



CHAPTER 2

Configuring VXLAN

This chapter describes how to configure the Virtual Extensible Local Area Network (VXLAN).

This chapter includes the following topics:

- [Information About VXLAN, page 2-1](#)
- [Prerequisites for VXLAN, page 2-1](#)
- [Default Settings, page 2-2](#)
- [Configuring VXLAN, page 2-2](#)
- [Verifying VXLAN Configuration, page 2-12](#)
- [Feature History for VXLAN, page 2-14](#)

Information About VXLAN

For detailed information about VXLAN, see [Chapter 1, “Overview”](#).

Prerequisites for VXLAN

VXLAN has the following prerequisites:

- The Cisco Nexus 1000V uplink port profiles and all interconnecting switches/routers in between the ESX hosts must have their supported MTU set to at least 50 bytes larger than the MTU of the VMs. For example, the VMs default to using a 1500 byte MTU (same as the uplinks and physical devices), so in this case they must be set to at least 1550 bytes. If this isn't possible, then all VM's VNICs should have their MTU lowered to be 50 bytes smaller than what the physical network supports, for example 1450 bytes. For more information, see the *Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.2(1)SV1(5.1)*.
- If the Cisco Nexus 1000V is using a port channel for its uplinks, then the load distribution algorithm should be set to use a 5-tuple hash (IP/L4/L4 Ports). The same should be used for any port channels on the physical switches. For more information, see the *Cisco Nexus 1000V Interface Configuration Guide, Release 4.2(1)SV1(5.1)*.
- If VEMs requiring VXLAN connectivity are separated by a router
 - Proxy ARP must be enabled on the SVIs connected to the Cisco Nexus 1000V's VXLAN transport VLANs (the ones the “capability vxlan” port profiles are connected to).
 - Multicast routing must be enabled on the routers.

Send document comments to nexus1k-docfeedback@cisco.com.

- VXLAN makes use of MAC in IP (UDP) with a destination port of 8472. You must allow this through any firewall.
- Your upstream switch, from the VEMs of the Cisco Nexus 1000V, needs to provide an IGMP querier function.

Default Settings

Table 2-1 lists the default settings for VXLAN parameters.

Table 2-1 Default VXLAN Parameters

Parameters	Default
VXLAN	Disabled

Configuring VXLAN

This section includes the following topics:

- [Initial Enabling of VXLANs, page 2-2](#)
- [Creating a VXLAN, page 2-5](#)
- [Creating a Port Profile Configured to Use a VXLAN, page 2-6](#)
- [Removing Ports from a VXLAN, page 2-8](#)
- [Deleting a VXLAN, page 2-9](#)
- [Disabling Segmentation, page 2-10](#)

Initial Enabling of VXLANs

To enable a VXLAN, you must to perform the following two procedures when first configuring VXLAN.

- [Configuring vmknics for VXLAN Encapsulation, page 2-2](#)
- [Enabling VXLANs, page 2-4](#)

Configuring vmknics for VXLAN Encapsulation

You can configure vmknics for VXLAN encapsulation by running the following procedure.

BEFORE YOU BEGIN

- Identify a VLAN to be used for transporting VXLAN encapsulated traffic.
- Ensure it is configured on the uplink port profile for all VEMs on which VXLAN can be configured.

SUMMARY STEPS

1. **configure terminal**
2. **port-profile** *profilename*

Send document comments to nexus1k-docfeedback@cisco.com.

3. **vmware port-group** *name*
4. **switchport mode access**
5. **switchport access vlan** *id*
6. **capability vxlan**
7. **no shutdown**
8. **state enabled**
9. **show port-profile name** *profilename*
10. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	port-profile <i>profilename</i> Example: switch(config)# port-profile vmknic-pp switch(config-port-prof)#	Enters port profile configuration mode for the named port profile. If the port profile does not already exist, it is created using the following characteristics: <ul style="list-style-type: none"> • profilename—The port profile name can be up to 80 characters and must be unique for each port profile on the Cisco Nexus 1000V. Note If a port profile is configured as an Ethernet type, it cannot be used to configure VMware virtual ports.
Step 3	vmware port-group <i>name</i> Example: switch(config-port-prof)# vmware port-group switch(config-port-prof)#	Designates the port profile as a VMware port group. The port profile is mapped to a VMware port group of the same name unless you specify a name here. When you connect the VSM to vCenter Server, the port group is distributed to the virtual switch on the vCenter Server.
Step 4	switchport mode access Example: switch(config-port-prof)# switchport mode access switch(config-port-prof)#	Designates the interfaces as switch access ports (the default).
Step 5	switchport access vlan <i>id</i> Example: switch(config-port-prof)# switchport access vlan 100 switch(config-port-prof)#	Assigns a VLAN ID to this port profile.

