Configuring Advanced OTV Features

This chapter describes the advanced configuration for Overlay Transport Virtualization (OTV) on Cisco NX-OS devices.

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 Advanced OTV Features

OTV uses an overlay control-plane protocol to learn and propagate MAC routing information across the overlay network. The OTV control-plane protocol uses Intermediate-System-to-Intermediate-System (IS-IS) messages to build adjacencies to remote sites and to send MAC route updates to remote sites.

Building Adjacencies

OTV builds Layer 2 adjacencies to remote sites on the overlay network through the following modes:

- Autodiscovery based on OTV control-planel hello messages over a common multicast group.
- OTV adjacency server operational mode that manages and distributes a list of all peer edge devices in an overlay

OTV also builds adjacencies with other edge devices in the local site. OTV sends OTV control-plane hello messages on a dedicated VLAN, which is the site VLAN, to detect other edge devices in the same local site. These edge devices communicate to elect the Authoritative Edge Device (AED) for each configured overlay network.

Autodiscovery on the Overlay Network

The overlay routing protocol uses the IS-IS hello messages that are sent to the multicast group address to detect and build adjacencies to remote sites on the overlay network. You configure each site in the overlay network with the same multicast group address. When local and remote sites exchange hellos, a control protocol adjacency is established between the edge devices of both sites. The overlay routing protocol optionally authenticates the remote edge device before building an adjacency to the edge device.

OTV Adjacency Server

Each OTV node provides multicast send capability by replicating at the head-end itself. Each OTV node that sends a multicast packet on a nonmulticast-capable network will unicast replicate the packet. Each OTV node takes a multicast packet that is originated by the upper layers and makes a copy to send to each OTV neighbor that is interested in the multicast packet.

To be able to unicast replicate, each OTV node must know a list of neighbors to replicate to. Rather than configuring the list of all neighbors in each OTV node, you can dynamically identify the neighbors. The set of OTV neighbors might be different for different multicast groups, but the mechanism supports a unicast-replication-list (URL) per multicast group address.

The OTV does not use a replication server, so there are no choke points or longer path delays due to the lack of multicast capability. The multicast data packets, even though they are sent as a unicast message, travel on the same path from the source OTV edge device to each interested party for the group address the multicast is sent to. The only difference is that there are multiple copies being sent from the OTV edge device source.

You must configure which OTV edge device acts as an adjacency server. The OTV edge devices are configured with the IPv4 or IPv6 address of the adjacency server. All other adjacency addresses are discovered dynamically.

When a new site is added, you must configure only the OTV edge devices for the new site with the adjacency server addresses. No other sites in this VPN or other VPNs need additional configuration.

You can have more than one adjacency server per virtual private network (VPN). An adjacency server can serve multiple VPNs.

When an OTV edge device is configured with one or more adjacency server addresses, they are added to the unicast-replication-list (URL). An OTV edge device does not process an alternate server's type length value (TLV) until it believes the primary adjacency server has timed out. The primary and secondary adjacency servers are configured in each OTV edge device. An adjacency server can also be an OTV edge device that connects an OTV site to one or more VPNs.

OTV pushes the secondary adjacency server in the replication list based on the configuration with the primary server.

When you gracefully deconfigure an adjacency server, the client starts using the replication list from the secondary adjacency server and pushes the difference to OTV. If you also deconfigure the secondary adjacency server, the client deletes the replication list entries from OTV immediately.

If you reboot the primary adjacency server, the client starts using the replication list from the secondary adjacency server and pushes the difference to OTV. If the secondary and the primary adjacency servers crash or rebooted, the client makes the replication list entries stale with a timer of 10 minutes. The replication list entries are deleted after 10 minutes in case there is no adjacency server in the network that is advertising the same entries in the replication list.

If you deconfigure or reboot the adjacency server client, the client stops sending hellos to the adjacency server. Consequently, the adjacency server deletes the replication list entry for that client and advertises the deletion to all client nodes. All the nodes delete the adjacency to that client immediately.

If the OTV adjacency is lost with a unicast-only adjacency server client, but the adjacency server continues to advertise the unicast-only node, the other nodes continue to send hellos to that node until the adjacency server specifically deletes it from its own list.

Related Topics

Configuring OTV Adjacency Servers, on page 19

Authoritative Edge Device

The AED is responsible for all MAC address reachability updates for a VLAN. The overlay routing protocol sends out hello messages on the edge device internal interfaces and over a designated site VLAN to discover other OTV edge devices in the local site. OTV uses a VLAN hashing algorithm to select the AED from one of these local site edge devices.

OTV load balances traffic for the overlay network by sending MAC address reachability updates on different AEDs, depending on the VLAN of the reachability update.

If the local site has only one edge device, that edge device becomes the AED for the VLANs in the configured advertise VLAN range and does not send updates for VLANs that are outside of the configured extended VLAN range.

Related Topics

Configuring the Site VLAN and Site Identifier

Assigning the Extended VLAN Range

Dual Site Adjacency and AED Election

OTV uses the dual site adjacency state to determine the AED election. A change in the dual site adjacency state also triggers an immediate AED reelection.

Dual site adjacency state considers the following individual state changes for AED election:

Site adjacency and overlay adjacency down

Neighbors remove this edge device from consideration in the AED election.

Site adjacency down but overlay adjacency up

Neighbors continue to use this edge device in any AED elections.

Site adjacency up but overlay adjacency down

Neighbors continue to use this edge device in any AED elections if the neighbor site IS-IS hello messages still include the OTV group address.

Related Topics

Feature History for OTV

Configuring the Site VLAN and Site Identifier

AED Election

The AED is elected for each VLAN based on a VLAN ID-based hash computation. The VLAN hash algorithm assigns ordinal numbers from zero to maximum to each edge device in the local site, based on the system ID (based on the system MAC address, by default). The hash algorithm uses the following equation:

f (VLAN ID) = (VLAN ID) % edges

where edges indicates the number of OTV edge devices in the local site.

If f (VLAN ID) equals the ordinal number for the local edge device, the edge device is authoritative for that VLAN ID. In a site with two edge devices, the VLANs are split as odd and even VLAN IDs on each edge device.

MAC Address Reachability Updates

The OTV control plane uses IS-IS Link State Packets (LSPs) to propagate MAC address to IP address mappings to all edge devices in the overly network. These address mappings contain the MAC address, VLAN ID, and associated IP address of the edge device that the MAC address is reachable from.

The AED uses IGMP snooping to learn all multicast MAC addresses in the local site. OTV includes these MAC addresses in a special group-membership LSP (GM-LSP) that is sent to remote edge devices on the overlay network.

ARP Neighbor Discovery Cache

OTV can suppress unnecessary ARP messages from being sent over the overlay network. OTV builds a local Layer 3 to Layer 2 mapping for remote hosts. Any ARP requests from local hosts are served by this ARP Neighbor Discovery cache.

Related Topics

Disabling the ARP Neighbor Discovery Cache, on page 22

Selective Unicast Flooding for OTV

Normally, unknown unicast Layer 2 frames are not flooded between OTV sites, and MAC addresses are not learned across the overlay interface. Any unknown unicast messages that reach the OTV edge device are blocked from crossing the logical overlay, allowing OTV to prevent Layer 2 faults from spreading to remote sites.

