

Configuring VXLAN Gateway for VOSS

Release 8.1.5 (VOSS) 9035889-01 Rev AA April 2020 © 2017-2020, Extreme Networks, Inc. All Rights Reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/ owners.

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

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: www.extremenetworks.com/support/ policies/software-licensing

Contents

| Chapter 1: About this Document | 5 |
|---|----|
| Purpose | 5 |
| Conventions | 5 |
| Text Conventions | 5 |
| Documentation and Training | 7 |
| Getting Help | 8 |
| Providing Feedback | 9 |
| Chapter 2: New in this Document | 10 |
| Notice about Feature Support | 10 |
| 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 Laver 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 | 23 |
| Chapter 5: VXLAN Gateway configuration using the CLI | 26 |
| Configuring the VXLAN Gateway boot flag | 26 |
| Configuring VXLAN Gateway | 27 |
| Flushing the MAC forwarding table | 29 |
| VXLAN Gateway show commands | 30 |
| Displaying VTEP source information | 30 |
| Displaying remote VTEP information | 30 |
| Displaying the name of the remote VTEP | 31 |
| Displaying VNID information | 31 |
| Displaying the MAC Addresses in the VNID FDB | 32 |
| Displaying the VXLAN Running Configuration | 32 |
| Chapter 6: OVSDB protocol support for VXLAN Gateway configuration using CLI | 34 |
| Configuring OVSDB Managed Interfaces | 34 |
| Configuring OVSDB protocol support for VXLAN Gateway | 35 |
| Configuring OVSDB Replication | 37 |
| Displaying the OVSDB configuration | 38 |
| Displaying the OVSDB controller status | 39 |
| Displaying the OVSDB replication state | 39 |
| Displaying the OVSDB managed interfaces | 40 |
| Chapter 7: VXLAN Gateway configuration using EDM | 41 |

| 41 |
|------|
| 41 |
| . 42 |
| . 43 |
| 44 |
| 45 |
| 46 |
| . 46 |
| 48 |
| . 48 |
| . 49 |
| . 50 |
| 51 |
| . 52 |
| 52 |
| 54 |
| - |

Chapter 1: About this Document

This section discusses the purpose of this document, the conventions used, ways to provide feedback, additional help, and information regarding other Extreme Networks publications.

Purpose

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

- VSP 7200 Series
- VSP 7400 Series
- 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.

Conventions

This section discusses the conventions used in this guide.

Text Conventions

The following tables list text conventions that can be used throughout this document.

Table 1: Notice Icons

| Icon | Alerts you to |
|------------|---|
| Important: | A situation that can cause serious inconvenience. |

Table continues...

| Icon | Alerts you to |
|------------------|---|
| Note: | Important features or instructions. |
| 🔁 Tip: | Helpful tips and notices for using the product. |
| A Danger: | Situations that will result in severe bodily injury; up to and including death. |
| Marning: | Risk of severe personal injury or critical loss of data. |
| Caution: | Risk of personal injury, system damage, or loss of data. |

Table 2: Text Conventions

| Convention | Description |
|------------------------|--|
| Angle brackets (< >) | Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command. |
| | <pre>If the command syntax is cfm maintenance- domain maintenance-level <0-7> , you can enter cfm maintenance-domain maintenance-level 4.</pre> |
| Bold text | Bold text indicates the GUI object name you must act upon. |
| | Examples: |
| | • Click OK . |
| | On the Tools menu, choose Options. |
| Braces ({ }) | Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command. |
| | For example, if the command syntax is ip address {A.B.C.D}, you must enter the IP address in dotted, decimal notation. |
| Brackets ([]) | Brackets ([]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command. |
| | For example, if the command syntax is show clock [detail], you can enter either show clock or show clock detail. |
| Ellipses () | An ellipsis () indicates that you repeat the last element of the command as needed. |
| | For example, if the command syntax is ethernet/2/1 [<parameter></parameter> |

Table continues...

