



Configuring Basic OTV Features

This chapter describes how to configure basic Overlay Transport Virtualization (OTV) features on Cisco NX-OS devices.

- [Finding Feature Information, on page 1](#)
- [Information About Basic OTV Features, on page 1](#)
- [Prerequisites for OTV, on page 3](#)
- [Guidelines and Limitations for OTV, on page 4](#)
- [Default Settings for OTV, on page 6](#)
- [Configuring Basic OTV Features, on page 7](#)
- [Verifying the OTV Configuration, on page 16](#)
- [Configuration Examples for OTV, on page 17](#)
- [Additional References, on page 18](#)
- [Feature History for OTV, on page 18](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your 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 "New and Changed Information" chapter or the Feature History table in this chapter.

Information About Basic OTV Features

The OTV control-plane creates adjacencies between remote sites to provide Layer 2 connectivity over a transport network. An OTV network performs the following functions:

- Discovers remote sites and builds a control-protocol adjacency
- Shares MAC routing information across the overlay network

The overlay network consists of one or more logical overlay interfaces that are created on the edge device in each remote site that connects to the physical transport network. You associate the logical overlay interface with a physical interface that connects to the transport network. The OTV control plane is responsible for discovering the edge devices in remote sites, creating control-protocol adjacencies to these sites, and establishing

protocol adjacencies among the sites. The OTV control-plane protocol uses the Intermediate-System-to-Intermediate-System (IS-IS) protocol to establish the adjacencies and exchange MAC reachability across the overlay network.



Note You do not need to configure IS-IS to use OTV. IS-IS runs in the background once OTV is enabled.

The OTV control-plane protocol also sends and receives MAC routing updates between remote sites and updates the OTV routing information base (ORIB) with these MAC to IP address pairs.

Overlay Interfaces

The overlay interface is a logical interface that connects to the remote edge devices on the overlay network through an associated physical interface on the transport network. From the perspective of MAC-based forwarding on the site, the overlay interface is simply another bridged interface. As a bridged interface, the overlay interface has unicast MAC addresses that are associated with it and is eligible for inclusion in the Outbound Interface List (OIL) for different multicast groups. However, no STP packets are forwarded over the overlay interface. Unknown unicast packets are also not flooded on the overlay interface. From the perspective of the IP transport, the overlay interface is not visible.

OTV encapsulates Layer 2 frames in IP packets and transmits them on the overlay interface.



Note The overlay interface does not come up until you configure a multicast group address or if the site-VLAN does not have at least an active port on the device.

Related Topics

[OTV Packet Flow](#)

[Creating an Overlay Interface](#), on page 8

MAC Address Learning

OTV learns MAC to IP address pairs from MAC address learning on the internal interfaces, the OTV control plane (IS-IS) updates over the overlay network, and through multicast IGMP snooping.

OTV edge devices snoop IGMP traffic and issue a Group Membership-Link State Packet (GM-LSP) to advertise the presence of receivers to remote edge devices. The remote edge devices include the overlay interface in the outbound interface list (OIL) for the corresponding multicast group. OTV does not program multicast MAC addresses in the forwarding tables, but rather updates OIL state as necessary.

- Layer 2 learning on the internal network
- IGMP snooping (for multicast MAC addresses)

All learned MAC addresses are stored in the OTV Routing Information Base (ORIB) with the VLAN ID and associated remote IP address.

Multicast Group Addresses and IGMP Snooping

OTV uses a multicast group address that is assigned from the transport network to create a unique multicast group between remote sites on the overlay network. Each edge device in the overlay network acts as a multicast host and sends an IGMP report message to join the multicast group. OTV sends encapsulated OTV control plane hello messages and MAC routing updates across this multicast group.

OTV uses IGMP snooping and group membership advertisements (GM-LSPs) to learn all multicast group members from remote sites. OTV also uses IGMP snooping to detect all multicast groups in the local site.

Related Topics

[Configuring the Multicast Group Address](#), on page 9

High Availability and ISSU

OTV supports stateful restarts and stateful switchovers. A stateful restart occurs when the OTV process fails and is restarted. A stateful switchover occurs when the active supervisor switches to the standby supervisor. The software applies the run-time configuration after the switchover.

