



VxLAN Support

This module contains information about VxLAN (Virtual eXtensible Local Area Network) Layer 2 gateway feature support on the Cisco ASR 1000 Series Routers. VxLAN is a technology that provides a Layer 2 overlay network, allowing for network isolation. The standard 802.1q VLAN implementation limits the number of tags to 4096. However, cloud service providers may want to operate more than 4096 virtual networks. VxLAN uses a 24-bit network ID, which allows for a much larger number of individual networks to be operated.

- [Finding Feature Information, on page 1](#)
- [Prerequisites for VxLAN Support, on page 1](#)
- [Information About VxLAN Support, on page 2](#)
- [Limitations of VxLAN Support, on page 3](#)
- [New Scale Number after Enhancements, on page 3](#)
- [Configuring VxLAN Layer 2 Gateway with Multicast, on page 3](#)
- [Configuring VxLAN Layer 2 Gateway with Unicast, on page 8](#)
- [Feature Information for VxLAN Support, on page 8](#)
- [Technical Assistance, on page 9](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table at the end of this module.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for VxLAN Support

The following are the prerequisites to configuring the Cisco ASR 1000 Routers as a VxLAN Layer 2 gateway:

1. Configure the loopback interface.
2. Configure the IP unicast reachability to remote VTEP's.
3. Configure Bidirectional Protocol Independent Multicast (PIM) or Protocol Independent Multicast-Sparse Mode (PIM-SM).

For more information, see the [IP Multicast: PIM Configuration Guide, Cisco IOS XE Release 3S](#).

Information About VxLAN Support

This feature enables the Cisco ASR 1000 Series Routers to act as a Layer 2 VxLAN gateway to provide support to bridge traffic across VxLAN segments in a hypervisor and on VLANs on physical servers. The operation of a VxLAN Layer 2 gateway is based on the data plane MAC address learning and flooding of multidestination traffic (such as unknown unicast, multicast, or broadcast frames) using IP multicast.

Acting as a VxLAN Layer 2 gateway, the Cisco ASR 1000 Routers can send and receive packets on multiple VxLAN networks, and provide connectivity between the hosts in a VLAN network and the virtual machines operating on a VxLAN network.

A VxLAN supports different modes for flood traffic:

- **Multicast Mode**—A VxLAN uses an IP multicast network to send broadcast, multicast, and unknown unicast flood frames. Each multicast mode VxLAN has an assigned multicast group IP address. When a new VM joins a host in a multicast mode VxLAN, the Virtual Tunnel Endpoint (VTEP) joins the assigned multicast group IP address by sending IGMP join messages. Flood traffic, broadcast, multicast and unknown unicast from the VM is encapsulated and is sent using the assigned multicast group IP address as the destination IP address. Packets sent to known unicast MAC addresses are encapsulated and sent directly to the destination server Virtual Tunnel Endpoint (VTEP) IP addresses.
- **Unicast-Only Mode**—A VxLAN uses each VEM's single unicast IP address as the destination IP address to send broadcast, multicast, and unknown unicast flood frames of the designated VTEP on each VEM that has at least one VM in the corresponding VxLAN. When a new VM joins the host in a unicast-mode VxLAN, a designated VTEP is selected for receiving flood traffic on that host. This designated VTEP is communicated to all other hosts through the Virtual Supervisor Module (VSM). Flood traffic (broadcast, multicast, and unknown unicast) is replicated on each VEM's designated VTEP in that VxLAN by encapsulating it with a VxLAN header. Packets are sent only to VEMs with a VM in that VxLAN. Packets that have a unicast MAC address are encapsulated and sent directly to the destination server's VTEP IP address.
- **MAC Distribution Mode** (supported only in unicast mode)—In this mode, unknown unicast flooding in the network is eliminated. The VSM learns all the MAC addresses from the VEMs in all the VxLANs and distributes those MAC addresses with VTEP IP mappings to other VEMs. Therefore, no unknown unicast MAC address exists in the network when the VMs on the VEMs are communicating and controlled by the same VSM.

The VxLAN Layer 2 gateway performs the following functions:

- Provides support to bridge traffic between a host in a VLAN domain and VMs behind a virtual switch (vSwitch) in a VxLAN domain. The VLAN and the virtual network identifier (VNI) on the VxLAN should be configured as member ports in the same bridge domain.
- Implements the Virtual Tunnel Endpoint (VTEP) function, which encapsulates the Layer 2 packet on the IP/UDP tunnel with the VxLAN header (VNI) information before sending it to a multicast group or particular virtual switch on the VxLAN domain.
- The VTEP function removes the VxLAN header, identifies the bridge domain under which the VNI is configured and then bridges the inner L2 packet to the VLAN side. The bridge function also learns the remote MAC address (the VM's MAC address behind the virtual switch).
- The Layer 2 gateway carries the inner payload of non-IP (Layer 2 traffic), IPv4, and IPv6 traffic over the VxLAN VNI member.