| Convention | Description |
|---------------------|---|
| | <pre><value>], you enter ethernet/2/1 and as many parameter-value pairs as you need.</value></pre> |
| Italic Text | Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links. |
| Plain Courier Text | Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages. |
| | Examples: |
| | • show ip route |
| | • Error: Invalid command syntax [Failed][2013-03-22 13:37:03.303 -04:00] |
| Separator (>) | A greater than sign (>) shows separation in menu paths. |
| | For example, in the Navigation tree, expand the Configuration > Edit folders. |
| Vertical Line () | A vertical line () separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command. |
| | For example, if the command syntax is access- policy by-mac action { allow deny }, you enter either access-policy by-mac action allow or access-policy by-mac action deny, but not both. |

Documentation and Training

Find Extreme Networks product information at the following locations:

Current Product Documentation Release Notes Hardware/software compatibility matrices for Campus and Edge products Supported transceivers and cables for Data Center products Other resources, like white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit <u>www.extremenetworks.com/education/</u>.

Getting Help

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

- Extreme Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.
- **The Hub** A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- <u>Call GTAC</u> For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: <u>www.extremenetworks.com/support/contact</u>

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

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- · A description of any actions 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

Subscribe to Service Notifications

You can subscribe to email notifications for product and software release announcements, Vulnerability Notices, and Service Notifications.

- 1. Go to www.extremenetworks.com/support/service-notification-form.
- 2. Complete the form (all fields are required).
- 3. Select the products for which you would like to receive notifications.

😵 Note:

You can modify your product selections or unsubscribe at any time.

4. Select Submit.

Providing Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. 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 we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information in the document.
- Broken links or usability issues.

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- · Access the feedback form at https://www.extremenetworks.com/documentation-feedback/.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Chapter 2: New in this Document

There are no feature changes in this document.

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 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. Fabric interoperations

You can also configure Fabric gateway functionality like SPB-PIM Gateway and VXLAN Gateway.

4. Operations and Management

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

Documentation Sources

Refer to the following documentation sources:

• For information on basic SPBM infrastructure and IS-IS configuration and Layer 2 services, see <u>Configuring Fabric Basics and Layer 2 Services for VOSS</u>.

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 <u>Configuring Fabric Layer 3</u> <u>Services for VOSS</u>.
- For information on IP Multicast over Fabric Connect configuration and services, see <u>Configuring Fabric Multicast Services for VOSS</u>. This document also contains information

about configuring the SPB-PIM Gateway (SPB-PIM GW), which provides multicast interdomain 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 <u>Troubleshooting VOSS</u>.
- For information on VXLAN Gateway configuration, see Configuring VXLAN Gateway for VOSS.

Chapter 4: VXLAN Gateway Fundamentals

| Product | Release introduced |
|--------------------------------|--|
| guring VXLAN Gateway for VOSS. | |
| VSP 4450 Series | Not Supported |
| VSP 4900 Series | Not Supported |
| VSP 7200 Series | VOSS 6.0 |
| VSP 7400 Series | VOSS 8.0 |
| VSP 8200 Series | VOSS 6.0 |
| VSP 8400 Series | VOSS 6.0 |
| VSP 8600 Series | Not Supported |
| XA1400 Series | Not Supported |
| VSP 4450 Series | Not Supported |
| VSP 4900 Series | Not Supported |
| VSP 7200 Series | VOSS 7.1 |
| VSP 7400 Series | VOSS 8.0 |
| VSP 8200 Series | VOSS 7.1 |
| VSP 8400 Series | VOSS 7.1 |
| VSP 8600 Series | Not Supported |
| XA1400 Series | Not Supported |
| VSP 4450 Series | Not Supported |
| VSP 4900 Series | Not Supported |
| VSP 7200 Series | VOSS 6.0 |
| VSP 7400 Series | VOSS 8.0 |
| VSP 8200 Series | VOSS 6.0 |
| VSP 8400 Series | VOSS 6.0 |
| VSP 8600 Series | Not Supported |
| XA1400 Series | Not Supported |
| | Product juring VXLAN Gateway for VOSS. VSP 4450 Series VSP 4900 Series VSP 7200 Series VSP 7400 Series VSP 8200 Series VSP 8400 Series VSP 8400 Series VSP 8400 Series VSP 8400 Series VSP 8600 Series VSP 4450 Series VSP 4900 Series VSP 7200 Series VSP 7400 Series VSP 7400 Series VSP 7400 Series VSP 7400 Series VSP 8200 Series VSP 8400 Series VSP 4450 Series VSP 4450 Series VSP 4450 Series VSP 4450 Series VSP 7200 Series VSP 7400 Series VSP 8400 Series VSP 8400 Series VSP 8600 Series |

Table 3: VXLAN Gateway product support

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.