Any upgrade from an image that is earlier than Cisco NX-OS 5.2(1) to an image that is Cisco NX-OS 5.2(1) or later in an OTV network is disruptive. A software image upgrade from Cisco NX-OS 5.2(1) or later to Cisco NX-OS 6.0 or 6.1 trains is not disruptive.

Any upgrade from an image that is earlier than Cisco NX-OS Release 6.2(2) to an image that is Cisco NX-OS Release 6.2(2) or later in an OTV network is disruptive. When you upgrade from any previous release, the OTV overlay needs to be shut down for ISSU to operate.

You must upgrade all edge devices in the site and configure the site identifier on all edge devices in the site before traffic is restored. You can prepare OTV for ISSU in a dual-homed site to minimize this disruption. An edge device with an older Cisco NX-OS release in the same site can cause traffic loops. You should upgrade all edge devices in the site during the same upgrade window. You do not need to upgrade edge devices in other sites because OTV interoperates between sites with different Cisco NX-OS versions.

Virtualization Support

The software supports multiple instances of OTV that run on the same system. OTV supports virtual routing and forwarding instances (VRFs) on the physical interface that is associated with the overlay interface. VRFs exist within virtual device contexts (VDCs). By default, the software places you in the default VDC and default VRF unless you specifically configure another VDC and VRF.

Only Layer 3 physical interfaces (and subinterfaces) or Layer 3 port channel interfaces (and subinterfaces) can be configured as join interfaces in Cisco NX-OS Release 5.0(3).

Prerequisites for OTV

OTV has the following prerequisites:

- Globally enable the OTV feature.
- Enable IGMPv3 on the join interfaces.
- Ensure connectivity for the VLANs to be extended to the OTV edge device.

Related Topics

- [Enabling the OTV Feature](#), on page 7
- [Extended VLANs and VLAN Interfaces](#)

Guidelines and Limitations for OTV

OTV has the following configuration guidelines and limitations:

- When the OTV VDC and the MPLS VDC share the same instance of the M2 forwarding engine (FE), there is a chance for traffic blackholing. The blackholing is because of the MPLS label in MPLS VDC overlap with the MPLS label, which is used to encode the OTV extended VLAN ID (OTV MPLS label = VLAN ID + 32) in the OTV VDC.

This traffic blackholing problem can be avoided by the following methods:

- You need to allocate the interfaces on the same M2 FE in such a way that the interfaces are not shared between multiple VDCs that utilize the MPLS.
 - For N7K-M224XP-23L (24-port 10GE): ports 1 to 12 are served by FE 0, and ports 13 to 24 are served by FE 1.
 - For N7K-M206FQ-23L (6-port 10/40GE): ports 1 to 3 are served by FE 0, and ports 4 to 6 are served by FE 1.
- Configure the **mpls label range**<lowest> <highest> command in the MPLS VDC to exclude all labels that can be used for OTV VLAN transport (top of the range is 4094 + 32 = 4196) from the dynamic allocation. For example: **mpls label range**4127 1028093



Note You need to reload the MPLS VDC to reallocate the existing labels within this range.

- If the same device serves as the default gateway in a VLAN interface and the OTV edge device for the VLANs being extended, configure OTV on a device (VDC or switch) that is separate from the VLAN interfaces (SVIs).
- Cisco Nexus 7000 Series Switches does not support OTV with VXLAN (VNI) feature on F3 line cards.
- The site VLAN must not be extended into the OTV. This configuration is not supported and this helps to avoid unexpected results.
- When possible, we recommend that you use a separate nondefault VDC for OTV to allow for better manageability and maintenance.
- An overlay interface will only be in an up state if the overlay interface configuration is complete and enabled (**no shutdown**). The join interface has to be in an up state.
- Configure the join interface and all Layer 3 interfaces that face the IP core between the OTV edge devices with the highest maximum transmission unit (MTU) size supported by the IP core. OTV sets the Don't Fragment (DF) bit in the IP header for all OTV control and data packets so the core cannot fragment these packets.
- Only one join interface can be specified per overlay. You can decide to use one of the following methods:

- Configure a single join interface, which is shared across multiple overlays.
- Configure a different join interface for each overlay, which increases the OTV reliability.