Send document comments to nexus1k-docfeedback@cisco.com.

	Command	Purpose
Step 6	capability vxlan Example: switch(config-port-prof)# capability vxlan switch(config-port-prof)	Assigns the VXLAN capability to the port profile to ensure that the interfaces that inherit this port profile are used as sources for VXLAN encapsulated traffic.
Step 7	no shutdown Example: switch(config-port-prof)# no shutdown switch(config-port-prof)	Administratively enables all ports in the profile.
Step 8	state enabled Example: switch(config-port-prof)# state enabled switch(config-port-prof)	Sets the operational state of a port profile.
Step 9	show port-profile name profilename Example: switch# show port-profile vmknic-pp	(Optional) Displays the port profile configuration.
Step 10	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

What to Do Next

- The vSphere administrator must create a new vmknic on each ESX/ESXi host and assign the previously created port profile to this vmknic.

Enabling VXLANs

You can enable VXLANs by performing the following procedure.

BEFORE YOU BEGIN

- Enter the **show system vem feature level** command to confirm that the feature level is 4.2(1)SV1(5.1) or later. If the feature level is not 4.2(1)SV1(5.1) or later, see the *Cisco Nexus 1000V Software Upgrade Guide, Release 4.2(1)SV1(5.1)*.

SUMMARY STEPS

- configure terminal
- feature segmentation
- show feature | grep segmentation
- show processes | grep seg_bd
- copy running-config startup-config

Send document comments to nexus1k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	feature segmentation Example: switch(config)# feature segmentation switch(config)	Enables the VXLAN feature.
Step 3	show feature grep segmentation Example: switch# show feature grep segmentation	(Optional) Displays if the VXLAN feature is enabled.
Step 4	show processes grep seg_bd Example: switch# show processes grep seg_bd	(Optional) Displays if the VXLAN process is running.
Step 5	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

EXAMPLES

The following example shows enabling the segmentation feature.

```
n1000v# configure terminal
n1000V(config)# feature segmentation
n1000v(config)# show feature | grep segmentation
network-segmentation 1 disabled
segmentation          1 enabled
n1000v(config)# show processes | grep seg_bd
4166   S   b7de9468      1      - seg_bd
n1000v(config)# copy running-config startup-config
```

Creating a VXLAN

You can create a VXLAN by running the following procedure.

RESTRICTIONS

- You are limited to creating a combination of 2048 VXLANs and VLANs.

SUMMARY STEPS

- configure terminal**
- bridge-domain** *name-string*
- segment id** [*number*]
- group** *ipaddr*

Send document comments to nexus1k-docfeedback@cisco.com.

5. **show bridge-domain** *name-string*
6. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	bridge domain <i>name-string</i> Example: switch(config)# bridge-domain tenant-red switch(config-bd)#	Creates a VXLAN and associates an identifying name to it.
Step 3	segment id [<i>number</i>] Example: switch(config-bd)# segment id 20480 switch(config-bd)#	Specifies the VXLAN Segment ID. Only one Bridge Domain can use a particular segment id value. Valid values are 4096 to 16777215. (1 - 4095 are reserved for VLANs.)
Step 4	group <i>ipaddr</i> Example: switch(config-bd)# group 239.1.1.1 switch(config-bd)#	Associates the multicast group for broadcasts and floods. Note Reserved multicast addresses are not allowed.
Step 5	show bridge-domain <i>name-string</i> Example: switch# show bridge-domain tenant-red switch(config-bd)#	(Optional)
Step 6	copy running-config startup-config Example: switch(config-bd)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Creating a Port Profile Configured to Use a VXLAN

You can create a port profile that is configured to use a VXLAN.

RESTRICTIONS

- Alternatively, you can associate ports with a bridge domain by modifying the configuration of an existing vEthernet port profile to use VXLANs instead of VLANs. To do so, enter the **switchport access bridge-domain name** command on a profile with **switchport mode access** configured.

SUMMARY STEPS

1. **configure terminal**
2. **port-profile** *profilename*

Send document comments to nexus1k-docfeedback@cisco.com.