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 <u>Release Notes for VOSS</u>.

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.

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.

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

Table 4: VXLAN Gateway Management Methods

| Function | Static management | Dynamic management |
|---|--|--|
| 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 VOSS.

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 **vxlan-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 Layer 3 VSN Fabric Connect traffic cannot be carried over into a VXLAN domain through a VXLAN Gateway. Layer 3 VSN Fabric Connect requires MAC-in-MAC encapsulation over VXLAN encapsulation, and the VXLAN Gateway data plane does not support double-encapsulation. However, IP Multicast over IP Shortcuts Fabric Connect can 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 <u>Configuring QoS and ACL-Based Traffic Filtering for VOSS</u>.
- 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).

😵 Note:

VMware NSX 6.2.4 has known issues that can cause VXLAN discovery delays and NVC connectivity issues in certain scenarios, for more information see <u>Release Notes for VOSS</u>.

- 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 loadbalancing.
- The NVC performs MAC learning on the VXLAN tunnel.
- ARP supression 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.

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 use the Segmented Management InstanceIP address to connect with the NVCs. For more information about configuring Segmented Management Instance, see <u>Administering VOSS</u>.
- 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 is different than NSX service node replication. Service node replication is configured and supported in NSX. OVSDB replication uses vIST and is configured and supported on the VXLAN Gateway.

- You require at least two NVCs to enable OVSDB replication.
- OVSDB replication operates in an ACTIVE-STANDBY server configuration.
- You must disable OVSDB on the standby vIST device before disabling OVSDB on the active vIST device.
- One active NVC and one standby NVC are supported by the VXLAN Gateway.
- 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 indicates that Full Interworking Mode is enabled.

Example

Switch:1>enable

```
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flags vxlan-gw-full-interworking-mode
```

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>

Example

For an example of this procedure, see <u>VXLAN Gateway configuration example in Base Interworking</u> <u>mode</u> on page 52 or <u>VXLAN Gateway configuration example in Full Interworking mode</u> on page 52.

Variable definitions

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

| Variable | Value |
|--|---|
| source-ip <a.b.c.d> [vrf WORD<1–16>]</a.b.c.d> | 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>]</a.b.c.d> | 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=""></ip> |
| | 😵 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. |
| | The switch reserves the following VLAN IDs: |
| | 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. |
| c-vid <1-4094> mlt <1-4094> | Specifies a value that uniquely identifies the customer VLAN ID and MLTs of this ELAN end point. |
| | The switch reserves the following VLAN IDs: |
| | 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. |
| untagged-traffic port <{slot/port[/sub-port] [-slot/port[/ sub-port]] [,]}> | Specifies the ports that support untagged traffic. |
| untagged-traffic mlt <1-4094> | Specifies the MLTs that support untagged traffic. |
| vtep <1-255> | 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

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

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

Example

| Swit | ch:1# show vt | ep remote | | | |
|--------------------------------------|-------------------|-------------------------------|-------------------------|--------------------|--------------------------------|
| Remote VTEP Info | | | | | |
| ID | NAME | VTEP I DEST-IP I | L3_VTEP_NEX PORT/MLT | T_HOP_INFO VLAN | VRF |
| 1 56 | vtep-1 vtep-56 | 198.51.100.6 198.51.100.22 | Mlt78 Mlt130 | 3851 3855 | GlobalRouter 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 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 11000060 11000065 11000066 | 120 1060 1065 1066 | c3:1/10 c60:1/10 c65:1/10 c66:1/48 | | 2,56 56 56 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> | remote]