The end points connected to the network are assumed to not be silent or unidirectional. However, some data center applications require the unknown unicast traffic to be flooded over the overlay to all the data centers, where end points may be silent. Beginning with Cisco NX-OS Release 6.2(2), you can configure selective unicast flooding to flood the specified destination MAC address to all other edge devices in the OTV overlay network with that unknown unicast traffic.

Related Topics

Configuring Selective Unicast Flooding, on page 22

Extended VLANs and VLAN Interfaces

A VLAN can either have Layer 3 connectivity through a VLAN interface (SVI) or the VLAN can be extended over OTV. If you have a design scenario that requires the VLAN to be both extended over OTV to a remote site and have Layer 3 connectivity through a VLAN interface, you must configure OTV in a separate VDC from the VDC that contains the VLAN interfaces.

Figure 1: OTV in a VDC

This figure shows one physical switch with a VDC for OTV configuration and a VDC for the Aggregation Layer (Agg VDC) configuration in a data center.

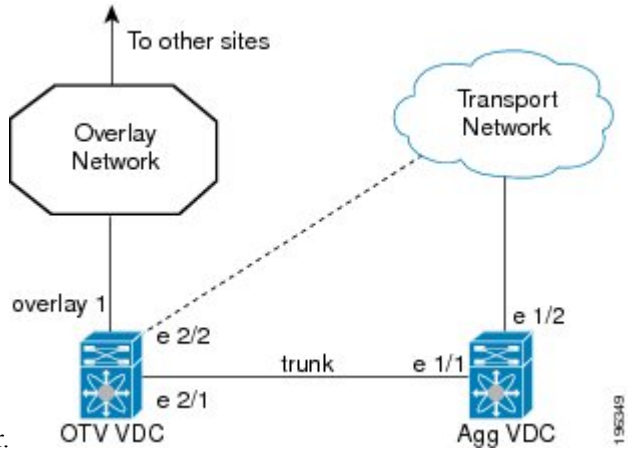

In this figure, the Agg VDC contains all the configuration and physical links for the Aggregation Layer of a data center. The Agg VDC also includes the VLAN interfaces (SVIs) for any VLANs that need Layer 3 connectivity. The Agg VDC is connected to the OTV VDC through a loopback cable over a trunk port. This trunk port carries any VLAN traffic that needs to be extended over the overlay network.

The OTV VDC also includes a trunk port that accepts this VLAN traffic. All OTV configuration exists in the OTV VDC. The overlay interface has an extended VLAN range that includes VLANs from the Agg VDC that have Layer 3 connectivity through VLAN interfaces. These extended VLANs are isolated in a separate VDC from the VLAN interfaces in the Agg VDC. The Agg VDC decides on whether a Layer 2 frame is forwarded to the local VLAN interface to Layer 3 or whether the Layer 2 frame is sent over the trunk port to the OTV VDC and encapsulated for the overlay network.

Note OTV is transparent to the Aggregation Layer and the rest of the data center site in this example.

OTV VLAN Mapping

You can extend VLANs over an OTV network in order to allow VLANs with the same VLAN ID to integrate seamlessly between local and remote sites. For example, when VLAN 1 on Site A is extended to Site B, VLAN 1 on Site A integrates seamlessly with VLAN 1 on Site B.

Beginning with Cisco NX-OS Release 6.2(2), you can map a VLAN on the local site to a VLAN with a different VLAN ID on the remote site. When you map two VLANs with different VLAN IDs across sites, they get mapped to a common VLAN called the transport VLAN. For example, when you map VLAN 1 on Site A to VLAN 2 on Site B, both VLANs are mapped to a transport VLAN. All traffic originating from

VLAN 1 on Site A is translated as going from the transport VLAN. All traffic arriving at Site B from the transport VLAN is translated to VLAN 2.

Note 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

Configuring OTV VLAN Mapping, on page 23

Forward Referencing of VLAN Maps

On the local site, you can map a VLAN that is not yet extended. OTV saves the mapping for this VLAN as a forward reference in its database. When you extend this VLAN later, the existing mapping is applied to the VLAN. The translation of traffic happens after the VLAN has been extended.

Consider a scenario where VLANs 1-10 are extended on Site A to Site B and you map VLANs 1 to 20 on Site A to VLANs on Site B. After the VLAN mapping, only VLANs 1 to 10 will be translated because they are extended. VLAN 11 to 20 mappings will be translated after you extend them to Site B. Until they are translated, the mappings are stored in the OTV database as a forward reference. The forward referencing is maintained in the OTV database even if a VLAN is unextended.

Dedicated Data Broadcast Forwarding

An OTV network that supports IP multicast uses the control-group address, which is a multicast address, to encapsulate and exchange OTV control-plane protocol updates. Each edge device that participates in a particular overlay network shares the same control-group address with all other edge devices of the same overlay network.

In addition to the control-group address, you can configure a dedicated broadcast-group address that can be used for all the broadcast traffic over the OTV cloud. If a broadcast-group address is not configured or the configuration is removed, OTV uses the configured control-group address for forwarding all broadcast packets.

Related Topics

Configuring a Dedicated Broadcast-Group Address, on page 25

OTV Fast Convergence

Cisco NX-OS Release 6.2(2) introduces the following enhancements to overcome the sources of convergence delays in an overlay network:

- VLAN AED synchronization
- Fast remote convergence by using the site ID and proactive advertisements
- Fast convergence on local edge devices by using prepopulation
- Fast detection of an edge device failure by using Bidirectional Forwarding and Detection (BFD) and route tracking
- Graceful insertion

- Graceful shutdown
- Prioritized processing of link-state packets (LSPs)

Related Topics

Configuring OTV Fast Convergence, on page 25

VLAN AED Synchronization

The election of an AED is triggered independently and is uncoordinated among the multiple edge devices in a site. Therefore, a short wait period is required to ensure that two or more edge devices are not simultaneously elected as the AED. A convergence delay can occur if there are failures at an edge device that is the AED for some VLANs.

VLAN AED synchronization in an overlay network ensures an orderly transition of the AED status from one edge device to another, prevents loops, and ensures rapid convergence.

Any edge device that needs to give AED status does so after it stops forwarding on the overlay. Any edge device that needs to take over as AED does so only after the previous AED has given up being the AED.

In AED synchronization, a backup AED is preassigned for each VLAN. The backup AED takes over immediately when an AED failure is detected.

AED Server Election

To aid in convergence improvement, the AED server and backup AED server are automatically elected per site for each overlay. All edge devices in a site elect both of these servers in a distributed manner. The eligible edge device with the highest system ID is selected as the AED server, and the edge device with the next highest system ID is selected as the AED backup server.

If an AED server is already elected and is active, a more eligible edge device is not designated as the AED server. Instead, that edge device becomes the new backup AED server. The backup AED server takes over only when the current AED server fails or declares itself ineligible.

AED Server Eligibility

An edge device indicates its eligibility to be elected as an AED server by using the AED server type, length, value (TLV). An edge device becomes eligible to be an AED server after it has completed graceful insertion, specifically after the edge device has completed synchronization and formed adjacencies with all edge devices in the site. An edge device loses its eligibility to be elected as an AED server when it loses its forwarding readiness due to events either in the site or in the overlay network.

