
3. The VTEP inserts an outer header comprising an outer MAC, outer IP address, and VXLAN header in front of the original MAC frame.
4. The VTEP transmits the final packet out to the destination IP address of the remote VTEP, connecting the destination VM addressed by the inner MAC destination address.

5. The remote VTEP verifies that the VNID is valid for the destination VM.
6. If the VNID is valid, the remote VTEP strips the packet outer header and passes the packet to the destination VM. The destination VM never knows about the VNID, or that the frame transmitted with VXLAN encapsulation.

VXLAN Gateway management methods

Two methods are available to manage VXLAN Gateway:

- Static management using CLI and local configuration files.
- Dynamic management using OVSDB protocol support for VXLAN Gateway (demo feature).

For static management, you use CLI or local configuration files to configure and manage the VXLAN Gateway hardware-based VTEP functions. You must manually configure the Source-VTEP, VNID, VNID to I-SID bindings, Remote-VTEP, and VNI to Remote-VTEP associations, and MAC learning on the VNID only occurs at the data-plane level.

For dynamic management, you use OVSDB protocol support for VXLAN Gateway to configure and manage the VXLAN Gateway hardware-based VTEP functions with Open vSwitch Database Management Protocol (OVSDB). You must configure at least one physical Network Virtualization Controller (NVC) running VMware NSX with OVSDB features. You can add a Hardware VTEP, add a logical switch, configure VNID to VTEP bindings, and configure a replication cluster in NSX. The VXLAN Gateway must have OVSDB enabled and the Network Virtualization Controller (NVC) must communicate on the OVSDB managed interface. You must manually configure the Source-VTEP and NVC IP addresses, and the VNID to I-SID binding in NSX. The NVC can manage the VNID, Remote-VTEP, and VNI to Remote-VTEP associations. The NVC distributes the hosts MAC and IP addresses learned by the VTEP over the VXLAN.

Important:

OVSDB protocol support for VXLAN Gateway is a demo feature. Do not use the OVSDB features in production environments.

The following table shows the functional differences between the two VXLAN Gateway management methods.

Table 1: VXLAN Gateway management methods

Function	Static management	Dynamic management (DEMO)
VNID provisioning	Manually from CLI or EDM	Automatically from NVC
VNID to I-SID binding	Manually from CLI or EDM	Automatically from NVC
Remote-VTEP	Manually from CLI or EDM	MAC learning from NVC
VNID to VTEP association for replicating BUM traffic	Manually from CLI or EDM	MAC learning from NVC
MAC learning at data-plane	Not distributed to other VXLAN Gateways	Distributed from NVC to other VXLAN Gateways
ARP suppression	Not supported	Not supported

VXLAN Gateway licensing

VXLAN Gateway requires a Premier or Premier with MACsec license. If a Premier License is not present, you cannot configure VXLAN Gateway.

For more information about licensing, see *Administering*.

VXLAN modes

The VXLAN Gateway implementation is available in two modes: Base Interworking Mode and Full Interworking Mode. By default, the Base Interworking Mode is enabled and Full Interworking Mode is disabled. You change modes through a boot configuration flag.

Important:

Changing the mode requires a reboot for the change to take effect, which can cause a loss of traffic.

Base Interworking Mode

Base Interworking Mode is the default mode. In this mode, VXLAN Gateway supports Layer 2 gateway communication at line rate between VXLAN and traditional VLAN environments.

Base Interworking Mode has the following limitations:

- No support for VXLAN-to-VXLAN communication.
- No support for VXLAN-to-SPB communication.
- No support for SMLT, vIST, or Simplified vIST.
- No support for OVSDB protocol support for VXLAN Gateway.

Each VNID represents a service instance for the VXLAN Gateway. To enable a VNID to extend into the SPB network, you must use Full Interworking Mode.

Note:

If your configuration does not have a VXLAN Gateway but has SMLT or vIST, you can enable VXLAN Gateway in the Full Interworking Mode *only*. SMLT and vIST are incompatible and mutually exclusive with the Base Interworking Mode.

Full Interworking Mode

Full Interworking Mode is **not** the default mode. To configure Full Interworking Mode, you must enable the `vxlان-gw-full-interworking-mode` boot flag, save the configuration, and then reboot the switch.

Full Interworking Mode supports the Base Interworking Mode communication between VXLAN and traditional VLAN environments as well as VXLAN-to-VXLAN communication and all SPB functionality including vIST and SMLT. Layer 2 and Layer 3 services extend across VXLAN, VLAN, and SPB domains using an internal loopback. In Full Interworking Mode, VXLAN traffic flows at less than full line rate and the total amount of available internal loopback bandwidth depends on the platform used.

Full Interworking Mode has the following limitations:

- No support for Simplified vIST.
- A VNID cannot have the same value as an I-SID. A consistency check ensures that this does not happen if you change the boot configuration flag to turn off Full Interworking Mode and go into Base Interworking Mode.

Types of VXLAN Gateway deployments

VXLAN Gateway provides the following deployment solutions:

- [VXLAN-to-VLAN Layer 2 gateway deployment](#) on page 17
- [VXLAN-to-VXLAN Layer 3 gateway deployment](#) on page 18
- [VXLAN-to-SPB gateway deployment](#) on page 19

These gateways provide VXLAN tunnel termination functions for switches that could be Top of Rack (ToR) switches, access switches, or switches higher up in the data center network topology such as a core switch or even WAN edge devices. The last case (WAN edge) could involve a Provider Edge (PE) router that terminates VXLAN tunnels in a hybrid cloud environment.

VXLAN-to-VLAN Layer 2 gateway deployment

The VXLAN-to-VLAN deployment is where the VXLAN Gateway communicates between VXLAN segments and VLAN segments. Use this scenario when nodes on a VXLAN overlay network need to communicate with nodes on legacy networks that are VLAN based as shown in the following figure.

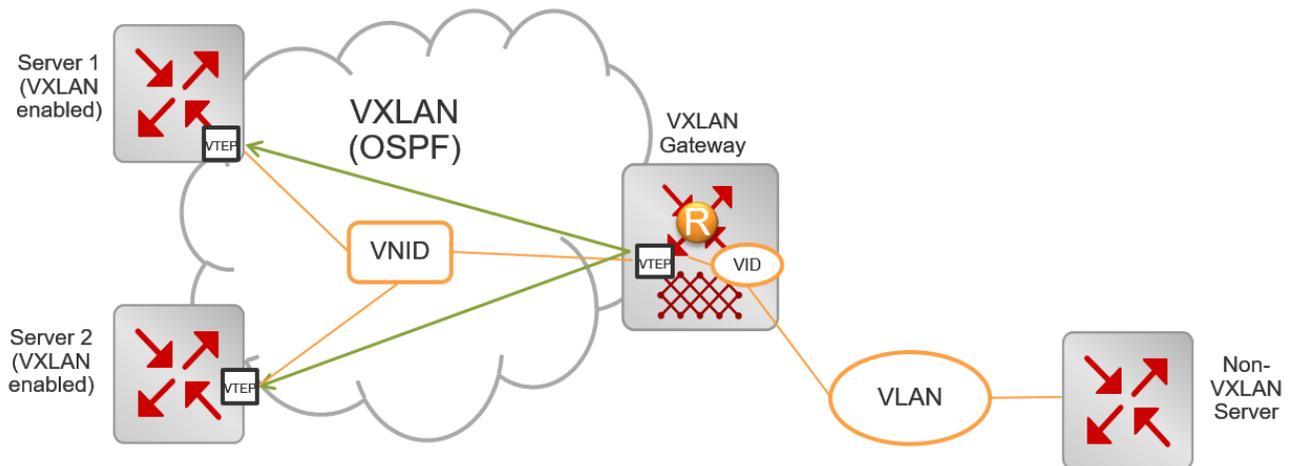


Figure 1: VXLAN-to-VLAN deployment

In this type of deployment, the VXLAN Gateway maps the VXLAN VNID to a VLAN in the following way:

- When packets traverse a VXLAN segment into the VLAN segment, the VXLAN VTEP removes the VXLAN header and forwards regular VLAN packets to a physical port based on the destination MAC address of the inner Ethernet frame.