3. **vmware port-group** *name*
4. **switchport mode access**
5. **switchport access bridge-domain** *name-string*
6. **no shutdown**
7. **state enabled**
8. **show port-profile name** *profilename*
9. **show running-config bridge-domain**
10. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	port-profile <i>profilename</i> Example: switch(config)# port-profile tenant-profile switch(config-port-prof)	Enters port profile configuration mode for the named port profile. If the port profile does not already exist, it is created using the following characteristics: <ul style="list-style-type: none"> • profilename—The port profile name can be up to 80 characters and must be unique for each port profile on the Cisco Nexus 1000V. Note If a port profile is configured as an Ethernet type, then it cannot be used to configure VMware virtual ports.
Step 3	vmware port-group <i>name</i> Example: switch(config-port-prof)# vmware port-group switch(config-port-prof)#	Designates the port profile as a VMware port group. The port profile is mapped to a VMware port group of the same name unless you specify a name here. When you connect the VSM to vCenter Server, the port group is distributed to the virtual switch on the vCenter Server.
Step 4	switchport mode access Example: switch(config-port-prof)# switchport mode access switch(config-port-prof)	Designates the interfaces as switch access ports (the default).
Step 5	switchport access bridge-domain <i>name-string</i> Example: switch(config-port-prof)# switchport access bridge-domain tenant-red switch(config-port-prof)	Assigns a VXLAN bridge domain to this port profile.

Send document comments to nexus1k-docfeedback@cisco.com.

	Command	Purpose
Step 6	no shutdown Example: switch(config-port-prof)# no shutdown switch(config-port-prof)#	Administratively enables all ports in the profile.
Step 7	state enabled Example: switch(config-port-prof)# state enabled switch(config-port-prof)	Sets the operational state of a port profile.
Step 8	show port-profile name <i>profilename</i> Example: switch(config-port-prof) # show port-profile name tenant-profile	(Optional) Displays the configuration of a port profile.
Step 9	show running-config bridge-domain Example: switch(config-port-prof) # show running-config bridge-domain	(Optional) Displays the segmentation configuration.
Step 10	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Removing Ports from a VXLAN

You can remove ports from a VXLAN by executing the following procedure.

RESTRICTIONS

- Executing this procedure moves the ports to the default VLAN.

SUMMARY STEPS

1. **configure terminal**
2. **port-profile *name***
3. **no switchport access bridge-domain**
4. **show port-profile usage**
5. **show bridge-domain *name***
6. **copy running-config startup-config**

Send document comments to nexus1k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	port-profile name Example: switch(config)# port-profile tenant-profile switch(config-port-prof)	Enters port profile configuration mode for the named port profile. If the port profile does not already exist, it is created using the following characteristics: <ul style="list-style-type: none"> <i>name</i>—The port profile name can be up to 80 characters and must be unique for each port profile on the Cisco Nexus 1000V. Note If a port profile is configured as an Ethernet type, then it cannot be used to configure VMware virtual ports.
Step 3	no switchport access bridge-domain Example: switch(config-port-prof)# no switchport access bridge-domain tenant-red switch(config-port-prof)	Removes the VXLAN bridge domain from this port profile.
Step 4	show port-profile usage Example: switch# show port-profile usage	(Optional) Displays a list of interfaces that inherited a port profile.
Step 5	show bridge-domain Example: switch# show bridge-domain	(Optional) Displays all bridge domains.
Step 6	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Deleting a VXLAN

You can delete a VXLAN domain by executing the following procedure.

RESTRICTIONS

- Deleting an existing bridge domain with ports on it moves all the ports to a **down** state. Traffic stops flowing.

SUMMARY STEPS

- configure terminal**
- no bridge-domain name-string**
- show bridge-domain**

Send document comments to nexus1k-docfeedback@cisco.com.

4. `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	<code>configure terminal</code> Example: switch# <code>configure terminal</code> switch(config)#	Enters global configuration mode.
Step 2	<code>no bridge-domain name-string</code> Example: switch(config)# <code>no bridge-domain group-red</code> switch(config-bd)	Deletes a VXLAN.
Step 3	<code>show bridge-domain</code> Example: switch# <code>show bridge-domain</code>	(Optional) Displays all bridge domains.
Step 4	<code>copy running-config startup-config</code> Example: switch# <code>copy running-config startup-config</code>	(Optional) Copies the running configuration to the startup configuration.

Disabling Segmentation

You can disable segmentation by executing the following procedure.

SUMMARY STEPS

1. `configure terminal`
2. `show bridge-domain`
3. `show running-config port-profile`
4. `port-profile name`
5. `no switchport access bridge-domain name-string`
6. `show port-profile usage`
7. `show bridge-domain name`
8. `no feature segmentation`
9. `show processes | grep seg_bd`
10. `copy running-config startup-config`