The AED server TLV is sent in hello messages on the overlay. The absence of a control group in the site hellos indicates that the edge device should not be considered eligible to be elected as an AED server.

VLAN Reassignment

The VLANs at an OTV site are distributed among the edge devices that exist at the site. The edge device carrying the traffic of a VLAN is designated as the AED for that VLAN. During AED election, the AED server uses procedures to avoid unnecessary reassignment of VLANs among the active edge devices. The AED server ensures that the amount of message processing on various edge devices is minimal.

The following mechanisms are also used to reduce VLAN reassignments:

- When an edge device fails, the VLANs belonging to other edge devices are not reassigned; therefore, the traffic for those VLANs is not affected.

- When an edge device is added to a site, the edge device is assigned VLANs. However, VLANs are not reassigned among the other edge devices.
- VLAN reassignments to rebalance VLAN distribution after edge device insertions and failures are scheduled and spread out over a period of time.
- The AED elections for reassigning VLANs are grouped so that only one edge device gives up ownership of its VLANs at a time.

Fast Remote Convergence

Fast remote convergence is a set of techniques used to optimize delays that are introduced during the learn-advertise cycle for a newly elected AED. When an AED fails, a newly elected AED learns the local routing information of the newly acquired VLANs and advertises it to the remote site. The learn-advertise cycle is dependent on the size of the MAC table. The MAC table does not need to be updated when a remote AED fails. The convergence is independent of the size of the MAC table and the number of MACs in the affected site.

Edge devices execute the fast cutover of traffic to the new remote AED. Fast remote convergence uses the remote site's exported VLAN-AED mapping.

Fast Failure Detection

AED Failure

If an AED has a local failure, it might become unable to forward traffic for all VLANs. The AED first ensures that it has disabled traffic forwarding for all VLANs. If the AED still has overlay or site reachability, the AED indicates this failure by bringing down its AED capability on either adjacency. If the AED does not have reachability or has shut down, other edge devices detect this failure by using a dual-adjacency timeout. In both cases, the preelected backup AEDs immediately assume authority for the respective VLANs after the AED failure has been determined.

Edge Device Failure

An edge device proactively informs neighbors about local failures. If an edge device shuts down before signaling its failure, the device's failure is determined by one or both of the following:

- Dual adjacency timeout—This method is used when both overlay and site adjacencies are lost. If only overlay adjacency is lost, the edge device is still deemed to be active. The VLAN AED status that was received previously from the edge device is maintained and is not deleted. Any AED transaction involving the edge device does not proceed until the edge device becomes reachable on the overlay or completely fails. If the edge device becomes completely isolated from the overlay, the edge device indicates a forwarding readiness failure on the site adjacency.
- Site edge device consensus—This method makes the failure detection more robust at the cost of extra latency and processing. All edge devices in the same site publish a list of edge devices to which they are adjacent, either on the overlay or on the site VLAN. When an edge device loses the overlay adjacency to another edge device, the first edge device immediately triggers a hello message with this list updated to exclude that edge device. If all edge devices in the site update the list, the edge device might have failed or is no longer reachable. All edge devices generate this list, but the list might not be used to determine the failure. At first, dual adjacency is used during AED election and transitions.

BFD over an SVI is used to detect neighbor failures within a site. Both site BFD and overlay route tracking must be configured for fast device failure detection within the site.

VLAN Failure

If an AED loses forwarding readiness for a VLAN, it generates a VLAN status update to disable both forwarding readiness and AED status bits. The backup AED can assume authority as soon as it receives the status update from the AED. The AED server-driven transition mechanism handles the failures of individual VLANs. The AED server processes the VLAN status update, runs the AED election, and generates a result that includes only the new AED value in its AED message. The backup AED then takes over as AED without waiting for any edge device's response.

Graceful Insertion

AED Server Insertion

When an AED server is elected or becomes active, it waits to become updated with the VLAN status of all the edge devices in the site. The AED server does this by synchronizing the VLAN AED database with the edge devices in the site. It then schedules and runs the first AED election for all the VLANs in the VLAN AED database and starts generating VLAN AED requests. These requests might reflect the current and backup AED state of the various VLANs, or they might affect a change based on VLAN status updates.

Backup AED Server Insertion

The backup AED server runs in cold standby mode and becomes active only after the active AED server fails. Before it can run AED elections, the backup AED server must ensure that it is up to date with the AED and backup AED status of all edge devices in the site. The backup AED server does this by synchronizing the VLAN AED database with the edge devices in the site. It then runs the AED election for all VLANs and starts generating requests. During this period, the preassigned backups handle any failures of the active AEDs. However, double failures or VLAN reassignments are not handled.

Edge Device Insertion

When an edge device is inserted or reinserted in a site, it must ensure that it has received the latest version of the AED computation result from the AED server, including any pending events that the AED server might be in the process of servicing. The edge device performs an explicit synchronization with the AED server to get the latest version of the VLAN AED results. It then generates the first VLAN status update and waits for the AED server to assign it VLANs in steady state.

Graceful Shutdown

The fast convergence enhancements ensure that edge devices that shut down proactively inform neighbors by using the fast failure notification system. The grace period is used when a VDC is shut down.

QoS and OTV

By default, OTV copies the QoS DSCP or 802.1p values from the original packet to the IP header of the OTV IP packet to propagate the QoS DSCP value across the overlay network. This action ensures that the encapsulated IP packet receives the same differentiated QoS policy that the original packet received before it was extended across the overlay network.

To override this default behavior, you must apply a QoS policy to the extended VLAN. This policy can set the OTV IP encapsulation DSCP values based on a chosen match criteria. At the remote site, OTV removes this VLAN QoS policy to maintain the QoS policy for the original packet.

Note

- For 802.1Q tagged IP traffic, the outer DSCP is derived from the original COS value during encapsulation. The original COS and DSCP values are preserved during decapsulation.
- For untagged IP traffic, the outer DSCP is derived from the original DSCP value during encapsulation. The original DSCP value is preserved during decapsulation.
- For non-IP packets, the DSCP is derived from the original COS value (COS is implicit 0 for untagged traffic) during encapsulation. The original COS value is preserved during decapsulation.

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).

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.

OTV Tunnel Depolarization with IP Pools

By default, OTV uses secondary IP addresses for route depolarization. If you have two edge devices in an overlay network and each edge device is configured with two IP addresses, then four different IP header values are supported for forwarding unicast traffic between the edge devices. You must configure secondary IP addresses on the existing join interface to use route depolarization for this overlay network. Secondary IP addresses can be selected from the same subnet as the primary IP address. You do not need to configure multiple overlay networks between the same edge devices. Use the **ip address** *ip-address mask* **secondary** command to assign a secondary IP address.

On some overlay networks, secondary IP addresses on the join interface might be reserved for a different application. In this scenario, you must disable route depolarization for an entire system and to signal the lack of support for the corresponding tunnels to remote overlay members.

For route depolarization, OTV gleans its local source IP addresses from the local interface and the remote IP addresses through Intermediate-System-to-Intermediate-System (IS-IS). OTV creates multiple unicast tunnels and any one of these tunnels is used for output. Through route depolarization, you can load balance traffic to these tunnels. Route depolarization programs forwarding to point to a group of all available tunnels and modifies the forwarding infrastructure to load balance based on the actual IP packet. This feature enables load balancing based on both source and destination MAC addresses, source and destination IP addresses, or on any other criteria available to the forwarding hardware.