The VTEP discards decapsulated frames with the inner VLAN ID unless you explicitly configure them to be passed on to the non-VXLAN interface. The VXLAN VNID to VLAN mapping is statically configured.

- In the reverse direction when packets traverse a VLAN segment into a VXLAN segment, the VXLAN VTEP maps the incoming frames from the non-VXLAN interfaces to a specific VXLAN overlay network based on the VLAN ID in the frame. This complies with the standard VXLAN RFC.

The VTEP also accepts both untagged and tagged inner frames and adds or replaces the tag of the inner frame with the local VLAN mapped to the VNID before sending out the frame.

The VTEP removes the VLAN ID of the original frame before it encapsulates the frame for VXLAN.

VXLAN-to-VXLAN Layer 3 gateway deployment

The VXLAN-to-VXLAN deployment is where the VXLAN Gateway communicates between VXLAN segments. Use this scenario when one VXLAN segment needs to communicate with another VXLAN segment.

When a packet traverses a VXLAN segment and enters the VXLAN Gateway, the VXLAN VTEP maps one VXLAN segment to another in the following way:

- Decapsulates the VXLAN header.
- Performs a route lookup of the customer packet.
- Identifies that the next hop VLAN is in another VXLAN segment.
- Encapsulates the packet with a new VXLAN header with a new VNID.
- Forwards the packet out.

The following figure shows two virtualized servers attached to a Layer 3 infrastructure. The servers could be on the same rack, on different racks, or across data centers within the same administrative domain.

There are four VXLAN overlay networks identified by the VNIDs: 22, 34, 74, and 98. Notice how the VMs are associated with a VNID. For example, VM1-1 in Server 1 and VM2-4 on Server 2 are on the same VXLAN overlay network identified by VNID 22.

The VMs do not know about the overlay networks and transport method because the encapsulation and decapsulation happen transparently at the VTEPs on Servers 1 and 2. The other overlay networks and the corresponding VMs are: VM1-2 on Server 1 and VM2-1 on Server 2 both on VNID 34, VM1-3 on Server 1 and VM2-2 on Server 2 on VNID 74, and finally VM1-4 on Server 1 and VM2-3 on Server 2 on VNID 98.

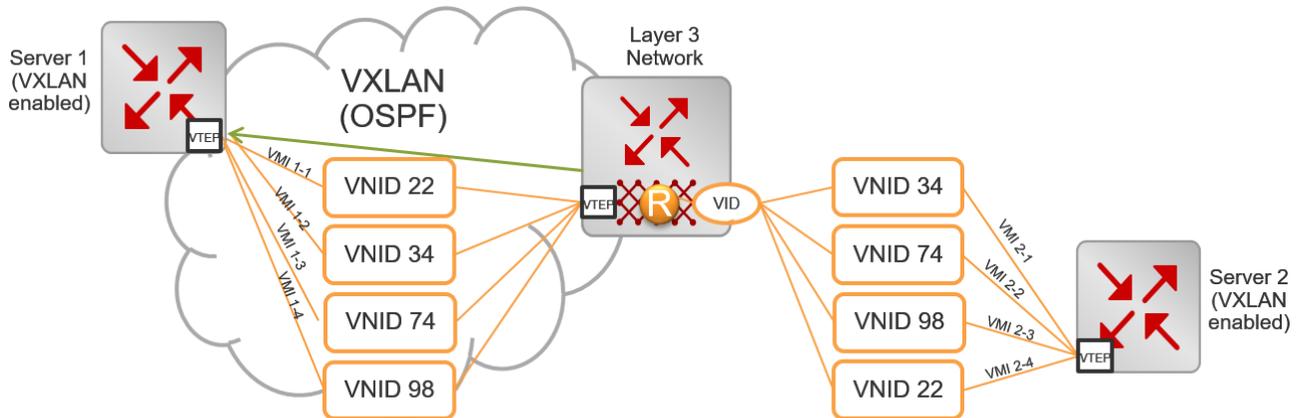


Figure 2: VXLAN-to-VXLAN deployment

VXLAN-to-SPB gateway deployment

The VXLAN-to-SPB deployment is where the VXLAN Gateway communicates between a VXLAN segment and an SPB domain. This type of deployment provides a solution for customers using VXLAN in their data-centers to interoperate with Fabric Connect as shown in the following figure.

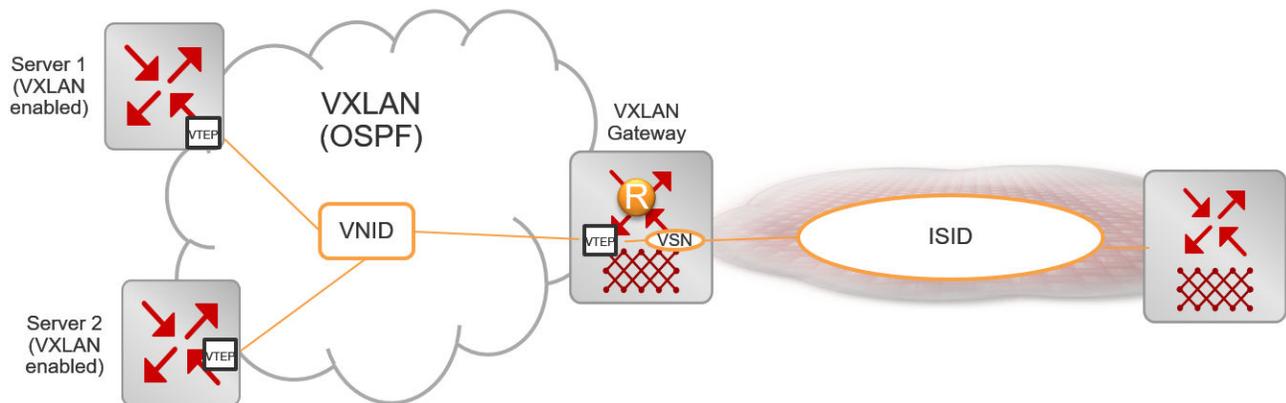


Figure 3: VXLAN-to-SPB deployment

The VXLAN Gateway maps the VXLAN VNID to I-SID in the following way:

- When packets traverse a VXLAN segment into an SPB domain, the VXLAN VTEP decapsulates the VXLAN packet, encapsulates it with MAC-in-MAC encapsulation, and then forwards it into the SPB domain.
- In the reverse direction, the VXLAN VTEP decapsulates the MAC-in-MAC packet from an SPB domain, encapsulates it with a VXLAN header, and then forwards it into a VXLAN segment.

VXLAN Gateway considerations and limitations

Review the following considerations, limitations, and behavioral characteristics associated with VXLAN Gateway.