For a higher resiliency, you can use a port channel, but it is not mandatory. There are no requirements for 1 Gigabit Ethernet versus 10 Gigabit Ethernet or dedicated versus shared mode.

- If your network includes a Cisco Nexus 1000V switch, ensure that switch is running 4.0(4)SV1(3) or later releases. Otherwise, disable Address Resolution Protocol (ARP) and Neighbor Discovery (ND) suppression for OTV.
- The transport network must support PIM sparse mode (ASM) or PIM-Bidir multicast traffic.
- OTV is compatible with a transport network configured only for IPv4. IPv6 is not supported.
- Do not enable PIM on the join interface.
- ERSPAN ACLs are not supported for use with OTV.
- Ensure the site identifier is configured and is the same for all edge devices on a site. OTV brings down all overlays when a mismatched site identifier is detected from a neighbor edge device and generates a system message.
- Any upgrade from an image that is earlier than Cisco NX-OS Release 5.2(1) to an image that is Cisco NX-OS Release 5.2(1) or later in an OTV network is disruptive. A software image upgrade from Cisco NX-OS Release 5.2(1) or later to Cisco NX-OS Release 6.0(1) is not disruptive.
- Any upgrade from an image that is earlier than Cisco NX-OS Release 6.2(2) to an image that is Cisco NX-OS Release 6.2(2) or later in an OTV network is disruptive. When you upgrade from any previous release, the OTV overlay needs to be shut down for ISSU to operate.
- You must upgrade all edge devices in the site and configure the site identifier on all edge devices in the site before traffic is restored. An edge device with an older Cisco NX-OS release in the same site can cause traffic loops. You should upgrade all edge devices in the site during the same upgrade window. You do not need to upgrade edge devices in other sites because OTV interoperates between sites with different Cisco NX-OS versions.
- Beginning with Cisco NX-OS Release 6.2, OTV supports the coexistence of F1 or F2e Series modules with M1 or M2 Series modules in the same VDC.
- For OTV fast convergence, remote unicast MAC addresses are installed in the OTV Routing Information Base (ORIB), even on non-AED VLANs.
- For OTV fast convergence, even non-AED OTV devices create a delivery source, delivery group (DS,DG) mapping for local multicast sources and send a join request to remote sources if local receivers are available. As a result, there are two remote data groups instead of one for a particular VLAN, source, group (V,S,G) entry.
- One primary IP address and no more than three secondary IP addresses are supported for OTV tunnel depolarization.
- F3 Series modules do not support the VLAN translation and traffic depolarization features in Cisco NX-OS Release 6.2(6).
- F3 Series modules support the OTV traffic depolarization feature in Cisco NX-OS Release 6.2(8).
- F2 Series modules in a specific VDC do not support OTV. F2e modules work only as internal interfaces in an OTV VDC.

- F3 Series modules in an OTV VDC should not have the VLAN mode configured as Fabricpath.
- F3 Series modules do not support data-group configurations for subnets larger than /27, in Cisco NX-OS Releases 6.2(14) / 7.2(x) and earlier. Starting from Release 6.2(16) / 7.3(0), the largest subnet mask supported is /24.
- NXOS does not support using FEX ports for OTV site or core facing interfaces.
- Beginning with Cisco NX-OS Release 7.3(0)DX(1), M3 Series modules are supported.
- The OTV VLAN mapping feature is not supported on the Cisco M3 Series and F3 Series modules, as explained in this chapter (using the **otv vlan mapping** command). In order to have VLAN translation on OTV devices using F3 or M3 line cards, you should use per-port VLAN translation on the OTV edge device internal interface (L2 trunk port), as described in the [Configuring OTV VLAN Mapping using VLAN Translation on a Trunk Port](#) document.