Example

| Switch:1# show vnid mac-address-entry | | | | |
|--|--|--|---|--------------------------------------|
| | | VNID Fdb Table | | |
| VNID | STATUS | MAC-ADDRESS | INTERFACE | TYPE |
| 11000065 14000100 14000100 14000100 | learned learned learned learned | 00:00:5e:00:01:01 a4:25:1b:51:a5:04 a4:25:1b:52:a5:04 a4:25:1b:52:bd:03 | vtep_56 vtep_67 vtep_67 vtep_2 | REMOTE REMOTE REMOTE REMOTE |
| c: customer | vid u: ur | ntagged-traffic | | |
| All 4 out of | 4 Total Nu | um 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 <u>Configuring the VXLAN Gateway boot flag</u> on page 26.
- You must create the required Flex UNI or CVLAN I-SIDs before you configure the OVSDB managed interface.

Procedure

1. Enter OVSDB Configuration mode:

enable configure terminal

ovsdb

2. Configure the OVSDB managed interface:

```
managed-interface i-sids WORD<1-1024>
```

Example

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 <u>Configuring OVSDB protocol</u> <u>support for VXLAN Gateway</u> on page 35.

Variable definitions

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

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

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 <u>Configuring the</u> VXLAN Gateway boot flag on page 26.
- You must configure and use the Segmented Management Instance IP address on the VXLAN Gateway to establish connectivity with the NVC. For more information about Segmented Management Instance, see <u>Administering VOSS</u>.
- You must configure an OVSDB management interface. For more information, see <u>Configuring</u> <u>OVSDB Managed Interfaces</u> on page 34.
- 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. Configure the source VTEP IP address:

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

3. Enter OVSDB Configuration mode:

ovsdb

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>]
```

Example

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.

If you change a previously configured VTEP source-ip and re-enable OVSDB, the controller sees a new VXLAN tunnel instead of updating the existing VXLAN. You must configure the VNID to I-SID binding on the controller for the new VXLAN tunnel associated with the new VTEP IP address.

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>]</a.b.c.d> | 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></a.b.c.d> | Specifies the IP address of the controller in IPv4 format. |
| protocol <ssl tcp></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.

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 OVSDB Configuration mode:

enable configure terminal

ovsdb

2. Configure the OVSDB replication:

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

Example

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></a.b.c.d> | Specifies the IP address of the secondary controller for OVSDB replication support. |
| local-ip <a.b.c.d></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. To enter User EXEC mode, log on to the switch.
- 2. Display the OVSDB configuration:

show ovsdb config

Example

The following is an example of displaying the OVSDB configuration.

 SWITCH:1>show ovsdb config

 Ovsdb 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

 Ovsdb Controller Config Info

 ID
 Protocol

 IP
 Port

 1
 ssl
 192.0.2.10

Displaying the OVSDB controller status

Use the following procedure to display the OVSDB controller status.

Procedure

- 1. To enter User EXEC mode, log on to the switch.
- 2. Display the OVSDB controller status:

show ovsdb controller status

Example

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

SWITCH:1>show ovsdb controller status

| | | OVSDB Contro | oller Status | |
|------------|------------------------|--------------|--------------|------------------|
| Protocol | IP | Port | status | source |
| ssl ssl | 192.0.2.2 192.0.2.3 | 6640 6640 | Up Down | Learnt 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. To enter User EXEC mode, log on to the switch.

2. Display the OVSDB controller status:

show ovsdb replication state

Example

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

| SWITCH:1>show ov | sdb replication state | | |
|------------------|-----------------------|----------------|--|
| | Ovsdb Repl | ication 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. To enter User EXEC mode, log on to the switch.
- 2. Display the OVSDB managed interfaces:

show ovsdb managed-interface

Example

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

SWITCH:1>show ovsdb 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 **Sourcelp** 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 | 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. |
| | 😵 Note: |
| | The VTEP source IP address must be on a loopback interface. |

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. |
| lpAddr | Specifies an IP address that uniquely identifies this remote VTEP. |
| | 😵 Note: |
| | The remote VTEP IP address cannot be a local, broadcast, or multicast IP address. |
| 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: |
| | 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. |
| Vtepld | 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 | Specifies a value that uniquely identifies the customer VLAN ID of this ELAN endpoint. |
| | The switch reserves the following VLAN IDs: |
| | 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. |
| lfIndex | Specifies the interface index for the MLT or port for this ELAN endpoint. |
| lsid | 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 |
|---------|--|
| Vtepld | Specifies an index value in the range from 1 to 500 that uniquely identifies this remote VTEP. |
| lp | Specifies an IP address that uniquely identifies this remote VTEP tunnel. |
| lfindex | 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. |
| Туре | 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.
- 12. In the ReplicationPeerlpAddr field, type the IP address.
- 13. In the ReplicationLocallpAddr 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. |
| ReplicationLocallpAddr | 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. |
| lpAddr | 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 OVSDB controller status

Use the following procedure to display status information about the OVSDB 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. |
| 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: Configuration Examples

This section provides network configuration examples for VXLAN Gateway and OVSDB protocol support for VXLAN Gateway.

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
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
```