By default, route depolarization is enabled. Use the **otv depolarization disable** command to disable the route depolarization feature. OTV displays the secondary IP addresses that are used by the overlay interfaces and adjacencies.

Related Topics

Disabling Tunnel Depolarization with IP Pools, on page 32

OTV UDP Encapsulation

The OTV UDP header encapsulation mode is introduced in the Nexus 7000 series (7000 and 7700) devices having F3 or M3 line cards and the NX-OS 7.2.0 software version. In this version, the forwarding engine for control plane and data plane packets supports UDP encapsulation over IP over Ethernet. The control and data paths will use UDP headers for the multicast and unicast core routing. The IANA assigned UDP and TCP port number for OTV is port 8472. The header format aligns bit by bit with the header format used for the VXLAN header defined in IETF RFC 7348.

UDP encapsulation helps utilize more links in the core network as the UDP source port is varied automatically.

By default, the encapsulation format is MPLS-GRE. You can configure the OTV encapsulation format as UDP using the **otv encapsulation-format ip udp** command.

Note Only Nexus 7000 series devices having F3 or M3 line cards support OTV UDP header encapsulation mode. OTV sites across a network should have the same encapsulation format configured.

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
Extended VLANs and VLAN Interfaces, on page 5

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
Configuring the Multicast Group Address
Assigning a Physical Interface to the Overlay Interface
Extended VLANs and VLAN Interfaces, on page 5

Guidelines for OTV Multicast

OTV has the following guidelines for multicast configuration:

- OTV does not require Protocol Independent Multicast (PIM) to be configured on an edge device. If you configure PIM on the edge device, ensure that the rendezvous point (RP) is also configured on the edge device. The reverse-path forwarding (RPF) interface for (*.PG) should be join interface.
- Do not configure PIM on a join interface of the edge device.
- You should configure IGMP version 3 on both sides of the join interface link. The OTV edge devices send IGMP (S,G) joins to the edge devices in other sites in the same VPN. If you must configure IGMPv2, you must configure the last-hop router to do an ssm-translate, and the data-group range for the overlay interface must be SSM.
- You can directly connect edge devices in different sites.
- If there is no router in the site, you must configure the **ip igmp snooping querier** command in VLAN configuration mode on the switch.
- IGMP snooping for VLANs extended over the overlay network is enabled by default and should not be disabled. IGMP reports that are originated in the site are not sent across the core. Enough multicast state is built in the edge devices and core routers so that traffic can be sent from the source in the source site to a destination in the destination site.
- You do not need to configure a unicast routing protocol on join interfaces, although in most situations, one will be configured.
- You must disable optimized multicast forwarding (OMF) in IGMP snooping in OTV edge devices for IPv6 unicast or multicast traffic to flow across an OTV overlay network.
- The IGMP snooping timer needs to be set to four (using the **ip igmp snooping max-gq-miss 4** command) on all L2 switches in a site that runs OTV. If there is an AED failover and the snooping timer is set to

the default of three, snooped groups on the aggregation switches may prematurely expire. This may delay multicast convergence.

- When you assign an IP address to a loopback interface for Anycast RP configuration on an OTV (edge) device, ensure that you do not use the same IP address as the multicast source IP address for the device.

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
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

Parameters	Default
Site BFD	Disabled
Tunnel depolarization with IP pools	Enabled

Configuring Advanced OTV Features

This section describes the tasks for configuring advanced 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.

Configuration Modes

The following sections show how to enter each of the configuration modes. From a mode, you can enter the question mark (?) command to display the commands available in that mode.

Interface Configuration Mode Example

The following example shows how to enter the overlay interface configuration mode:

```
switch# configure terminal
switch(config)# interface overlay 2
switch(config-if-overlay)#
```

OTV IS-IS VPN Configuration Mode Example

The following example shows how to enter OTV IS-IS VPN configuration mode:

```
switch# configure terminal
switch(config)# otv-isis default
switch(config-router)# vpn overlay 2
switch(config-router-vrf)#
```

Configuring Authentication for Edge Devices

You can configure authentication for the OTV control-plane protocol hello messages. OTV use hello authentication to authenticate a remote site before OTV creates an adjacency to that remote site. Each overlay network uses a unique authentication key. An edge device only creates an adjacency with a remote site that shares the same authentication key and authentication method.

OTV supports the following authentication methods:

- Clear text

- Message Digest (MD5) authentication

Before you begin

- Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **interface overlay** *interface*
3. **otv isis authentication-check**
4. **otv isis authentication-type** {**cleartext** | **md5**}
5. **otv isis authentication keychain** *keychain-name*
6. (Optional) **show otv overlay** [*interface*]
7. (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	**interface overlay** *interface*
	Example:
	`switch(config)# interface overlay 1`
	`switch(config-if-overlay)#`
		Creates an OTV overlay interface and enters interface configuration mode.

Step 3	Required: **otv isis authentication-check** **Example:** `switch(config-if-overlay)# otv isis authentication-check`	Enables authentication of hello messages between OTV edge devices. The default is enabled.
Step 4	Required: **otv isis authentication-type** {**cleartext** | **md5**} **Example:** `switch(config-if-overlay)# otv isis authentication-type md5`	Configures the authentication method.
Step 5	Required: **otv isis authentication keychain** *keychain-name* **Example:** `switch(config-if-overlay)# otv isis authentication keychain OTVKeys`	Configures the authentication keychain for edge device authentication. The *keychain-name* can be any case-sensitive alphanumeric string up to 16 characters. See the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide* for information about key chains.
Step 6	(Optional) **show otv overlay** [*interface*] **Example:**	Displays the OTV overlay interface configuration.

	Command or Action	Purpose
	`switch(config-if-overlay)# show otv overlay 1`	
Step 7	(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

Configuring OTV PDU Authentication, on page 18

Configuring OTV PDU Authentication

You can configure OTV to authenticate all incoming OTV control-plane protocol data units (PDUs). OTV supports the following authentication methods:

- Clear text
- Message Digest (MD5) authentication

Note OTV control-plane protocol hello authentication is configured separately.

Before you begin

Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **otv-isis default**
3. **vpn** *overlay-name*
4. **authentication-check**
5. **authentication-type** {**cleartext** | **md5**}
6. **authentication keychain** *keychain-name*
7. (Optional) **show otv isis hostname vpn** [*overlay-name* | **all**]
8. (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	**otv-isis default**
	Example:
	`switch(config)# otv-isis default`
	`switch(config-router)#`
		Enters OTV router configuration mode.

Step 3	**vpn** *overlay-name*
	Example:
	`switch(config-router)# vpn overlay 2`
	`switch(config-router-vrf)#`
		Enters OTV virtual private network (VPN) configuration mode. The *overlay-name* should match with the overlay interface.