Unsupported features

VXLAN Gateway does not support the following features:

- IP Multicast over Fabric Connect traffic cannot be carried over into a VXLAN domain through a VXLAN Gateway.
- SPB-PIM Gateway and VXLAN Gateway interoperability is not supported.
- VXLAN Gateway does not support Simplified vIST in any mode.

QoS and VXLAN Gateway

- If present in the packet, VXLAN Gateway honors user-configured ingress dot1p and DSCP mappings to derive the internal CoS.
- Customer packets are not remarked.
- Control packets are handled by high priority CoS queues.
- DSCP bits in the outer IP header of VXLAN-encapsulated packets are always derived from the internal QoS, irrespective of the ingress port DiffServ configuration. Customer packet IP DSCP bits are not modified as part of VXLAN encapsulation. For more information about QoS, see *Configuring QoS and ACL-Based Traffic Filtering*.
- DSCP bits in the inner IP header of VXLAN-encapsulated packets change when traffic comes in on a Layer 2 Trusted or a Layer 3 Untrusted port with DiffServ enabled.

ECMP support for VXLAN Gateway and Fabric Extend

VXLAN Gateway requires ECMP support to communicate with remote VTEPs. The software extended this ECMP support to Fabric Extend Layer 3 core tunnels. Therefore, if your switch supports VXLAN Gateway, it also supports ECMP for both VXLAN Gateway and Fabric Extend.

Base Interworking Mode considerations and limitations

- After associating a VNID with an I-SID (`vnid <vnid value> i-sid <isid value>`), you cannot create an I-SID with the same value as the `<vnid value>` and vice versa.
- There is no support for SMLT, vIST, or Simplified vIST in Base Interworking Mode.
- Neither Base Interworking Mode nor Full Interworking Mode support Fabric Attach endpoints.
- Base Interworking Mode does not allow SMLT and vIST configurations. If you change from Full Interworking Mode to Base Interworking Mode, make sure there are none of these configurations.
- In Base Interworking Mode, the `untagged-traffic <port >< mlt>` support is limited to untagged traffic forwarding only. No control packet forwarding is supported on this port or MLT. LACP MLT is not supported as part of `untagged-traffic <mlt>`.

Full Interworking Mode considerations and limitations

- Whenever you change from Base Interworking Mode to Full Interworking Mode, you have to reboot the switch. The software automatically changes the configuration to be in line with Full Interworking Mode.
- Full Interworking Mode does not have VXLAN endpoints. Configure ELAN endpoints under the ELAN I-SID only.
- Neither Base Interworking Mode nor Full Interworking Mode support Fabric Attach endpoints.

General configuration guidelines

- In VXLAN environments that have an underlying IP network with SMLT deployment, you must enable RSMLT in the underlying network on vIST VTEP devices.
- If you use more than one VXLAN Gateway node for the same I-SID, you must ensure that loops are prevented.
- The VTEP remote destination IP table in the datapath hardware is shared with the Fabric Extend IP core remote tunnel destination IP table.
 - The maximum number of FE tunnel destination IP addresses is 256.
 - The maximum number of VTEP IP addresses is 500.

For every FE tunnel destination IP you configure, you must reduce the number of VTEP IPs by one. For example, if you configure the maximum number of FE tunnel destination IPs then the maximum of VTEP remote destination that you can configure is $500-256=244$.

- The total number of MACs that the VTEP learns in the switch is constant and includes the MACs learned on VNID, ISID, and on VLANs (normal and CVLANs).
 - MAC addresses that the VTEP learns from the VNID table are also learned when an I-SID is associated with the VNID. Therefore, the maximum number of MAC addresses that the VTEP can support is reduced accordingly.
 - If a VLAN is associated with an I-SID, and that I-SID is associated with a VNID, then the VTEP learns each customer MAC address from three different tables (VLAN, I-SID, and VNID). This consumes three MAC entry records. In a vIST scenario where all I-SIDs have associated VNIDs, then the max number of C-MACs that the VTEP learns is reduced by $224K$ divided by $3 = 74K$ MAC addresses.
 - If a Switched-UNI endpoint is associated with an I-SID, then the VTEP learns the C-MAC only once in the I-SID table. If this I-SID is associated with a VNID, then the VTEP learns the C-MAC twice: once in the I-SID table and once in the VXLAN VNID table.
- IP filter rules that you create do not work for VXLAN tunnel terminated packets. This restriction also applies to the Fabric Extend Layer 3 core.
- All ports in an MLT configured as a Switched-UNI endpoint or a VXLAN endpoint have the same properties as configured under the MLT interface.
- A VXLAN tunnel VTEP source IP cannot be a broadcast or multicast IP address. Configure it as a loopback IP address and not a brouter IP address.
- The VXLAN source address (`vtep-source-ip`) should be different from the SPBM IP Shortcut (`ip-source-address`) and Fabric Extend (`ip-tunnel-source-address`) source addresses.

- If you configure the loopback IP address under a VRF, the `vtep-source-ip` address must be set up under the same VRF.
- The remote VTEP IP address (`vtep <id> ip <ip address>`) cannot be local to the system.
- You can configure a maximum of 500 remote VTEPs per VNID.
- You can change the VTEP source IP and remote VTEP IP dynamically.
- VTEPs have to be reachable for the tunnel to be up.
- The tunnel ID for a VTEP has to be the same on both vIST peers.
- If the VTEP source IP is configured under a VRF, then the VRF cannot have an I-SID associated with it or vice versa.
- Two VNIDs cannot be mapped to the same I-SID.
- An I-SID associated with a VNID cannot be a T-UNI or E-Tree I-SID.
- To prevent routing loops, the platform VLANs used for VXLAN interworking should NOT be in the same VRF or GRT as the IP interfaces in the underlay network.
- You cannot delete a VTEP source IP address (`vtep-source-ip`) when VTEPs are configured.
- If an ELAN I-SID attached to a VNID has a platform VLAN associated with it, then you cannot delete the platform VLAN or change the I-SID associated with that VLAN.
- You cannot delete the loopback IP associated with a VTEP source IP.
- If the VTEP source IP is configured under a VRF, then you cannot delete the VRF.
- You cannot delete or modify the VLAN with I-SID that is associated with a VNID.
- If an ELAN I-SID attached to a VNID also has a platform VLAN associated with it, then you cannot delete the platform VLAN or change the I-SID associated with that VLAN.

OVSDB protocol support for VXLAN Gateway considerations and limitations

Review the following considerations, limitations, and behavioral characteristics associated with OVSDB protocol support for VXLAN Gateway.