Related Topics

[Creating an Overlay Interface](#), on page 8

[Configuring the Multicast Group Address](#), on page 9

[Assigning a Physical Interface to the Overlay Interface](#), on page 10

[Extended VLANs and VLAN Interfaces](#)

Default Settings for OTV

This table lists the default settings for OTV parameters.

Table 1: Default OTV Parameter Settings

Parameters	Default
OTV feature	Disabled
Advertised VLANs	None
ARP and ND suppression	Enabled
Graceful restart	Enabled
Site VLAN	1
Site identifier	0x0
IS-IS overlay hello interval	20 seconds (Cisco NX-OS Release 6.2 or later) 4 seconds (Cisco NX-OS Release 5.2 through Cisco NX-OS Release 6.1) 10 seconds (Cisco NX-OS releases prior to 5.2)
IS-IS overlay hello multiplier	3

Parameters	Default
IS-IS site hello interval	3 seconds (Cisco NX-OS Release 6.2 or later) 1 second (Cisco NX-OS releases prior to 6.2)
IS-IS site hello multiplier	20 (Cisco NX-OS Release 6.2 or later) 10 (Cisco NX-OS releases prior to 6.2)
IS-IS CSNP interval	10 seconds
IS-IS LSP interval	33 milliseconds
Overlay route tracking	Disabled
Site BFD	Disabled
Tunnel depolarization with IP pools	Enabled

Configuring Basic OTV Features

This section describes how to configure basic OTV features.



Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Enabling the OTV Feature

By default, the OTV feature is disabled on the device. You must explicitly enable the OTV feature to access the configuration and verification commands.

SUMMARY STEPS

1. **configure terminal**
2. **feature otv**
3. (Optional) **show feature | include otv [interface]**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	feature otv Example: switch(config)# feature otv	Enables OTV.
Step 3	(Optional) show feature include otv [interface] Example: switch(config)# show feature include otv	Displays the enable/disable status for the OTV feature.
Step 4	(Optional) copy running-config startup-config Example: switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Creating an Overlay Interface

You can create a logical OTV overlay interface. Once you create the overlay interface, you must configure a multicast group address and associate the interface with a physical interface.

Before you begin

- Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **interface overlay *interface***
3. (Optional) **description [*dstring*]**
4. (Optional) **show otv overlay [*interface*]**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.

	Command or Action	Purpose
Step 2	interface overlay <i>interface</i> Example: switch(config)# interface overlay 1 switch(config-if-overlay) #	Creates an OTV overlay interface and enters interface configuration mode. The range is from 0 to 65535.
Step 3	(Optional) description [<i>dstring</i>] Example: switch(config-if-overlay) # description site 4	Configures a description for the overlay network. The <i>dstring</i> is any case-sensitive, alphanumeric string up to 80 characters.
Step 4	(Optional) show otv overlay [<i>interface</i>] Example: switch(config-if-overlay) # show otv overlay 1	Displays the OTV overlay interface configuration. The range is from 0 to 65535.
Step 5	(Optional) copy running-config startup-config Example: switch(config-if-overlay) # copy running-config startup-config	Copies the running configuration to the startup configuration.

Related Topics

[Enabling the OTV Feature](#), on page 7

[Configuring the Multicast Group Address](#), on page 9

[Assigning a Physical Interface to the Overlay Interface](#), on page 10

Configuring the Multicast Group Address

You can configure a unique multicast group address for each overlay network.

OTV uses the following multicast groups in the Transport Network:

- An any source multicast (ASM) group for neighbor discovery and to exchange MAC reachability.
- A specific source multicast (SSM) group range to map internal multicast groups in the sites to the multicast groups in the core, which will be leveraged to extend the multicast data traffic across the overlay.