Step 4	Required: **authentication-check** **Example:** `switch(config-router-vrf)# authentication-check`	Enables authentication of OTV PDUs. The default is enabled.
Step 5	Required: **authentication-type** {**cleartext** | **md5**} **Example:** `switch(config-router-vrf)# authentication-type md5`	Configures the authentication method.
Step 6	Required: **authentication keychain** *keychain-name* **Example:** `switch(config-router-vrf)# authentication keychain OTVKeys`	Configures the authentication keychain for PDU authentication. The *keychain-name* can be any case-sensitive, alphanumeric string up to 16 characters. For more information about key chains, see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.
Step 7	(Optional) **show otv isis hostname vpn** [*overlay-name* | **all**] **Example:** `switch(config-router-vrf)# show otv isis hostname vpn Marketing`	Displays the OTV VPN configuration. The *overlay-name* can be any case-sensitive, alphanumeric string up to 32 characters.
Step 8	(Optional) **copy running-config startup-config** **Example:** `switch(config-router-vrf)# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Related Topics

Configuring Authentication for Edge Devices, on page 16

Configuring OTV Adjacency Servers

You can either configure the local edge device to act as an adjacency server, or you can configure a remote adjacency server.

Before you begin

Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **interface overlay** *interface*
3. (Optional) **otv adjacency-server unicast-only**
4. (Optional) **otv use-adjacency-server** *primary-ip-address* [*secondary-ip-address*] **unicast-only**
5. (Optional) **show otv adjacency** [**overlay** *if-number* | **vpn** *vpn-name*] [**detail**]
6. (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	**interface overlay** *interface*
	Example:
	`switch(config)# interface overlay 1`
	`switch(config-if-overlay)#`
		Creates an OTV overlay interface and enters interface configuration mode.

Step 3	(Optional) **otv adjacency-server unicast-only** **Example:** `switch(config-if-overlay)# otv adjacency-server unicast-only`	Configures the local edge device to act as an adjacency server. **Note** If the two overlay interface numbers do not match between the two OTV sites configured to use unicast adjacency servers, the OTV adjacencies will not form and OTV will not come up until the overlay interface numbers are changed to match.
Step 4	(Optional) **otv use-adjacency-server** *primary-ip-address* [*secondary-ip-address*] **unicast-only** **Example:** `switch(config-if-overlay)# otv use-adjacency-server 192.0.2.1 unicast-only`	Configures the local edge device to use a remote adjacency server. The IP address format is in dotted decimal notation. The *secondary-ip-address* argument is the IP address of the backup adjacency server, if you have configured a backup adjacency server.
Step 5	(Optional) **show otv adjacency** [**overlay** *if-number*	**vpn** *vpn-name*] [**detail**] **Example:** `switch(config-if-overlay)# show otv adjacency overlay 1`	Displays the OTV adjacency information. The *if-number* range is from 0 to 65503. The *vpn-name* is any case-sensitive, alphanumeric string up to 80 characters.
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.

Configuring the ARP Neighbor Discovery Timeout for an Overlay

Beginning with NX-OS Release 6.1(1), you can configure how long a dynamically learned IP address and its corresponding MAC address remain in the OTV ARP and ND cache. This command applies to all IP addresses learned for this overlay regardless of whether they were learned on the overlay interface or on an associated access interface.

SUMMARY STEPS

1. **configure terminal**
2. **interface overlay** *interface*
3. **otv arp-nd timeout** *seconds*
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	**interface overlay** *interface*
	Example:
	`switch(config)# interface overlay 1`
	`switch(config-if-overlay)#`
		Creates an overlay interface and enters interface configuration mode.

Step 3	Required: **otv arp-nd timeout** *seconds* **Example:** `switch(config-if-overlay)# otv arp-nd timeout 70`	Configures the time, in seconds, that an entry remains in the ARP-ND cache. The time is in seconds varying from 60 (1 minute) to 86400 (24 hours). The default timeout value is 480 seconds. Use the **no** form of this command to disable this feature.
Step 4	(Optional) **copy running-config startup-config** **Example:** `switch(config-if-overlay)# copy running-config startup-config`	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure the ARP Neighbor Discovery timeout for an overlay:

```
switch # configure terminal
switch(config)# interface overlay 1
switch(config-if-overlay)# otv arp-nd timeout 70
switch(config-if-overlay)# copy running-config startup-config
```

Disabling the ARP Neighbor Discovery Cache

An ARP cache is maintained by every OTV edge device and is populated by snooping ARP replies. Initial ARP requests are broadcasted to all sites, but subsequent ARP requests are suppressed at the edge device and answered locally. OTV edge devices can reply to ARPs on behalf of remote hosts. Use the following procedure to disable this functionality.

SUMMARY STEPS

1. **configure terminal**
2. **interface overlay** *interface*
3. **no otv suppress-arp-nd**
4. (Optional) **show otv arp-nd-cache** [*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.

Step 2	**interface overlay** *interface*
	Example:
	`switch(config)# interface overlay 1`
	`switch(config-if-overlay)#`
		Creates an OTV overlay interface and enters interface configuration mode.

Step 3	Required: **no otv suppress-arp-nd** **Example:** `switch(config-if-overlay)# no otv suppress-arp-nd`	Suppresses the sending of ARP and ND packets on an overlay network. This command supports both IPv4 and IPv6.
Step 4	(Optional) **show otv arp-nd-cache** [*interface*] **Example:** `switch(config-if-overlay)# show otv arp-nd-cache`	Displays the Layer 2 and Layer 3 address mapping for remote MAC addresses.
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.

Configuring Selective Unicast Flooding

You can configure selective unicast flooding for OTV.

Before you begin

Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **otv flood mac** *mac-address* **vlan** *vlan-id*
3. (Optional) **show otv mroute vlan** *vlan-id* **startup**
4. (Optional) **show otv route vlan** *vlan-id*
5. (Optional) **show forwarding distribution otv multicast route vlan** *vlan-id*
6. (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	**otv flood mac** *mac-address* **vlan** *vlan-id* **Example:** `switch(config)# otv flood mac 0000.ffaa.0000 vlan 328`	Enables selective unicast OTV flooding.
Step 3	(Optional) **show otv mroute vlan** *vlan-id* **startup** **Example:** `switch(config)# show otv mroute vlan 328 startup`	Displays the OTV multicast route information for a specific VLAN from the OTV Routing Information Base (ORIB).
Step 4	(Optional) **show otv route vlan** *vlan-id* **Example:** `switch(config)# show otv route vlan 328`	Displays OTV Intermediate System-to-Intermediate System (IS-IS) route information from ORIB for a specific VLAN.
Step 5	(Optional) **show forwarding distribution otv multicast route vlan** *vlan-id* **Example:** `switch(config)# show forwarding distribution otv multicast route vlan 328`	Displays Forwarding Information Base (FIB) OTV multicast route information for a specific VLAN.
Step 6	(Optional) **copy running-config startup-config** **Example:** `switch# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Configuring OTV VLAN Mapping

You can configure OTV VLAN mapping to allow VLANs with different VLAN IDs to communicate across sites.

Note The OTV VLAN mapping feature is not supported on the Cisco M3 Series and F3 Series modules. 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.

Before you begin

Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **interface overlay** *interface-number*
3. **otv vlan mapping** [**add** | **remove**] {*vlan-range*}
4. (Optional) **show otv vlan-mapping** [**overlay** *interface-number*]
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.