Send document comments to nexus1k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	show bridge-domain Example: switch(config)# show bridge-domain switch(config)#	Displays all bridge domains. Note You must identify all bridge domains with non-zero port counts.
Step 3	show running-config port-profile Example: switch(config)# show running port-profile	Displays the running configuration for all port-profiles. Note You must use this command to identify which port profiles have bridge domains identified in Step 2 configured.
Step 4	port-profile name Example: switch(config)# port-profile tenant-profile switch(config-port-prof)	Names the port profile and enters port profile configuration mode. If the port profile does not already exist, it is created using the following characteristics: <ul style="list-style-type: none"> <i>name</i>—The port profile name can be up to 80 characters and must be unique for each port profile on the Cisco Nexus 1000V. Note If a port profile is configured as an Ethernet type, then it cannot be used to configure VMware virtual ports.
Step 5	no switchport access bridge-domain name-string Example: switch(config-port-prof)# no switchport access bridge-domain tenant-red switch(config-port-prof)	Removes the VXLAN bridge domain from this port profile.
Step 6	show port-profile usage Example: switch# show port-profile usage	(Optional) Displays a list of interfaces that inherited a port profile.
Step 7	show bridge-domain Example: switch# show bridge-domain	(Optional) Displays all bridge domains.
Step 8	no feature segmentation Example: switch(config)# no feature segmentation switch(config)#	Removes the segmentation feature.

Send document comments to nexus1k-docfeedback@cisco.com.

	Command	Purpose
Step 9	show processes grep seg_bd Example: switch(config)# show processes grep seg_bd switch(config)#	Displays the processes to determine that the segmentation feature is not running.
Step 10	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Verifying VXLAN Configuration

To display VXLAN configuration information, enter one of the following commands:

Command	Purpose
show processes grep seg_bd	Displays that the VXLAN process is running.
show bridge-domain	Displays all bridge domains.
show interface brief	Displays a short version of the interface configuration.
show interface switchport	Displays information about switchport interfaces.

EXAMPLES

This example shows how to display if the VXLAN process is running.

```
switch (config)# show processes | grep seg_bd
-      NR      -          1      - seg_bd
```

This example shows how to display all bridge domains.

```
switch (config)# show bridge-domain

Bridge-domain tenant-red (2 port in all)
Segment ID: 5000 (manual/Active)
Group IP: 239.1.1.1
-      NR      -          1      - seg_bd
```

This example shows how to display a short version of the interface table.

```
switch(config)# show interface brief

-----
Port      VRF      Status  IP Address                               Speed      MTU
-----
mgmt 0    --      up      172.23.233.117                           1000       1500

-----
Ethernet  VLAN    Type Mode  Status Reason      Speed      Port
Interface                                     Ch #
-----
Eth3/5    1       eth trunk up      none       1000
```

Send document comments to nexus1k-docfeedback@cisco.com.

```

Vethernet  VLAN  Type Mode  Status Reason      Speed
-----
Veth1      --   virt access up    none      auto
Veth1      --   virt access up    none      auto
Veth1      100  virt access up    none      auto

-----

Port      VRF      Status  IP Address      Speed  MTU
control0  --       up      --              1000  1500
switch#(config)#

```

This example shows how to display information about switchport interfaces.

```

switch#(config)# show int switchport
Name: Ethernet3/5
  Switchport: Enabled
  Switchport Monitor: Not enabled
  Operational Mode: Trunk
  Access Mode VLAN: 1 (default)
  Trunking Native Mode: trunk
  Trunking VLANs Enabled: 180-181,231-233,571-574
  Administrative private-vlan primary host-association: none
  Administrative private-vlan secondary host-association: none
  Administrative private-vlan primary mapping: none
  Administrative private-vlan secondary mapping: none
  Administrative private-vlan trunk native VLAN: none
  Administrative private-vlan trunk encapsulation: dot1q
  Administrative private-vlan trunk normal VLANs: none
  Administrative private-vlan trunk private VLANs:
  Operational private-vlan: none

ifindex 0x1c000000 swbd 4096
Name Vethernet1
  Switchport: Enabled
  Switchport Monitor: Not enabled
  Operational Mode: access
  Access Mode VLAN: 0 (none)
  Access BD name: tenant-red
  Trunking Native ModeVLAN: 1 (default)
  Trunking VLANs Enabled: 1-3967,4048-4093
  Administrative private-vlan primary host-association: none
  Administrative private-vlan secondary host-association: none
  Administrative private-vlan primary mapping: none
  Administrative private-vlan secondary mapping: none
  Administrative private-vlan trunk native VLAN: none
  Administrative private-vlan trunk encapsulation: dot1q
  Administrative private-vlan trunk normal VLANs: none
  Administrative private-vlan trunk private VLANs:
  Operational private-vlan: none

```

For detailed information about the fields in the output from these commands, refer to the *Cisco Nexus 1000V Command Reference, Release 4.2(1)SV1(5.1)*.

Send document comments to nexus1k-docfeedback@cisco.com.

Feature History for VXLAN

Table 2-2 lists the release history for this feature. Only features that were introduced or modified in Release 4.2(1)SV1(5.1) or a later release appear in the table.

Table 2-2 ***Feature History for VXLAN***

Feature Name	Releases	Feature Information
VXLAN	4.2(1)SV1(5.1)	Introduced the Virtual Extensible Local Area Network (VXLAN) feature.