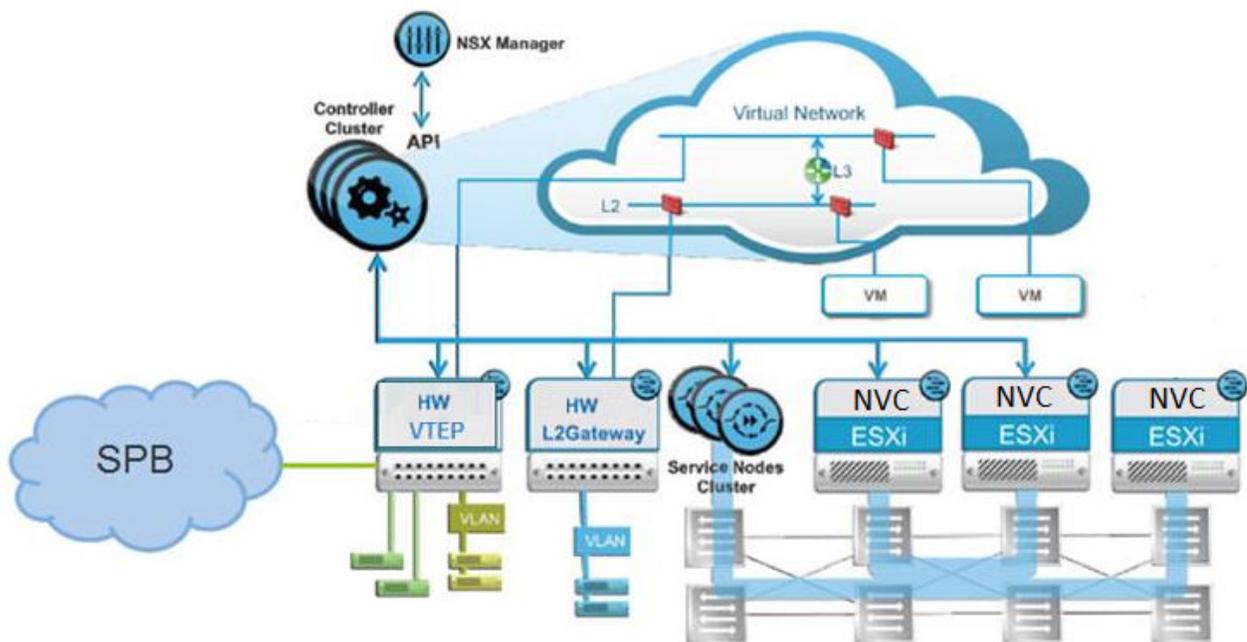
Network Virtualization Controllers dependencies and restrictions

OVSDB protocol support for VXLAN Gateway requires at least one Network Virtualization Controller (NVC) that runs OVSDB:

- OVS is an open source implementation of a virtual switch and OVSDB is a standard defined in RFC 7047.
- You require at least one NVC. The NVC is an ESXi host server that runs the VMware NSX Network Virtualization and Security Platform (minimum version 6.2.4).
- You must configure VMware NSX. You can add a Hardware VTEP, add a logical switch, configure VNID to VTEP bindings, and configure a replication cluster in NSX.

- The NSX logical switch on the NVC is assigned a VNID. The VXLAN Gateway is assigned an OVSDB managed interface I-SID. You must configure the VNID to I-SID binding in NSX.
- SSL is the default protocol for NVC to VXLAN Gateway communications.
- You can configure more than one NVC in a controller cluster for high availability and load-balancing.
- The NVC performs MAC learning on the VXLAN tunnel.
- ARP suppression is not supported.
- One NVC manages all encapsulated Broadcast, Unknown Unicast, and Multicast (BUM) traffic. This NVC replicates the BUM traffic to all other VTEPs in the network. If connectivity to the NVC managing BUM traffic fails, a loss of BUM traffic can occur. NVC cluster NSX service node replication is not applicable for BUM traffic.
- If the NVC does not populate the OVSDB tables for a VXLAN Gateway, you can remove and add the Hardware Gateway service for the HW VTEP from NSX.
- Under high stress scenarios, NVC connectivity might disconnect and reconnect.

The following diagram shows an example of NVCs managing a distributed virtual network.



OVSDB protocol support for VXLAN Gateway dependencies and restrictions

- The VXLAN Gateway must operate in Full Interworking mode.
- You must configure at least one NVC connection to the VXLAN Gateway.
- You can configure up to three NVCs connections to one VXLAN Gateway in this release.
- The VXLAN Gateway must connect with the NVCs on the NLS managed interface for OVSDB.

- You cannot change or delete the VTEP source IP address when OVSDB is enabled.
- The OVSDB managed interface I-SIDs are communicated to the NVC as physical ports of the switch.
- EDM support for OVSDB is limited in this release. You can use CLI where EDM configuration is unavailable.
- The VXLAN Gateway general considerations and limitations also apply to OVSDB protocol support for VXLAN Gateway.

*** Note:**

NVC connectivity issues for SSL communication failures are not logged on the switch. You can use the logs on the NVC to troubleshoot SSL communications with the VXLAN Gateway.

OVS modules on VXLAN Gateway

OVSDB protocol support for VXLAN Gateway uses OVS modules on the VXLAN Gateway. The switch runs the following OVS modules:

- OVSDB server — Manages the database.
- Database — Manages the OVS schema and HW VTEP schema and JSON content.

OVSDB replication considerations and limitations

*** Note:**

OVSDB replication configuration is available for testing only. OVSDB replication is not supported in this release. OVSDB replication is different than NSX service node replication. Service node replication is supported and configured in NSX.

- You require at least two NVCs to enable OVSDB replication.
- OVSDB replication operates in an ACTIVE-STANDBY server configuration.
- One active and multiple standby NVCs are supported by the VXLAN Gateway. However one VXLAN Gateway can support a maximum of three NVCs.
- VNID to I-SID binding, Remote VTEPs, VNI to Remote VTEPs for BUM, Remote MACs, and Local MACs are replicated across NVCs.

Chapter 5: VXLAN Gateway configuration using the CLI

This section provides procedures to configure VXLAN Gateway using the CLI.

Configuring the VXLAN Gateway boot flag

There are two configuration modes: Base Interworking Mode and Full Interworking Mode. The Base Interworking Mode is the default mode. To switch to the Full Interworking Mode, use this procedure.

Important:

After you change this boot flag, you must save the change to the configuration file and reboot the switch.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable Full Interworking Mode:

```
boot config flags vxlan-gw-full-interworking-mode
```

3. Save the changed configuration.

```
save config
```

4. Restart the switch.

```
reset
```

5. Verify that VXLAN Gateway is in Full Interworking Mode:

```
show boot config flags
```

The **true** parameter in the output shows that Full Interworking Mode is enabled.

```
flags vxlan-gw-full-interworking-mode true
```

Configuring VXLAN Gateway

Use the following procedure to configure VXLAN Gateway parameters to allow SPBM to operate on the switch.

Before you begin

Use the `show boot config flags` to verify that you are in the mode you want. The `vxlan-gw-full-interworking-mode` should be set to `false` for the Base Interworking Mode (default) and `true` for the Full Interworking Mode.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the VTEP source IP address:

```
vtep source-ip <A.B.C.D> [vrf WORD<1-16>]
```

3. Configure the remote VTEP IP address:

```
vtep <1-500> ip <A.B.C.D> [name WORD<1-64>]
```

4. Create a VNID instance:

```
vnid <1-16777215> i-sid <1-16777215>
```

* Note:

The command prompt changes to **#vxlan** to indicate that you are now in VXLAN Configuration mode for the VNID specified in `vnid <1-16777215>`.

5. Associate VLANs in a port or MLT list to this VNID instance:

```
c-vid <1-4094> port <{slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]}>
```

or

```
c-vid <1-4094> mlt <1-4094>
```

6. Specify the ports or MLT in this VNID instance that will support untagged traffic:

```
untagged-traffic port <{slot/port[/sub-port]}>
```

or

```
untagged-traffic mlt <1-4094>
```

7. Associate VTEPs to this VNID instance:

```
vtep <1-255>
```

Examples

For an example of this procedure, see [VXLAN Gateway configuration example in Base Interworking mode](#) on page 51 or [VXLAN Gateway configuration example in Full Interworking mode](#) on page 51.