Before you begin

- Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **interface overlay** *interface*
3. **otv control-group** *mcast-address*
4. **otv data-group** *mcast-range1* [*mcast-range2...*]
5. (Optional) **show otv data-group** [**local** | **remote**] [**detail**]
6. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface overlay interface Example: <pre>switch(config)# interface overlay 1 switch(config-if-overlay)#</pre>	Creates an OTV overlay interface and enters interface configuration mode.
Step 3	Required: otv control-group mcast-address Example: <pre>switch(config-if-overlay)# otv control-group 239.1.1.1</pre>	Configures the multicast group address used by the OTV control plane for this OTV overlay network. The multicast group address is an IPv4 address in dotted decimal notation.
Step 4	Required: otv data-group mcast-range1 [mcast-range2...] Example: <pre>switch(config-if-overlay)# otv data-group 232.1.1.0/28</pre>	Configures one or more ranges of local IPv4 multicast group prefixes used for multicast data traffic. Use SSM multicast groups 232.0.0.0/8. The multicast group address is an IPv4 address in dotted decimal notation. A subnet mask is used to indicate ranges of addresses. You can define up to eight data-group ranges.
Step 5	(Optional) show otv data-group [local remote] [detail] Example: <pre>switch(config-if-overlay)# show otv data-group</pre>	Displays the advertised multicast groups.
Step 6	(Optional) copy running-config startup-config Example: <pre>switch(config-if-overlay)# copy running-config startup-config</pre>	Copies the running configuration to the startup configuration.

Related Topics

[Enabling the OTV Feature](#), on page 7

[Creating an Overlay Interface](#), on page 8

[Assigning a Physical Interface to the Overlay Interface](#), on page 10

Assigning a Physical Interface to the Overlay Interface

You must define a physical Layer 3 interface as the join interface for the overlay.

Before you begin

- Enable the OTV feature.
- Configure IGMPv3 on the physical Layer 3 interface that will become the join interface.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface overlay <i>interface</i> Example: <pre>switch(config)# interface overlay 1 switch(config-if-overlay)#</pre>	Creates an OTV overlay interface and enters interface configuration mode.
Step 3	Required: otv join-interface <i>interface</i> Example: <pre>switch(config-if-overlay)# otv join-interface ethernet 2/1</pre>	Joins the OTV overlay interface with a physical Layer 3 interface. You must configure an IP address on the physical interface. You can specify only one join interface per overlay. You can decide to use one of the following methods: <ul style="list-style-type: none"> • A single join interface, which is shared across multiple overlays. • A different join interface for each overlay, which increases the OTV reliability.
Step 4	(Optional) show otv overlay [<i>interface</i>] Example: <pre>switch(config-if-overlay)# show otv overlay 1</pre>	Displays the OTV overlay interface configuration.
Step 5	(Optional) copy running-config startup-config Example: <pre>switch(config-if-overlay)# copy running-config startup-config</pre>	Copies the running configuration to the startup configuration.

Related Topics

[Enabling the OTV Feature](#), on page 7

[Creating an Overlay Interface](#), on page 8

[Configuring the Multicast Group Address](#), on page 9

Assigning the Extended VLAN Range

You can configure OTV to advertise MAC address updates for a range of VLANs on an OTV overlay interface. OTV does not forward Layer 2 packets for VLANs that are not in the extended VLAN range for the overlay interface. You can add or remove VLANs from an existing extended VLAN range.



Note You can assign a VLAN to only one overlay interface. Ensure that the VLANs do not overlap across the configured overlay interfaces.

Before you begin

- Enable the OTV feature.
- Enable the VLANs in the extended VLAN range.
- Ensure that you do not extend a VLAN that is also associated with a VLAN interface (SVI) in this VDC.

SUMMARY STEPS

1. **configure terminal**
2. **interface overlay** *interface*
3. **otv extend-vlan** *vlan-range*
4. (Optional) **otv extend-vlan** {**add** | **remove**} *vlan-range*
5. (Optional) **show otv vlan** [*vlan-range*] [**detail**]
6. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface overlay <i>interface</i> Example: <pre>switch(config)# interface overlay 1 switch(config-if-overlay)#</pre>	Creates an overlay interface and enters interface configuration mode.
Step 3	Required: otv extend-vlan <i>vlan-range</i> Example: <pre>switch(config-if-overlay)# otv extend-vlan 2,5-34</pre>	Extends a range of VLANs over this overlay interface and enables OTV advertisements for these VLANs. The <i>vlan-range</i> is from 1 to 3967, and from 4048 to 4093.
Step 4	(Optional) otv extend-vlan { add remove } <i>vlan-range</i> Example: <pre>switch(config-if-overlay)# otv extend-vlan add 3</pre>	Adds or removes VLANs to the existing range of VLANs over this overlay interface. The <i>vlan-range</i> is from 1 to 3967, and from 4048 to 4093.
Step 5	(Optional) show otv vlan [<i>vlan-range</i>] [detail] Example: <pre>switch(config-if-overlay)# show otv vlan 2</pre>	Displays the VLAN information for the overlay network