Limitations of VxLAN Support

1. Platforms that support a new scale number (8192 or 16000) require an 8G RP memory. Scale number for RP memory that is less than 8G is unchanged.
2. Scale number on platform RP+ESP5 and ASR1002F is unchanged.
3. VxLAN is not supported on ISR4000 series platforms before Cisco IOS XE Everest 16.5.1.
4. The maximum NVE interface number is unchanged on all platforms.
5. The NVE source is supported for lookback interface before Cisco IOS XE Denali 16.3. After Cisco IOS XE Denali 16.3, it can support physical interfaces as well.
6. The scale enhancement is applicable only for the VxLAN layer 2 and layer 3 gateway feature. Other bridge-domain related features are not impacted.
7. RP switchover for VxLAN is not supported on these platforms before Cisco IOS XE Denali 16.3.
8. Only one VNI ID on every bridge-domain is supported.

New Scale Number after Enhancements

The following table lists new VxLAN scale numbers on different platforms after enhancements. All platforms that support a new scale number (8192 or 16000) require an 8G RP memory.

Platform	MAX BD per system	MAX BDI interface per system	MAX VNI per system
RP+ESP200	16000	16000	16000
RP+ESP100	16000	16000	16000
RP+ESP40	16000	16000	16000
RP+ESP20	16000	16000	16000
RP+ESP10	16000	16000	16000
ASR1002-X	16000	16000	16000
ASR1001-X	16000	16000	16000
ASR 1001	8192	8192	8192
CSR1000v	8192	8192	8192

Configuring VxLAN Layer 2 Gateway with Multicast

- [Configuring the VxLAN UDP Destination Port \(Optional\), on page 4](#)
- [Creating the Network Virtualization Endpoint \(NVE\) Interface, on page 4](#)
- [Creating the Access Ethernet Flow Point \(EFP\), on page 5](#)
- [Mapping the VLAN to the Bridge Domain, on page 6](#)

Configuring the VxLAN UDP Destination Port (Optional)

The default VxLAN UDP destination is 4789. If you want to change the VxLAN UDP destination port value, you must change it before configuring the network virtualization endpoint (NVE) interface.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **vxlan udp port** *number*

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>router> enable</pre>	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: <pre>router# configure terminal</pre>	Enters global configuration mode.
Step 3	vxlan udp port <i>number</i> Example: <pre>Router(config)# vxlan udp port 1000</pre>	Configures the VxLAN UDP destination port number. The default value is 4789.

Creating the Network Virtualization Endpoint (NVE) Interface

You create the network virtualization endpoint (NVE) interface and then assign member virtual network identifiers (VNIs) to it. The mapping between the VNI range and the multicast group range is either one-to-one or many-to-one.

SUMMARY STEPS

1. **interface nve** *number*
2. **source-interface loopback** *number*
3. **member vni** {*range* | *startnumber-endnumber*} **multicast-group** *startip-address endip-address*
4. **member vni** *range*
5. **ingress-replication** *Unicast IP Addresses*
6. **no shutdown**

DETAILED STEPS

	Command or Action	Purpose
Step 1	interface nve <i>number</i> Example: <pre>Router(config)# interface nve 1</pre>	Creates a network virtualization endpoint (NVE) interface and enters NVE interface configuration mode.
Step 2	source-interface loopback <i>number</i> Example: <pre>Router(config-if)# source-interface loopback 0</pre>	Assigns the previously-created loopback interface to the NVE interface.
Step 3	member vni { <i>range</i> <i>startnumber-endnumber</i> } multicast-group <i>startip-address endip-address</i> Example: <pre>Router(config-if)# member vni 7115 multicast-group 225.1.1.1</pre>	Creates a VNI member or a range of VNI members. Repeat this step for each VNI to be added to the NVE interface. The valid values for the VNI number are from 4096 to 16777215.
Step 4	member vni <i>range</i> Example: <pre>Router(config-if)# member vni 7115</pre>	Creates a VNI member or a range of VNI members. Repeat this step for each VNI to be added to the NVE interface. The valid values for the VNI number are from 4096 to 16777215.
Step 5	ingress-replication <i>Unicast IP Addresses</i> Example: <pre>Router(config-if-nve-vni)# ingress-replication 225.1.1.1 ingress-replication 225.1.1.2</pre>	Sets up ingress-replication unicast addresses which enables the headend replication functionality.
Step 6	no shutdown Example: <pre>Router(config-if)# no shutdown</pre>	Enables the NVE interface.

Creating the Access Ethernet Flow Point (EFP)

After the member VNI is created, you must create the access Ethernet Flow Point (EFP) for the VLAN interface.

SUMMARY STEPS

1. **interface** **GigabitEthernet** *number*
2. **service instance** *id* **ethernet**
3. **encapsulation dot1q** *vlan-ID*

4. rewrite ingress tag pop 1 symmetric

DETAILED STEPS

	Command or Action	Purpose
Step 1	interface <i>GigabitEthernet number</i> Example: <code>Router(config)# interface GigabitEthernet1</code>	Enters interface configuration mode.
Step 2	service instance <i>id</i> <i>ethernet</i> Example: <code>Router(config-if)# service instance 20 ethernet</code>	Configures an Ethernet service instance on the overlay interface being configured and enters service instance configuration mode. The service instance identifier range is from 1 to 8000.
Step 3	encapsulation dot1q <i>vlan-ID</i> Example: <code>Router(config-if-srv)# encapsulation dot1q 100</code>	Defines the VLAN encapsulation format as IEEE 802.1Q and specifies the VLAN identifier.
Step 4	rewrite ingress tag pop 1 symmetric Example: <code>Router(config-if-srv)# rewrite ingress tag pop 1 symmetric</code>	Removes the VLAN tag in the Layer 2 traffic before switching to the outgoing VxLAN interface. Note This command is required to remove the VLAN tag before sending the VLAN traffic to VxLAN and adding the VLAN tag in the reverse direction.