Variable definitions

Use the data in the following table to use the `vte` command.

Variable	Value
source-ip <A.B.C.D> [vrf WORD<1–16>]	Specifies the VXLAN tunnel end point (VTEP) source IP address, which can be on the GRT or a VRF. * Note: The VTEP source IP address must be on a loopback interface.
<1–500> ip <A.B.C.D> [name WORD<1–64>]	Specifies an index value and an IP address that uniquely identifies this remote VTEP. Optionally, you can assign a specific name to this tunnel. By default, the switch assigns a name in this format: VTEP-#ID-<IP address> * Note: The remote VTEP IP address cannot be a local, broadcast, or multicast IP address.

Use the data in the following table to use the `vnid` command.

Variable	Value
<1–16777215> i-sid <1–16777215>	Uses this VNID and I-SID information to create a VNID instance and enter VXLAN Configuration Mode in the CLI. * Note: A VNID must not have the same value as an I-SID.

Use the data in the following table to configure commands in the VXLAN Configuration Mode (`#vxlan`). These commands apply to the VNID instance that you specified in the `vnid <1–16777215> i-sid <1–16777215>` command.

Variable	Value
c-vid <1-4094> port <{slot/port[/sub-port] [-slot/port[/sub-port]] [...]}>	Specifies a value that uniquely identifies the customer VLAN ID and ports of this ELAN end point.

Table continues...

Variable	Value
	<p>The switch reserves the following VLAN IDs:</p> <ul style="list-style-type: none"> • VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • VLAN IDs 4060 to 4094 are reserved for internal use. • VLAN IDs 3500 to 3999 are reserved if you enable VRF scaling and SPBM mode. • VLAN ID 4095 is not used. • VLAN ID 4096 is reserved for untagged traffic.
<code>c-vid <1-4094> mlt <1-4094></code>	<p>Specifies a value that uniquely identifies the customer VLAN ID and MLTs of this ELAN end point.</p> <p>The switch reserves the following VLAN IDs:</p> <ul style="list-style-type: none"> • VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • VLAN IDs 4060 to 4094 are reserved for internal use. • VLAN IDs 3500 to 3999 are reserved if you enable VRF scaling and SPBM mode. • VLAN ID 4095 is not used. • VLAN ID 4096 is reserved for untagged traffic.
<code>untagged-traffic port <{slot/port[/sub-port] [-slot/port[/sub-port]] [...]}></code>	Specifies the ports that support untagged traffic.
<code>untagged-traffic mlt <1-4094></code>	Specifies the MLTs that support untagged traffic.
<code>vtep <1-255></code>	Lists the remote VTEP destinations to associate with the specified VNID.

Flushing the MAC forwarding table

Use this procedure to flush all the learned MAC addresses from the forwarding database of the selected VNID.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Flush the MAC FDB:

```
vnid mac-address-entry <1-16777215> flush
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# vnid mac-address-entry 100 flush
```

VXLAN Gateway show commands

Use the procedures in this section to display specific information about VXLAN Gateway on the switch.

Displaying VTEP source information

Use the following procedure to display information about the VTEP source.

Procedure

1. Enter Privileged EXEC mode:


```
enable
```
2. Display the VTEP's source IP address and the name of the VRF:


```
show vtep local
```

Example

```
Switch:1# show vtep local
=====
                          VTEP General Info
=====
vtep source-ip-address: 198.51.100.1
vtep vrf : GlobalRouter
```

Displaying remote VTEP information

Use the following procedure to display information about the remote VTEP.

Procedure

1. Enter Privileged EXEC mode:


```
enable
```
2. Display information about all of the remote VTEPs or use the option to specify a particular VTEP:


```
show vtep remote [<1-500>]
```

Example

```
Switch:1# show vtep remote
=====
Remote VTEP Info
=====
ID      NAME      VTEP      L3_VTEP_NEXT_HOP_INFO
DEST-IP  PORT/MLT  VLAN      VRF
-----
1       vtep-1    198.51.100.6  M1t78      3851      GlobalRouter
56      vtep-56   198.51.100.22 M1t130     3855      underlay-vxlan
-----
2 out of 2 Total Num of remote vteps
=====
```

Displaying the name of the remote VTEP

Use the following procedure to display the name of the remote VTEP.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the names of all remote VTEPs or use the option to specify a particular VTEP:
show vtep remote name [<vtepId>]

Example

```
Switch:1# show vtep remote name
=====
Remote VTEP Name
=====
ID      NAME
-----
1       vxlan_tunnel_1
2       --
3       --
-----
3 out of 3 Total Num of remote vteps
=====
```

Displaying VNID information

Use the following procedure to display information about the VNID.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display information about all the VNIDs or use the option to specify a particular VNID:

```
show vnid i-sid [<1-500>]
```

Example

```
Switch:1# show vnid i-sid
```

```
=====
                               vnid Info
=====
```

VNID ID	ISID ID	PORT INTERFACES	MLT INTERFACES	VTEP IDS
110	120	c3:1/10	-	2,56
11000060	1060	c60:1/10	-	56
11000065	1065	c65:1/10	-	56
11000066	1066	c66:1/48	-	2
.				
.				
.				

Displaying the MAC addresses in the VNID FDB

Use the following procedure to display the MAC addresses in the VNID forwarding database.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display all the VNIDs in the FDB table or use one of the options to specify a particular VNID:

```
show vnid mac-address-entry [<1-16777215> | port <{slot/port [-slot/  
port]} [,...]]> | mac <0x00:0x00:0x00:0x00:0x00:0x00> | remote]
```

Example

```
Switch:1# show vnid mac-address-entry
```

```
=====
                               VNID Fdb Table
=====
```

VNID	STATUS	MAC-ADDRESS	INTERFACE	TYPE
11000065	learned	00:00:5e:00:01:01	vtep_56	REMOTE
14000100	learned	a4:25:1b:51:a5:04	vtep_67	REMOTE
14000100	learned	a4:25:1b:52:a5:04	vtep_67	REMOTE
14000100	learned	a4:25:1b:52:bd:03	vtep_2	REMOTE

c: customer vid u: untagged-traffic

All 4 out of 4 Total Num of vnid FDB Entries displayed

```
=====
```

Displaying the VXLAN running configuration

Use the following procedure to display the currently running VXLAN configuration.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the running configuration:

```
show running-config module vxlan
```

Example

This example shows the VXLAN information only. It does not show the hardware-specific information that also appears in the output for this command.

The example shows the VXLAN Gateway in Base Interworking Mode. It also shows an external loopback (see **c-vid** commands), which enables communication between the VNID and an I-SID in an SPB network.

```
Switch:1# show running-config module vxlan

#
# VTEP CONFIGURATION
#
vtep source-ip 1.1.1.1 vrf underlay-vxlan

#
# REMOTE VTEP CONFIGURATIONS
#
vtep 2 ip 1.1.1.1 name "vtep-2"
vtep 67 ip 1.1.1.67 name "vtep-67"

#
# VNID CONFIGURATION
#
vnid 12002105 i-sid 102105
c-vid 105 port 1/11
c-vid 2105 port 1/10
vtep 67
exit
```

Chapter 6: OVSDB protocol support for VXLAN Gateway configuration using CLI

This section provides procedures to configure OVSDB protocol support for VXLAN Gateway using the Command Line Interface (CLI).