	Command or Action	Purpose
Step 6	(Optional) copy running-config startup-config Example: switch(config-if-overlay)# copy running-config startup-config	Copies the running configuration to the startup configuration.

Related Topics

[Multihomed Sites and Load Balancing](#)

[Verifying Load Balancing](#)

Configuring the Site VLAN and Site Identifier

You can configure the site VLAN. OTV uses the site VLAN to communicate with other edge devices in the local site. OTV sends hello messages on the site VLAN to determine if there are other edge devices on the local site. Ensure that the site VLAN is active on at least one of the edge device ports.



Note You must configure the site identifier in Cisco NX-OS release 5.2(1) or later releases. The overlay network will not become operational until you configure the site identifier.

OTV uses the site identifier to support dual site adjacency. Dual site adjacency uses both site VLAN and site identifier to determine if there are other edge devices on the local site and if those edge devices can forward traffic. Ensure that the site identifier is the same on all neighbor edge devices in the site.



Note The site VLAN must not be extended into the OTV. This configuration is not supported and this helps to avoid unexpected results.

The site VLAN and site identifier must be configured before entering the **no shutdown** command for any interface overlay and must not be modified while any overlay is up within the site.

Before you begin

Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **otv site-vlan** *vlan-id*
3. **otv site-identifier** *id*
4. (Optional) **show otv site**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	Required: otv site-vlan <i>vlan-id</i> Example: <pre>switch(config)# otv site-vlan 10</pre>	Configures a VLAN that all local edge devices communicate on. You must configure this VLAN ID to match on all local edge devices. We recommend that you use the same VLAN ID across all sites. The range is from 1 to 3967, and from 4048 to 4093. The default is 1.
Step 3	Required: otv site-identifier <i>id</i> Example: <pre>switch(config)# otv site-identifier 0018.g957.6rk0</pre>	Configures the site identifier. You should configure this same site identifier on all local OTV edge devices. The site identifier should be unique across different sites. The range is from 0x1 to 0xffffffff. The default is 0x0. The format is either MAC address or hex format. Note This configuration step is required for Cisco NX-OS Release 5.2(1) and later releases.
Step 4	(Optional) show otv site Example: <pre>switch(config)# show otv site</pre>	Displays the OTV site information.
Step 5	(Optional) copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	Copies the running configuration to the startup configuration.

Related Topics

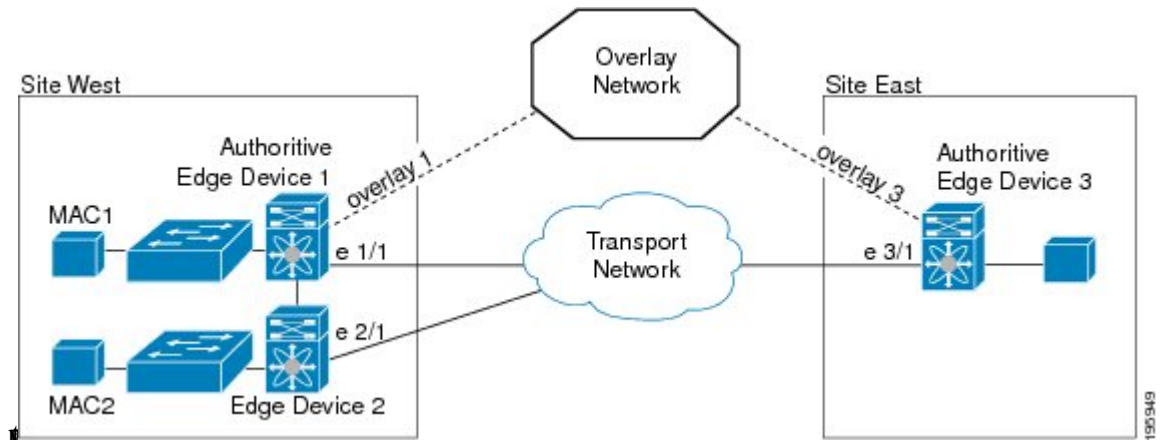
[Verifying Load Balancing](#)
[Dual Site Adjacency](#)