Step 2	**interface overlay** *interface-number* **Example:** `switch(config)# interface overlay 1`	Creates an OTV overlay interface and enters overlay interface configuration mode.
Step 3	**otv vlan mapping** [**add** | **remove**] {*vlan-range*} **Example:** `switch(config-if-overlay)# otv vlan mapping 1-5 to 7-11`	Creates translation mappings of VLANs on a local site to VLANs on a remote site in an OTV network.
Step 4	(Optional) **show otv vlan-mapping** [**overlay** *interface-number*] **Example:** `switch(config-if-overlay)# show otv vlan-mapping`	Displays VLAN translation mappings from a local site to a remote site.
Step 5	(Optional) **copy running-config startup-config** **Example:** `switch# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Configuring a Dedicated Broadcast-Group Address

You can configure a dedicated broadcast-group address for an OTV network.

Before you begin

Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **interface overlay** *interface-number*
3. **otv broadcast-group** *multicast-address*
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.

Step 2	**interface overlay** *interface-number* **Example:** `switch(config)# interface overlay 1`	Creates an OTV overlay interface and enters overlay interface configuration mode.
Step 3	**otv broadcast-group** *multicast-address* **Example:** `switch(config-if-overlay)# otv broadcast-group 224.1.1.10`	Configures an IP multicast address as the dedicated broadcast-group address for the specified OTV network.
Step 4	(Optional) **show otv** [**overlay** *interface*] **Example:** `switch(config-if-overlay)# show otv`	Displays the OTV overlay interface configuration.
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.

Configuring OTV Fast Convergence

You can enable OTV fast convergence by configuring a switched virtual interface (SVI) on an OTV site VLAN.

Before you begin

Enable the OTV feature.

Enable the BFD feature.

Ensure that the IP addresses of all OTV switches in a site are in the same subnet as the site VLAN SVI.

Ensure that the site VLAN is not extended on the OTV overlay.

SUMMARY STEPS

1. **configure terminal**
2. **feature interface-vlan**
3. **interface vlan**
4. **no ip redirects**
5. **ip address** *ip-address mask*
6. **no shutdown**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**feature interface-vlan** **Example:** `switch(config)# feature interface-vlan`	Enables the creation of VLAN interfaces.

Step 3	**interface vlan**
	Example:
	`switch(config)# interface vlan 2500`
	`switch(config-if)#`
		Creates an SVI and enters interface configuration mode.

Step 4	**no ip redirects** **Example:** `switch(config-if)# no ip redirects`	Disables IP redirects.
Step 5	**ip address** *ip-address mask* **Example:** `switch(config-if)# ip address 172.1.2.1 255.255.255.0`	Sets a primary or secondary IP address for the interface.
Step 6	**no shutdown** **Example:** `switch(config-if)# no shutdown`	Enables the interface.

Configuring Fast Failure Detection

You can configure fast failure detection in an OTV site VLAN.

Before you begin

Enable the OTV feature.

Enable the BFD feature.

SUMMARY STEPS

1. **configure terminal**
2. **otv-isis default**
3. **track-adjacency-nexthop**
4. **exit**
5. **otv site-vlan** *vlan-id*
6. **otv isis bfd**
7. (Optional) **show otv isis track-adjacency-nexthop**
8. (Optional) **show bfd neighbors**
9. (Optional) **show otv isis site**
10. (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	**otv-isis default**
	Example:
	`switch(config)# otv-isis default`
	`switch(config-router)#`
		Enters OTV router configuration mode.

Step 3	**track-adjacency-nexthop** **Example:** `switch(config-router)# track-adjacency-nexthop`	Enables overlay route tracking. **Note** This command tracks only the site-adjacent edge device. The site-adjacent device must be reachable only by IGP and not by any static routes or default routes.

Step 4	**exit**
	Example:
	`switch(config-router)# exit`
	`switch(config)#`
		Exits OTV router configuration mode.

Step 5	**otv site-vlan** *vlan-id* **Example:**	Configures a VLAN on which all local edge devices can communicate.

	Command or Action	Purpose
	`switch(config)# otv site-vlan 10` `switch(config-site-vlan)#`	**Note** You must configure this VLAN ID on all local edge devices.
Step 6	**otv isis bfd** **Example:** `switch(config-site-vlan)# otv isis bfd`	Enables BFD on an OTV site VLAN for failure detection and notification. The OTV IS-IS instance brings down site adjacency when a BFD failure notification occurs.
Step 7	(Optional) **show otv isis track-adjacency-nexthop** **Example:** `switch(config-site-vlan)# show otv isis track-adjacency-nexthop`	Displays the OTV IS-IS next-hop adjacencies.
Step 8	(Optional) **show bfd neighbors** **Example:** `switch(config-site-vlan)# show bfd neighbors`	Displays a line-by-line listing of existing BFD adjacencies.
Step 9	(Optional) **show otv isis site** **Example:** `switch(config-site-vlan)# show otv isis site`	Displays the BFD configuration state on both local and neighboring edge devices.
Step 10	(Optional) **copy running-config startup-config** **Example:** `switch(config-site-vlan)# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Configuring Redistribution

You can configure a route map to filter OTV updates on an overlay network. The route map can use the following match options:

match mac-list

List of MAC addresses to match against. Only MAC addresses that match a mac-list entry are redistributed across the overlay network.

match vlan

VLAN ID to match against. OTV redistributes the MAC routes that match this VLAN ID.

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide* for more information on route maps and MAC address lists.

Before you begin

- Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**

2. **otv-isis default**
3. **vpn** *overlay-name*
4. **redistribute filter route-map** *map-name*
5. (Optional) **show otv isis redistribute route** [**vpn** *overlay-name* | **summary**]
6. (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	**otv-isis default**
	Example:
	`switch(config)# otv-isis default`
	`switch(config-router)#`
		Enters OTV router configuration mode.

Step 3	**vpn** *overlay-name*
	Example:
	`switch(config-router)# vpn Marketing`
	`switch(config-router-vrf)#`
		Enters OTV virtual private network (VPN) configuration mode. The *overlay-name* can be any case-sensitive, alphanumeric string up to 32 characters.

Step 4	Required: **redistribute filter route-map** *map-name* **Example:** `switch(config-router-vrf)# redistribute filter route-map otvFilter`	Assigns a route map that OTV uses to filter OTV updates that are sent to remote sites. The *map-name* can be any case-sensitive, alphanumeric string up to 63 characters.

Step 5	(Optional) **show otv isis redistribute route** [**vpn** *overlay-name*	**summary**]
		Example:
		`switch(config-router-vrf)# show otv isis redistribute route`
		`vpn Marketing`
			Displays the OTV VPN redistribution information. The *overlay-name* can be any case-sensitive, alphanumeric string up to 32 characters.

Step 6	(Optional) **copy running-config startup-config** **Example:** `switch(config-router-vrf)# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Verifying Load Balancing

You can load balance overlay network traffic across different edge devices in a local site. OTV uses the site VLAN to discover all edge devices in the local site. OTV then dynamically assigns VLANs to an AED for each VLAN, based on the VLAN ID, the number of edge devices in the local site, and the system ID of the edge device. Load balancing is achieved because each edge device is authoritative for a subset of all VLANs that are transported over the overlay.