Configuring OVSDB managed interfaces

Use the following procedure to configure OVSDB managed interfaces on the switch.

Before you begin

- You must enable VXLAN Gateway Full Interworking Mode. For more information, see [Configuring the VXLAN Gateway boot flag](#) on page 25.
- You must create the required Flex UNI or CVLAN I-SIDs before you configure the OVSDB managed interface.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enter OVSDB configuration mode:
`ovsdb`
3. Configure the OVSDB managed interface:
`managed-interface i-sids WORD<1-1024>`

Examples

The following is an example configuring two I-SIDs as OVSDB management interfaces.

```
SWITCH:1>enable
SWITCH:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SWITCH:1 (config)#ovsdb
SWITCH:1 (config-ovsdb)#managed-interface i-sids 300
SWITCH:1 (config-ovsdb)#managed-interface i-sids 500
```

Next steps

Configure the OVSDB protocol support for VXLAN Gateway, see [Configuring OVSDB protocol support for VXLAN Gateway](#) on page 34.

Variable definitions

Use the data in the following table to use the `managed-interface` command.

Variable	Value
i-sids WORD<1–1024>	<p>Specifies the I-SID information to create an interface to manage OVSDB protocol support for VXLAN Gateway.</p> <p> Note:</p> <p>You must create the required Flex UNI or CVLAN I-SIDs before you configure the managed interface I-SID.</p>

Configuring OVSDB protocol support for VXLAN Gateway

Use the following procedure to configure OVSDB protocol support for VXLAN Gateway on the switch.

Before you begin

- You must enable VXLAN Gateway Full Interworking Mode. You can use `show boot config flags` to verify the current VXLAN Gateway mode. For more information, see [Configuring the VXLAN Gateway boot flag](#) on page 25.
- You must configure an OVSDB management interface to establish a session with the NVC. For more information, see [Configuring OVSDB managed interfaces](#) on page 33.
- You must transfer an OVSDB certificate file and private-key file to the flash storage of the switch. You can use an ovs-pki utility with SSL libraries to generate the private keys and certificates. You can use `boot config flags FTPD` and then an SCP utility to transfer the private key and certificate file to the flash storage of the switch.
- If the switch is an aggregation switch, the IST peer must support OVSDB, have the same source VTEP-IP and OVSDB managed-interface, and must communicate with the NVC management IP.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enter OVSDB configuration mode:

```
ovsdb
```

3. Configure the source VTEP IP address:

```
vtep source-ip <A.B.C.D> [vrf WORD<1-16>]
```

4. Install the OVSDB certificate file:

```
install-cert-file WORD<1-128>
```

5. Install the OVSDB private key:

```
private-key WORD<1-128>
```

6. Enable OVSDB protocol support for VXLAN Gateway:

```
enable
```

7. Configure the NVC:

```
controller <1-100> ip address <A.B.C.D> protocol <ssl|tcp> [port <1-65535>]
```

Examples

The following is an example of configuring a VTEP source IP, installing an OVSDB certificate file, installing an OVSDB private key, enabling OVSDB protocol support for VXLAN Gateway, and configuring the IP address, protocol, and port for one NVC.

```
SWITCH:1>enable
SWITCH:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SWITCH:1(config)#vtep source-ip 192.0.2.1 vrf vxlan-underlay
SWITCH:1(config)#ovsdb
SWITCH:1(config-ovsdb)#install-cert-file /intflash/tom/sc-cert.pem
SWITCH:1(config-ovsdb)#private-key /intflash/tom/sc-privkey.pem
SWITCH:1(config-ovsdb)#enable
SWITCH:1(config-ovsdb)#controller 1 ip address 192.0.2.2 protocol ssl port 6640
```

* Note:

You can configure multiple controllers for high availability. One VXLAN Gateway can support a maximum of three controllers.

! Important:

If you add or delete a controller, or modify the OVSDB managed interface when a controller is configured, the existing controller connections reset. Log messages generate to indicate the status changes as the controllers disconnect and reconnect.

Variable definitions

Use the data in the following table to use the `vtep source-ip` command.

Variable	Value
<A.B.C.D> [vrf WORD<1–16>]	Specifies the VXLAN tunnel end point (VTEP) source IP address in IPv4 format. Optionally you can specify a VRF.  Note: The VTEP source IP address must be on a loopback interface.

Use the data in the following table to use the `install-cert-file` command.

Variable	Value
WORD<1–128>	Specifies the path and file name of the OVSDB certificate file.

Use the data in the following table to use the `private-key` command.

Variable	Value
WORD<1–128>	Specifies the path and file name of the OVSDB private key.

Use the data in the following table to use the `controller` command.

Variable	Value
<1–100>	Specifies the ID of the controller.
ip address <A.B.C.D>	Specifies the IP address of the controller in IPv4 format.
protocol <ssl tcp>	Specifies the networking protocol as SSL or TCP for controller communications.
port <1–65535>	Specifies the networking port of the controller.

Configuring OVSDB replication

Use the following procedure to configure OVSDB replication to support vIST.

 **Note:**

OVSDB replication configuration is available for testing only. OVSDB replication is not supported in this release.

Before you begin

- You must enable VXLAN Gateway Full Interworking Mode on the switch.
- You must configure an OVSDB managed interface on the switch.
- You must configure and enable OVSDB protocol support for VXLAN Gateway on the switch.

- You must have at least two NVCs configured to communicate on the managed interface network.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enter OVSDB configuration mode:

```
ovsdb
```

3. Configure the OVSDB replication:

```
replication peer-ip <A.B.C.D> local-ip <A.B.C.D>
```

Examples

The following is an example configuring OVSDB replication to support VIST.

```
SWITCH:1>enable
SWITCH:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SWITCH:1(config)#ovsdb
SWITCH:1(config-ovsdb)#replication peer-ip 192.0.2.3 local-ip 192.0.2.2
```

Variable definitions

Use the data in the following table to use the **replication** command.

Variable	Value
peer-ip <A.B.C.D>	Specifies the IP address of the secondary controller for OVSDB replication support.
local-ip <A.B.C.D>	Specifies the IP address of the primary controller for OVSDB replication support.

Displaying the OVSDB configuration

Use the following procedure to display the OVSDB protocol support for VXLAN Gateway configuration.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the OVSDB configuration:

```
show ovsdb config
```

Examples

The following is an example of displaying the OVSDB configuration.

```
SWITCH:1>show ovbdb config
=====
                        Ovbdb Config Info
=====
                        Status - Enabled
                        Certificate - Installed
                        Private Key - Installed
                        Replication Peer NLS-IP - 192.0.2.3
                        Replication Local NLS-IP - 192.0.2.2
=====
                        Ovbdb Controller Config Info
=====
ID          Protocol      IP          Port
1           ssl             192.0.2.2   6640
```

Displaying the OVSDB controller status

Use the following procedure to display the OVSDB controller status.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the OVSDB controller status:

```
show ovbdb controller status
```

Examples

The following is an example of displaying the OVSDB controller status.

```
SWITCH:1>show ovbdb controller status
=====
                        OVSDB Controller Status
=====
Protocol      IP          Port      status      source
-----
ssl           192.0.2.2   6640      Up          Learnt
ssl           192.0.2.3   6640      Down        Learnt
2 out of 2 Total Num of Controllers displayed
```

Displaying the OVSDB replication state

Use the following procedure to display the OVSDB replication state.