Mapping the VLAN to the Bridge Domain

You must map the VLAN created in the previous procedure to the bridge domain.

SUMMARY STEPS

1. **bridge-domain** *bridge-id*
2. **member interface** *service-instance id*
3. **member vni** *vni-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	bridge-domain <i>bridge-id</i> Example: <code>Router(config)# bridge-domain 10</code>	Creates a bridge domain and enters bridge domain configuration mode. The valid range for bridge-id is 1-4096.
Step 2	member interface <i>service-instance id</i> Example:	Binds the bridge domain to the service instance.

	Command or Action	Purpose
	Router(config-bdomain)# member gigabitEthernet 1 service-instance 1	
Step 3	member vni vni-id Example: Router(config-bdomain)# member vni 1010	Maps the VNI to the bridge domain.

What to do next

The following example displays the NVE VNIs configured on the router:

```
Router# show nve vni
```

```
Interface  VNI          mcast          VNI state
nve1       5000           230.1.1.1      UP          L2DP 2 N/A
```

The following example displays the NVE VNIs assigned to NVE interface 1:

```
Router(config)# show nve vni interface nve1
Interface  VNI          mcast          VNI state
nve1       5000           230.1.1.1      UP          L2DP 2 N/A
```

The following example shows the status of NVE interface 1:

```
Router(config)# show nve interface nve1
Interface: nve1, State: Admin Up, Oper Up Encapsulation: Vxlan
source-interface: Loopback0 (primary:11.11.11.11 vrf:0)
```

The following example shows a detailed display for NVE interface 1:

```
Router(config)# show nve interface nve1 detail
Interface: nve1, State: Admin Up, Oper Up Encapsulation: Vxlan
source-interface: Loopback0 (primary:11.11.11.11 vrf:0)
Pkts In   Bytes In   Pkts Out   Bytes Out
0          0          0          0
```

The following example shows the NVE peers configured on the router:

```
Router(config)# show nve peers
Interface Peer-IP          VNI      Up Time
nve1      230.1.1.1             5000     UP          L2DP 2 N/A
nve2      1.1.1.3               2030     20h
```

The following example shows the bridge domain configuration with the entry in bold displaying the VM's MAC address that was learned on the VxLAN VNI:

```
Router# show bridge-domain 1000
Bridge-domain 1000 (3 ports in all)
State: UP                               Mac learning: Enabled
Aging-Timer: 300 second(s)
  GigabitEthernet1 service instance 1000
  GigabitEthernet3 service instance 1000
  vni 7639335
  MAC address      Policy Tag      Age      Pseudoport
```

```

0050.56A4.ECD2 forward dynamic 297 nve1.VNI7639335 VxLAN
src:10.0.0.1 dst:10.0.0.2
0050.56A4.257A forward dynamic 297 GigabitEthernet3.EFP1000

```

Configuring VxLAN Layer 2 Gateway with Unicast

The following example shows VxLAN with unicast ingress-replication which is a point-to-point (unicast) configuration.

```

interface Loopback0
ip address 11.11.11.11 255.255.255.255
!
interface nve1
no ip address
member vni 5001
  ingress-replication 22.22.22.22 < Remote L2 GW loopback ip>
!
source-interface Loopback0
!
bridge-domain 1
member vni 5001
member GigabitEthernet0/2/0 service-instance 1
interface GigabitEthernet0/2/0
service instance 1 ethernet
encapsulation dot1q 100
rewrite ingress tag pop 1 symmetric

```

Feature Information for VxLAN Support

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for VxLAN Support

Feature Name	Releases	Feature Configuration Information
VxLAN Support	Cisco IOS XE Release 3.13.1S Cisco IOS XE Fuji 16.9	This feature was introduced on the Cisco ASR 1000 Series Routers. This feature was introduced on the following: • Cisco ISR 1000 Series Integrated Services Routers. • Cisco ISR 4000 Series Integrated Services Routers.
Protocol Independent Multicast-Sparse Mode (PIM-SM) Support	Cisco IOS XE Release 3.17S	This feature was introduced on the Cisco ASR 1000 Series Routers. No commands were introduced or modified for this feature.

Feature Name	Releases	Feature Configuration Information
Support for multiple ingress replication peers	Cisco IOS XE Everest 16.5.1b	The VXLAN feature was modified to support multiple ingress replication peers on the Cisco ASR 1000 Series Routers. The ingress-replication command was modified to support multiple replication peers for every VNI up to 32 nodes.

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