Before you begin

- Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **otv site-vlan** *vlan-id*
3. (Optional) **show otv site** [**all**] [**detail**]
4. (Optional) **show otv** [*overlay-interface*] **vlan** *vlan-id* **authoritative** [**detail**]

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**otv site-vlan** *vlan-id* **Example:** `switch(config)# otv site-vlan 10`	Configures a VLAN that all local edge devices communicate on. You must configure this VLAN ID on all local edge devices. The range is from 1 to 3967 and from 4048 to 4093. The default is 1.
Step 3	(Optional) **show otv site** [**all**] [**detail**] **Example:** `switch(config)# show otv site`	Displays all the edge devices for the local site.
Step 4	(Optional) **show otv** [*overlay-interface*] **vlan** *vlan-id* **authoritative** [**detail**] **Example:** `switch(config)# show otv vlan authoritative detail`	Displays all the VLANs that this edge device is the AED for. Use this command on each edge device in the local site to show which is the AED for each VLAN.

Example

This example shows the output for the **show otv vlan authoritative detail** command:

```
switch(config)# show otv vlan authoritative detail
OTV VLAN Configuration Information
Legend: F - Forwarding B - Blocked
VLAN-ID  VlanState           Switchport/    External       Overlay
                             Forward Count  Interface      Group
```

Related Topics

Multihomed Sites and Load Balancing
Authoritative Edge Device, on page 3
Configuring the Site VLAN and Site Identifier

Tuning OTV

You can tune parameters for the overlay routing protocol.

Note We recommend that only very experienced users of OTV perform these configurations.

Before you begin

- Enable the OTV feature.

SUMMARY STEPS

1. **configure terminal**
2. **interface overlay** *interface*
3. (Optional) **otv isis csnp-interval** *seconds*
4. (Optional) **otv isis hello-interval** *seconds*
5. (Optional) **otv isis hello-multiplier** *mulltiplier*
6. (Optional) **otv isis hello-padding**
7. (Optional) **otv isis lsp-interval** *msec*
8. (Optional) **otv isis metric** *metric*
9. (Optional) **otv isis priority** *dis-priority*
10. (Optional) **show otv isis** [*isis-tag*] [**interface** *interface*]
11. (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	**interface overlay** *interface*
	Example:
	`switch(config)# interface overlay 1`
	`switch(config-if-overlay)#`
		Creates an OTV overlay interface and enters interface configuration mode.

Step 3	(Optional) **otv isis csnp-interval** *seconds* **Example:** `switch(config-if-overlay)# otv isis csnp-interval 100`	Specifies the interval between CSNP PDUs on an interface. The *seconds* range is from 1 to 65535. The default is 10 seconds.

	Command or Action	Purpose
Step 4	(Optional) **otv isis hello-interval** *seconds* **Example:** `switch(config-if-overlay)# otv isis hello-interval 30`	Specifies the interval between hello PDUs on an interface. The *seconds* range is from 1 to 65535. The default is 10 seconds.
Step 5	(Optional) **otv isis hello-multiplier** *mulltiplier* **Example:** `switch(config-if-overlay)# otv isis hello-multiplier 30`	Specifies the multipler that is used to calculate the interval within which hello PDUs must be received to keep the OTV adjacency up. The *multiplier* range is from 3 to 1000. The default is 3.
Step 6	(Optional) **otv isis hello-padding** **Example:** `switch(config-if-overlay)# otv isis hello-padding`	Pads OTV hello PDUs to the full MTU length.
Step 7	(Optional) **otv isis lsp-interval** *msec* **Example:** `switch(config-if-overlay)# otv isis lsp-interval 30`	Specifies the interval between LSP PDUs on an interface during flooding. The *msec* range is from 10 to 65535. The default is 33 milliseconds.
Step 8	(Optional) **otv isis metric** *metric* **Example:** `switch(config-if-overlay)# otv isis metric 30`	Configures the OTV metric on an interface. The *metric* range is from 1 to 16777215.
Step 9	(Optional) **otv isis priority** *dis-priority* **Example:** `switch(config-if-overlay)# otv isis lsp-interval 30`	Configures the OTV priority for DIS election on the interface. The *priority* range is from 1 to 127. The default is 64.
Step 10	(Optional) **show otv isis** [*isis-tag*] [**interface** *interface*] **Example:** `switch(config-if-overlay)# show otv isis interface overlay 2`	Displays the overlay routing protocol information for the OTV overlay interface.
Step 11	(Optional) **copy running-config startup-config** **Example:** `switch(config-if-overlay)# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Disabling Tunnel Depolarization with IP Pools

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# **feature otv**	Enables OTV.
Step 3	switch(config)# **otv depolarization disable**	Disables route depolarization. By default, route depolarization is enabled on the device.
Step 4	(Optional) switch(config)# **show otv** [**adjacency**]	Displays secondary addresses and information about the adjacencies on the overlay network.
Step 5	(Optional) switch(config)# **show otv adjacency detail**	Displays information about a secondary tunnel on the overlay network.
Step 6	(Optional) switch(config)# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Verifying the OTV Configuration

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

Command	Purpose
show running-configuration otv [**all**]	Displays the running configuration for OTV.
show otv overlay [*interface*]	Displays information about overlay interfaces.
show otv adjacency [**detail**]	Displays information about the adjacencies on the overlay network.
show otv [**overlay** *interface*] [**vlan** [*vlan-range*] [**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** [*interface* [**neighbor-address** *ip-address*]] [**vlan** *vlan-range*] [*mac-address*]]	Displays information about the OTV routes.
show otv mroute vlan *vlan-id* **startup**	Displays the OTV multicast route information for a specific VLAN from the OTV Routing Information Base (ORIB).

Command	Purpose
show forwarding distribution otv multicast route vlan *vlan-id*	Displays Forwarding Information Base (FIB) OTV multicast route information for a specific VLAN.
show otv vlan-mapping [**overlay** *interface-number*]	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

Configuration Example for Load Balancing

Basic OTV Network

The following example displays how to configure load balancing on two edge devices in the same site:

```
Edge Device 1
interface ethernet 2/1
 ip address 192.0.2.1/24
 ip igmp version 3
 no shutdown

vlan 5-10

feature otv
otv site-identifier 0018.g957.6rk0
interface overlay 1
 otv control-group 239.1.1.1
 otv data-group 239.1.1.0/29
 otv join-interface ethernet 2/1
 otv extend-vlan 5-10
 no shutdown



Edge Device 2
interface ethernet 1/1
 ip address 192.0.2.16/24
 ip igmp version 3
 no shutdown

vlan 5-10
```

```
feature otv
otv site-identifier 0018.g957.6rk0
interface overlay 2
 otv control group 239.1.1.1
 otv data-group 239.1.1.0/29
 otv join-interface ethernet 1/1
 otv extend-vlan 5-10
 no shutdown
```

Configuration Example for OTV Selective Unicast Flooding

The following example shows the configuration and verification of the flooding of the 0000.ffaa.0000 destination MAC address to all other edge devices in the OTV overlay network for VLAN 328:

```
switch# configure terminal
switch(config)# otv flood mac 0000.ffaa.0000 vlan 328
switch(config)# show otv mroute vlan 328 startup
switch(config)# show otv route vlan 328
switch(config)# show forwarding distribution otv multicast route vlan 328
switch(config)# show otv mroute vlan 328 startup
OTV Multicast Routing Table For Overlay1
(328, *, 255.255.255.253), metric: 0, uptime: 00:00:46, site - New entry
Outgoing interface list: (count: 1)
Overlay1, uptime: 00:00:46, otv
switch(config)# show otv route vlan 328
OTV Unicast MAC Routing Table For Overlay2
VLAN MAC-Address Metric Uptime Owner Next-hop(s)
---- -------------- ------ -------- --------- -----------
328 0000.ffaa.0000 0 00:00:15 static Overlay2
switch(config)# show forwarding distribution otv multicast route vlan 328
Vlan: 100, Group: 255.255.255.253, Source: 0.0.0.0
OTV Outgoing Interface List Index: 6
Reference Count: 1
Number of Outgoing Interfaces: 2
External interface:
Delivery group IP: 255.255.255.253
Delivery source IP: 0.0.0.0
Interface Index: Overlay1
External interface: Ethernet3/11
Delivery group IP: 239.1.1.1
Delivery source IP: 10.10.10.10
Interface Index: Overlay1
```