Procedure

1. Log on to the switch to enter User EXEC mode.

2. Display the OVSDB controller status:

```
show ovldb replication state
```

Examples

The following is an example of displaying the OVSDB replication state.

```
SWITCH:1>show ovldb replication state
```

```
=====
                               Ovldb Replication Status
=====
```

Local NLS IP	Peer NLS IP	State
192.0.2.2	192.0.2.3	ENABLED

Displaying the OVSDB managed interfaces

Use the following procedure to display the OVSDB managed interfaces.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the OVSDB managed interfaces:

```
show ovldb managed-interface
```

Examples

The following is an example of displaying the OVSDB managed interfaces.

```
SWITCH:1>show ovldb managed-interface
```

```
=====
                               OVSDB Managed Interfaces - ISID
=====
```

```
50001
1 out of 1 Total Num of Managed Interfaces displayed
```

Chapter 7: VXLAN Gateway configuration using EDM

This section provides procedures to configure VXLAN Gateway using the Enterprise Device Manager (EDM).

Configuring the VXLAN Gateway boot flag

There are two configuration modes: Base Interworking Mode and Full Interworking Mode. The Base Interworking Mode is the default mode. To switch to the Full Interworking Mode, use this procedure.

Important:

After you change this boot flag, you must save the change to the configuration file and reboot the switch.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Chassis** folders.
2. Click the **Boot Config** tab.
3. Select **EnableVxlanGwFullInterworkingMode** to enable Full Interworking Mode.
4. Click **Apply**.

Configuring a VTEP source IP address

Use the following procedure to configure the VXLAN Gateway source IP address.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Globals** tab.
4. In the **SourceIp** field, enter the VTEP source IP address.

5. In the **Vrf** field, select `GlobalRouter` or a VRF ID.
6. Click **Apply**.

Globals field descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Sourcelp	Specifies the VXLAN tunnel end point (VTEP) source IP address in an A.B.C.D format.
Vrf	<p>Specifies the VRF name used for the source IP address, which you select from the drop down menu. This IP address can be on the GRT or a VRF.</p> <p> Note: The VTEP source IP address must be on a loopback interface.</p>

Configuring a remote VTEP

Use the following procedure to configure the remote VTEP destination. The total number of VTEP remote destination IPs that you can configure is 500.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Vtep** tab.
4. In the **Id** field, assign a unique index value to this remote VTEP.
5. In the **IpAddr** field, enter the IP address of this remote VTEP.
6. In the **Name** field, assign a name to this remote VTEP.
7. Click **Apply**.

Vtep field descriptions

Use the data in the following table to use the **Vtep** tab.

Name	Description
Id	Specifies an index value in the range from 1 to 500 that uniquely identifies this remote VTEP.
IpAddr	Specifies an IP address that uniquely identifies this remote VTEP. <div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;">* Note:</div> <div>The remote VTEP IP address cannot be a local, broadcast, or multicast IP address.</div> </div>
Name	Assigns a name to this remote VTEP tunnel.
NextHopVrfName	Displays the VRF name of the next hop to reach the remote VTEP.
OvsdbConfigured	Specifies if the VTEP is configured with OVSDB protocol support for VXLAN Gateway.

Configuring a VNID

Use the following procedure to configure a VNID, which identifies one of the potential *16 million* VXLAN segments that can coexist within the same administrative domain.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Vnid** tab.
4. In the **Identifier** field, assign a unique value to this VNID.
5. In the **Isid** field, specify the I-SID associated with this VNID.
6. In the **Action** field, select one of the available actions: none or flush.
7. Click **Apply**.

Vnid field descriptions

Use the data in the following table to use the **Vnid** tab.

Name	Description
Identifier	Specifies a value that uniquely identifies the VNID of this VXLAN tunnel end point.

Table continues...

Name	Description
Isid	Specifies a value that uniquely identifies the I-SID associated with this VXLAN tunnel end point. * Note: A VNID must not have the same value as an I-SID.
Action	Choose one of the following actions to take with the MAC FDB: <ul style="list-style-type: none"> • none: Don't take any action. This is the default. • flushMacFdb: Flush all the learned MAC addresses from the forwarding database of the selected VNID.
OvsdbConfigured	Specifies if the VNID is configured with OVSDB protocol support for VXLAN Gateway.

Configuring VNID endpoints

Use the following procedure to configure a VNID and the VXLAN tunnel end point.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Vnid End Points** tab.
4. In the **Vnid** field, enter the VNID ID.
5. In the **VtepId** field, enter the ID of the remote VTEP.
6. Click **Apply**.

Vnid End Points field descriptions

Use the data in the following table to use the **Vnid End Points** tab.

Name	Description
Vnid	Specifies a value that uniquely identifies the VNID of this VXLAN tunnel end point.
VtepId	Specifies an index value in the range from 1 to 500 that uniquely identifies this remote VTEP.
Isid	Displays a value that uniquely identifies the I-SID associated with this VXLAN tunnel end point.

Configuring ELAN endpoints

Use the following procedure to configure information associated with ELAN end points.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Elan End Points** tab.
4. In the **Vnid** field, specify the VNID associated with this ELAN.
5. In the **Cvid** field, specify the customer VLAN ID.
6. In the **IfIndex** field, specify the interface index for this ELAN endpoint.
7. Click **Apply**.

Elan End Points field descriptions

Use the data in the following table to use the **Elan End Points** tab.

Name	Description
Vnid	Specifies a value that uniquely identifies the VNID of this ELAN tunnel end point.
Cvid	<p>Specifies a value that uniquely identifies the customer VLAN ID of this ELAN endpoint.</p> <p>The switch reserves the following VLAN IDs:</p> <ul style="list-style-type: none"> • VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • VLAN IDs 4060 to 4094 are reserved for internal use. • VLAN IDs 3500 to 3999 are reserved if you enable VRF scaling and SPBM mode. • VLAN ID 4095 is not used. • VLAN ID 4096 is reserved for untagged traffic.
IfIndex	Specifies the interface index for the MLT or port for this ELAN endpoint.
Isid	Specifies the I-SID associated with this VNID.

Displaying the VTEP next hop

Use the following procedure to display information about the remote VTEP.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Next Hop** tab.

Next Hop field descriptions

Use the data in the following table to use the **Next Hop** tab.

Name	Description
VtepId	Specifies an index value in the range from 1 to 500 that uniquely identifies this remote VTEP.
Ip	Specifies an IP address that uniquely identifies this remote VTEP tunnel.
IfIndex	Specifies the interface index of the next hop to reach the remote VTEP.
Vid	Specifies the VLAN ID of the next hop to reach remote VTEP.

Displaying the VNID forwarding database

Use the following procedure to display information about the VNID forwarding database (FDB).

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Vnid FDB** tab.

Vnid FDB field descriptions

Use the data in the following table to use the **Vnid FDB** tab.

Name	Description
Vnid	Displays the VNID IDs of the VXLAN tunnel endpoints.
Address	Displays all the learned MAC addresses in the forwarding database of the selected VNID.
Status	Displays whether or not the MAC address was learned.
InterfaceIndex	Displays the interface index name for the VNID.
Type	Displays whether the VNID is local or remote.

Chapter 8: OVSDB protocol support for VXLAN Gateway configuration using EDM

This section provides procedures to configure OVSDB protocol support for VXLAN Gateway using the Enterprise Device Manager (EDM).