Preparing OTV for ISSU to Cisco NX-OS 5.2(1) or Later Releases in a Dual-Homed Site

Performing an ISSU for OTV from a release earlier than Cisco NX-OS 5.2(1) to Cisco NX-OS 5.2(1) or later releases is not supported. However, you can minimize the OTV traffic disruption in a dual-homed OTV site.

Figure 1: Dual-Homed Site

This example figure shows that Edge Device 1 is the Authoritative Edge Device (AED) for the dual-homed site in the OTV

**Before you begin**

- Step 1** Shut down all overlay interfaces on the OTV VDC for the Cisco Nexus 7000 Series chassis that you want to upgrade. ISSU does not proceed unless the overlay interfaces are administratively down.

Example:

```
edge-device-1(config)# interface overlay1
edge-device-1(config-if-overlay)# shutdown
```

Once the overlay interface on the OTV VDC is down, the other OTV edge device should become the AED for all VLANs and no major traffic disruption should occur at this point. In this example, Edge Device 2 becomes the AED.

- Step 2** Initiate the ISSU on this Cisco Nexus 7000 Series chassis. In this example, ISSU occurs on Edge Device 1. Wait until the chassis upgrade completes.

- Step 3** Configure the site identifier on this upgraded device.

Example:

```
edge-device-1(config)# otv site-identifier 0018.g957.6rk0
```

You should configure this same site identifier on all local OTV edge devices. The site identifier should be unique across different sites. The range is from 0x1 to 0xffffffff. The default is 0x0. The format is either MAC address or hex format.

Note This step is required for Cisco NX-OS Release 5.2(1) and later releases.

- Step 4** Switch back to the default VDC and enter configuration mode.

Example:

```
edge-device-1(config)# switchback
switch# configure terminal
switch(config)#
```

- Step 5** Apply the default CoPP policy.

Example:

```
switch(config)# copp profile strict
```

When you upgrade to Cisco NX-OS 5.2(1) releases, you must configure the default CoPP policy.

Step 6 Switch to the OTV VDC and enter configuration mode.

Example:

```
switch(config)# switchto vdc edge-device-1
edge-device-1# configure terminal
edge-device-1(config)#
```

Step 7 Bring the overlay interface back up on the upgraded OTV VDC.

Example:

```
edge-device-1(config)# interface overlay1
edge-device-1(config-if-overlay)# no shutdown
```

The overlay interface becomes operational.

Step 8 Shut down the overlay interface of the other OTV VDC in the dual-homed site. This action causes a disruption in OTV traffic.

Example:

```
edge-device-2(config)# interface overlay1
edge-device-2(config-if-overlay)# shutdown
```

OTV traffic is disrupted until your upgraded OTV edge device becomes the AED. In this example, Edge Device 1 becomes AED.

Step 9 Repeat [Step 1, on page 15](#) to [Step 7, on page 16](#) on the other Cisco Nexus 7000 Series chassis in the dual-homed site that runs Cisco NX-OS 5.1 or earlier releases.

You must configure the same site identifier on all edge devices in this site.

The upgraded OTV VDC comes up and becomes the AED for a subset of VLANs. OTV load balances VLANs across the two upgraded edge devices in this site.

Verifying the OTV Configuration

To display the OTV configuration, perform one of the following tasks:

Command	Purpose
<code>show running-configuration otv [all]</code>	Displays the running configuration for OTV.
<code>show otv overlay [interface]</code>	Displays information about overlay interfaces.
<code>show otv adjacency [detail]</code>	Displays information about the adjacencies on the overlay network.