OVSDB protocol support for VXLAN Gateway configuration example

This example shows how to configure OVSDB protocol support for VXLAN Gateway:

Important:

OVSDB protocol support for VXLAN Gateway also requires an ESXi host running VMware NSX configured with a HW-VTEP, and a Logical Switch configured with VNID to HW-VTEP bindings to function. See VMware NSX documentation for more information.

Important:

OVSDB requires full interworking mode on the VXLAN Gateway. Changing the mode requires a reboot for the change to take effect, which can cause a loss of traffic.

```
#VXLAN GATEWAY UNDERLAY IP NETWORK CONFIGURATION HW-VTEP 1
enable
config terminal
ip vrf underlay-vxlan vrfid 1
router vrf underlay-vxlan
ip ospf
ip osfp admin-state
ip osfp router-id 203.0.113.1
vlan create 10 type port-mstprstp 1
vlan mlt 10 10
interface vlan 10
vrf underlay-vxlan
ip address 198.51.100.122 255.255.255.0
ip ospf enable
exit
#VXLAN GATEWAY UNDERLAY IP NETWORK CONFIGURATION HW-VTEP 2
ip vrf underlay-vxlan vrfid 1
router vrf underlay-vxlan
ip ospf
ip osfp admin-state
ip osfp router-id 203.0.113.2
vlan create 12 type port-mstprstp 1
vlan members add 12 3/1 portmember
interface vlan 12
vrf underlay-vxlan
```

```
ip address 198.51.100.124 255.255.255.0
ip ospf enable
exit
#NLS and VTEP CONFIGURATIONS HW-VTEP 1
vlan create 4059 type port-mstprstp 0
vlan members 4059 1/24 portmember
mgmt vlan 4059
ip address 192.0.2.101/24
ip route 192.0.2.0/16 next-hop 192.0.2.1 weight 1
enable
exit
interface loopback 10
ip address 10 122.122.122.1/255.255.255.255 vrf underlay-vxlan
ip ospf 10 vrf underlay-vxlan
vtep source-ip 122.122.122.1 vrf underlay-vxlan
#NLS and VTEP CONFIGURATIONS HW-VTEP 2
vlan create 4059 type port-mstprstp 0
vlan members 4059 3/24 portmember
mgmt vlan 4059
ip address 192.0.2.102/24
ip route 192.0.2.0/16 next-hop 192.0.2.1 weight 1
enable
exit
interface loopback 10
ip address 10 124.124.124.124/255.255.255.255 vrf underlay-vxlan
ip ospf 10 vrf underlay-vxlan
vtep source-ip 124.124.124.1 vrf underlay-vxlan
#LAYER2 VNI SERVICE AND MANAGED-INTERFACE CONFIGURATIONS HW-VTEP 1
vlan create 1001 type port-mstprstp 0
vlan members 1001 1/10 portmember
vlan i-sid 1001 1001
ovsdb
managed-interface i-sids 1001
exit
##LAYER2 VNI SERVICE AND MANAGED-INTERFACE CONFIGURATIONS HW-VTEP 2
vlan create 1001 type port-mstprstp 0
vlan members 1001 3/10 portmember
vlan i-sid 1001 1001
ovsdb
managed-interface i-sids 1001
exit
#OVSDB CERTIFICATE AND PRIVATE KEY CONFIGURATION HW-VTEP 1
ovsdb
install-cert-file /intflash/tom/vtep1-cert.pem
private-key /intflash/tom/vtep1-privkey.pem
```