Configuring OVSDB globally

Use the following procedure to configure the OVSDB certificates, private key, and replication global settings.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **OVSDB**.
3. Click the **Globals** tab.
4. Select **Enable**.
5. In the **CertificateFilename** field, type the `filepath/filename`.
6. In the **CertFileInstallAction** field, select **install**.
7. Click **Apply**.
8. In the **PrivateKeyFilename** field, type the `filepath/filename`.
9. In the **PrivateKeyInstallAction** field, select **install**.
10. Click **Apply**.
11. Select **ReplicationEnable**.

 **Note:**

OVSDB replication configuration is available for testing only. OVSDB replication is not supported in this release.

12. In the **ReplicationPeerIpAddr** field, type the IP address.

13. In the **ReplicationLocalIpAddr** field, type the IP address.
14. Click **Apply**.
15. Verify the **ReplicationState**.

Globals field descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Enable	Enables or disables the feature. Default is disabled (cleared).
CertificateFilename	Specifies the filename of the certificate.
CertFileInstallAction	Specifies to install or uninstall the certificate. You must have a certificate file on the switch.
PrivateKeyFilename	Specifies the filename of the private key.
PrivateKeyInstallAction	Specifies to install or uninstall the private key. You must have a private key file on the switch.
ReplicationEnable	Enables or disables replication support. Default is disabled (cleared).
ReplicationPeerIpAddr	Specifies the replication server peer IP address in A.B.C.D format.
ReplicationLocalIpAddr	Specifies the replication local IP address in A.B.C.D format.
ReplicationState	Displays the current replication status.

Configuring an OVSDB controller

Use the following procedure to configure an OVSDB controller ID, IP address, protocol, and port.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **OVSDB**.
3. Click the **Controller** tab.
4. Click **Insert**.
5. In the **Id** field, assign a unique index value to this controller.
6. In the **IpAddr** field, enter the IP address of this controller.
7. In the **Protocol** field, select **tcp** or **ssl**.

8. In the **Port** field, enter a port number.
9. Click **Insert**.

Controller field descriptions

Use the data in the following table to use the **Controller** tab.

Name	Description
Id	Specifies an ID number for a controller. Value range of 1 to 100.
IpAddr	Specifies the IP address of a controller. Type an IP address in A . B . C . D format.
Protocol	Specifies the communications protocol for the controller. Select tcp or ssl .
Port	Specifies the communications port for the controller. Value range of 1 to 65535.

Displaying the OVSDDB controller status

Use the following procedure to display status information about the OVSDDB controller.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **OSVDB**.
3. Click the **Status** tab.

Status field descriptions

Use the data in the following table to use the **Status** tab.

Name	Description
IpAddress	Displays the IP address of a controller.
Protocol	Displays the communications protocol of a controller as TCP or SSL.
Port	Displays the communications port for the controller. Value range of 1 to 65535.

Table continues...

Name	Description
Source	Displays the source of the controller as learned or configured.
OperStatus	Displays the operational status of the controller.

Configuring OVSDB managed interface

Use the following procedure to configure a I-SID interface for the OVSDB.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **OSVDB**.
3. Click the **Interface** tab.
4. Click **Insert**.
5. In the **Isid** field, type the I-SID identifier.
6. Click **Insert**.

Interface field descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
Isid	Specifies an I-SID to use as the managed interface.

Chapter 9: VXLAN Gateway configuration example

This example shows how to configure the VXLAN Gateway. Remember that there are two modes:

- *Base Interworking Mode* is the default mode so there is no need to set the mode boot flag. However, you should use the `show boot config flags` command to verify that the `vxlan-gw-full-interworking-mode` is set to `false`.
- *Full Interworking Mode* supports all SPB functionality including vIST and SMLT using an internal loopback. Full Interworking Mode is **not** the default mode. To configure Full Interworking mode, you must enable the `vxlan-gw-full-interworking-mode` boot flag.

! Important:

Changing the mode requires a reboot for the change to take effect, which can cause a loss of traffic.

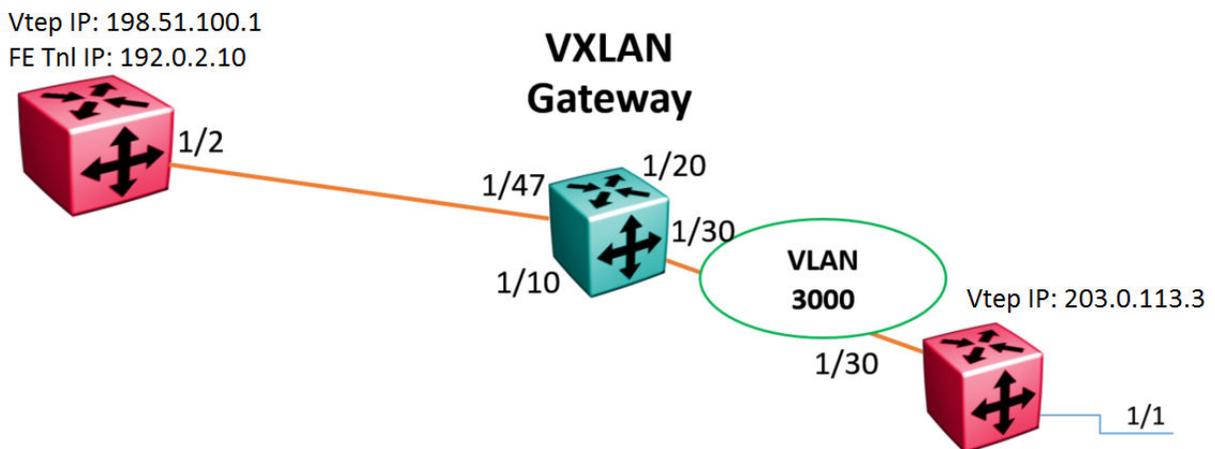


Figure 4: VXLAN Gateway configuration example

```
#  
#SPBM CONFIGURATION  
#  
spbm  
spbm ethertype 0x8100  
#  
#VLAN CONFIGURATION  
#  
vlan members remove 1 1/1,1/30,2/1/1,2/1/4
```

VXLAN Gateway configuration example

```
vlan create 3000 type port-mstprstp 0
vlan members 3000 1/30 portmember
interface Vlan 3000
ip address 192.168.30.1 255.255.255.0 0
ip ospf enable
exit

#
#PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 1/1
default-vlan-id 0
flex-uni enable

#
#CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
#
interface loopback 1
ip address 1 203.0.113.3/255.255.255.255
ip ospf 1
exit

#
#VTEP CONFIGURATION
#
vtep source-ip 198.51.100.1

#
#REMOTE VTEP CONFIGURATIONS
#
vtep 1 ip 203.0.113.3

#
#OSPF CONFIGURATION - GlobalRouter
#
router ospf enable
router ospf
exit

#
#IP REDISTRIBUTION CONFIGURATION - GlobalRouter
#
router ospf
redistribute direct
redistribute direct enable
exit

#
#I-SID CONFIGURATION
#
i-sid 300 elan
c-vid 10 port 1/1
exit

#
#VNID CONFIGURATION
#
vnid 100 i-sid 300
vtep 1
exit

#
#IP REDISTRIBUTE APPLY CONFIGURATIONS
#
ip ospf apply redistribute direct
```