Command	Purpose
show otv [<i>overlay interface</i>] [vlan [<i>vlan-range</i>] [authoritative detail]]	Displays information about VLANs that are associated with an overlay interface.
show otv isis site [database statistics]	Displays the BFD configuration state on both local and neighboring edge devices.
show otv site [all]	Displays information about the local site.
show otv [route [<i>interface</i> [neighbor-address <i>ip-address</i>]]] [vlan [<i>vlan-range</i>] [<i>mac-address</i>]]	Displays information about the OTV routes.
show otv mroute vlan <i>vlan-id</i> startup	Displays the OTV multicast route information for a specific VLAN from the OTV Routing Information Base (ORIB).
show forwarding distribution otv multicast route vlan <i>vlan-id</i>	Displays Forwarding Information Base (FIB) OTV multicast route information for a specific VLAN.
show otv vlan-mapping [overlay <i>interface-number</i>]	Displays VLAN translation mappings from a local site to a remote site.
show mac address-table	Displays information about MAC addresses.
show otv internal adjacency	Displays information about additional tunnels on the overlay network.

Configuration Examples for OTV

This example displays how to configure a basic OTV network that uses the configuration default values:

```
!Configure the physical interface that OTV uses to reach the
! DCI transport infrastructure
interface ethernet 2/1
 ip address 192.0.2.1/24
 ip igmp version 3
 no shutdown
```

```
!Configure the VLAN that will be extended on the overlay network
! and the site-vlan
vlan 2,5-10
```

```

! Configure OTV including the VLANs that will be extended.
feature otv
otv site-vlan 2
otv site-identifier 0018.g957.6rk0
interface Overlay1
  otv control-group 239.1.1.1
  otv data-group 232.1.1.0/28
  otv join-interface ethernet 2/1
!Extend the configured VLAN
otv extend-vlan 5-10
no shutdown

```

Additional References

This section includes additional information related to implementing OTV.

Related Documents

Related Topic	Document Title
Cisco NX-OS licensing	<i>Cisco NX-OS Licensing Guide</i>
OTV commands	<i>Cisco Nexus 7000 Series NX-OS OTV Command Reference</i>
Configuring BFD	<i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i>
BFD commands	<i>Cisco Nexus 7000 Series NX-OS Interfaces Command Reference</i>

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

Feature History for OTV

This table lists the release history for this feature.

Table 2: Feature History for OTV

Feature Name	Releases	Feature Information
OTV	7.3(0)DX(1)	Added support for M3 modules
OTV	6.2(6)	Added support for F3 Series modules.
Tunnel depolarization with IP pools	6.2(6)	Introduced this feature.
Selective unicast flooding	6.2(2)	Introduced this feature.

Feature Name	Releases	Feature Information
OTV VLAN mapping	6.2(2)	Introduced this feature.
Dedicated data broadcast forwarding	6.2(2)	Introduced this feature.
OTV fast convergence	6.2(2)	Introduced this feature.
Fast failure detection	6.2(2)	Introduced this feature.
OTV	6.2(2)	Added the track-adjacency-nexthop command to enable overlay route tracking.
OTV	6.2(2)	Added support for F1 and F2e Series modules.
OTV	6.2(2)	Added a reverse timer to the show otv vlan command output to show the time remaining for the VLANs to become active after the overlay interface is unshut.
ARP neighbor discovery timeout	6.1(1)	Introduced this feature.
OTV adjacency server	5.2(1)	Introduced this feature.
Dual site adjacency	5.2(1)	Added site identifier support for dual site adjacency.
Extended VLAN range	5.2(1)	Added support to add or remove VLANs to the extended VLAN range.
IPv6 unicast forwarding and multicast flooding	5.2(1)	Added support for IPv6 unicast forwarding and multicast flooding across the OTV overlay.
Configuration limits	5.2(1)	Enhanced the OTV scalability limits.
OTV	5.0(3)	Introduced this feature.

Related Topics

[OTV Adjacency Server](#)

[Configuring the Site VLAN and Site Identifier](#), on page 13

[Assigning the Extended VLAN Range](#), on page 11