Configuration Examples

```
enable
exit
#OVSDB NETWORK VIRTUALIZATION CONTROLLER CONFIGURATION HW-VTEP 1
ovsdb
controller 1 ip address 192.0.2.2 protocol ssl port 6640
exit
#OVSDB CERTIFICATE AND PRIVATE KEY CONFIGURATION HW-VTEP 2
#
ovsdb
install-cert-file /intflash/tom/vtep2-cert.pem
private-key /intflash/tom/vtep2-privkey.pem
enable
exit
#OVSDB NETWORK VIRTUALIZATION CONTROLLER CONFIGURATION HW-VTEP 2
ovsdb
controller 1 ip address 192.0.2.2 protocol ssl port 6640
exit
#If HW-VTEP 1 needs VXLAN Gateway redundancy, the redundant gateway must be a VIST
peer(HW-VTEP 11).
#For VXLAN Gateway redundancy, the configs at HW-VTEP-1 must also have ovsdb replication
config:
#OVSDB NETWORK VIRTUALIZATION CONTROLLER REPLICATION CONFIGURATION
#
ovsdb
replication peer-ip 192.0.2.101 local-ip 192.0.2.111
exit.
#Configs at HW-VTEP-11 for reduncancy:
#VXLAN GATEWAY UNDERLAY IP NETWORK CONFIGURATION HW-VTEP 11
enable
config terminal
ip vrf underlay-vxlan vrfid 1
router vrf underlay-vxlan
ip ospf
ip osfp admin-state
ip osfp router-id 203.0.113.11
vlan create 10 type port-mstprstp 1
vlan mlt 10 10
interface vlan 10
vrf underlay-vlxan
ip address 198.51.100.111 255.255.255.0
ip ospf enable
exit
#NLS and VTEP CONFIGURATIONS HW-VTEP 11
#NOTE: HW-VTEP 1 and HW-VTEP 11's VTEP IP must be the same.
```

```
#
vlan create 4059 type port-mstprstp 0
vlan members 4059 1/24 portmember
mgmt vlan 4059
ip address 192.0.2.111/24
ip route 192.0.2.0/16 next-hop 192.0.2.1 weight 1
enable
exit
interface loopback 10
ip address 10 122.122.122.1/255.255.255.255 vrf underlay-vxlan
ip ospf 10 vrf underlay-vxlan
vtep source-ip 122.122.122.1 vrf underlay-vxlan
#LAYER2 VNI SERVICE AND MANAGED-INTERFACE CONFIGURATIONS HW-VTEP 11
#NOTE: HW-VTEP 1 and HW-VTEP 11 must have same managed-interface configured under 'ovsdb'
vlan create 1001 type port-mstprstp 0
vlan members 1001 1/10 portmember
vlan i-sid 1001 1001
ovsdb
managed-interface i-sids 1001
exit
#OVSDB CERTIFICATE AND PRIVATE KEY CONFIGURATION HW-VTEP 11
#NOTE: HW-VTEP 1 and HW-VTEP 11 must have same certificate and private keys configured.
ovsdb
install-cert-file /intflash/tom/vtep1-cert.pem
private-key /intflash/tom/vtep1-privkey.pem
enable
exit
#OVSDB NETWORK VIRTUALIZATION CONTROLLER CONFIGURATION HW-VTEP 11
ovsdb
controller 1 ip address 192.0.2.2 protocol ssl port 6640
exit
#OVSDB NETWORK VIRTUALIZATION CONTROLLER REPLICATION CONFIGURATION
ovsdb
replication peer-ip 192.0.2.101 local-ip 192.0.2.111
exit
```