Configuration Examples for OTV VLAN Mapping

The following example shows how to map VLANs 10, 14, 15, 16, and 18 on Site A with VLANs 20, 21, 25, 28, and 30 on Site B:

```
switch(config)# interface overlay 5
switch(config-if-overlay)# otv vlan mapping 10,14-16,18 to 20-21,25,28,30
switch(config-if-overlay)# show otv vlan-mapping
Original VLAN -> Translated VLAN
--------------------------------
10 -> 20
14 -> 21
15 -> 25
16 -> 28
18 -> 30
```

The following example shows how to overwrite the previous VLAN mapping translation configuration:

```
switch(config)# interface overlay 5
switch(config-if-overlay)# otv vlan mapping 40,41,42 to 50,51,52
switch(config-if-overlay)# show otv vlan-mapping
Original VLAN -> Translated VLAN
--------------------------------
40 -> 50
41 -> 51
42 -> 52
```

The following example shows how to add a VLAN map to an existing translation configuration:

```
switch(config)# interface overlay 5
switch(config-if-overlay)# otv vlan mapping add 43 to 53
switch(config-if-overlay)# show otv vlan-mapping
Original VLAN -> Translated VLAN
--------------------------------
40 -> 50
41 -> 51
42 -> 52
43 -> 53
```

The following example shows how to remove a VLAN map from an existing translation configuration:

```
switch(config)# interface overlay 5
switch(config-if-overlay)# otv vlan mapping remove 40 to 50
switch(config-if-overlay)# show otv vlan-mapping
Original VLAN -> Translated VLAN
--------------------------------
41 -> 51
42 -> 52
43 -> 53
```

The following example shows how to remove all VLAN translation mappings from the existing translation configuration:

```
switch(config)# interface overlay 5
switch(config-if-overlay)# no otv vlan mapping
Removing all translations
switch(config-if-overlay)# show otv vlan-mapping
Original VLAN -> Translated VLAN
--------------------------------
```

Configuration Examples for Dedicated Data Broadcast Forwarding

The following example shows how to configure a dedicated broadcast-group address for an OTV network:

```
switch# configure terminal
switch(config)# feature otv
switch(config)# interface overlay 5
switch(config-if-overlay)# otv broadcast-group 224.2.1.0
switch(config-if-overlay)# show otv
OTV Overlay Information
Site Identifier 0000.0000.0002
Overlay interface Overlay5
VPN name : Overlay5
VPN state : UP
Extended vlans : 25-150 251-327 (Total:203)
Control group : 224.1.1.0
```

```
Data group range(s) : 232.1.0.0/24
Broadcast group : 224.2.1.0
Join interface(s) : Po21 (2.100.21.1)
Site vlan : 1000(up)
AED-Capable : Yes
Capability : Multicast-Reachable
```

The following example shows that the broadcast-group address defaults to the control-group address when the broadcast-group address configuration is removed:

```
switch# configure terminal
switch(config)# feature otv
switch(config)# interface overlay 5
switch(config-if-overlay)# no otv broadcast-group 224.2.1.0
switch(config-if-overlay)# show otv
OTV Overlay Information
Site Identifier 0000.0000.0002
Overlay interface Overlay5
VPN name : Overlay5
VPN state : UP
Extended vlans : 25-150 251-327 (Total:203)
Control group : 224.1.1.0
Data group range(s) : 232.1.0.0/24
Broadcast group : 224.1.1.0
Join interface(s) : Po21 (2.100.21.1)
Site vlan : 1000(up)
AED-Capable : Yes
Capability : Multicast-Reachable
```

Configuration Example for OTV Fast Convergence

The following example shows how to enable OTV fast convergence by configuring an SVI on an OTV site VLAN:

```
switch# configure terminal
switch(config)# feature bfd
switch(config)# feature interface-vlan
switch(config)# interface vlan 2500
switch(config-if)# no ip redirects
switch(config-if)# ip address 172.1.2.1/24
switch(config-if)# no shutdown
```

Configuration Example for Fast Failure Detection

The following example shows how to configure fast failure detection in an OTV site VLAN. The output of the **show** commands displays that the BFD adjacency is "Up" between switches in the same site and the BFD configuration is applied on OTV switches in the same site:

```
switch# configure terminal
switch(config)# otv-isis default
switch(config-router)# track-adjacency-nexthop
switch(config-router)# exit
switch(config)# otv site-vlan 5
switch(config-site-vlan)# otv isis bfd
switch(config-site-vlan)# show bfd neighbors
OurAddr   NeighAddr LD/RD                RH/RS Holdown(mult) State Int      Vrf
172.1.1.1 172.1.1.2 1107296329/1107296399 Up    5462(3)       Up    Vlan2500 default
switch(config-site-vlan)# show otv isis track-adjacency-nexthop
OTV-IS-IS process: default
```

```
    OTV-ISIS adjs for nexthop: 10.0.1.1, VRF: default
       Hostname: 0022.557a.3040, Overlay: Overlay4
       Hostname: 0022.557a.3040, Overlay: Overlay3
       Hostname: 0022.557a.3040, Overlay: Overlay2
       Hostname: 0022.557a.3040, Overlay: Overlay1
switch(config-site-vlan)# show otv isis site
OTV-ISIS site-information for: default

  Level      Metric   CSNP  Next CSNP  Hello   Multi   Next IIH
  1        16777214     10  Inactive       3   20      0.292879

  Level  Adjs   AdjsUp Pri  Circuit ID         Since
  1         1        1  64  0022.557a.3043.01  00:15:01

  BFD: Enabled [IP: 5.5.5.11]


OTV-IS-IS site adjacency local database:

 SNPA           State Last Chg Hold     Fwd-state Site-ID        Version BFD
 0022.557a.3043 UP    00:15:01 00:01:00 DOWN      000a.000a.000a 3       Enabled [Nbr IP:
5.5.5.12]

OTV-IS-IS Site Group Information (as in OTV SDB):

  SystemID: 0022.557a.3040, Interface: site-vlan, VLAN Id: 5, VLAN: Up

  Overlay     State   Next IIH   Int   Multi
  Overlay1    Up      0.290956   3     20
  Overlay2    Up      0.289360   3     20
  Overlay3    Up      0.287777   3     20
  Overlay4    Up      0.286202   3     20

  Overlay    Active SG       Last CSNP            CSNP Int  Next CSNP
  Overlay1   239.1.1.1       ffff.ffff.ffff.ff-ff 01:15:21  Inactive
  Overlay2   239.1.1.2       ffff.ffff.ffff.ff-ff 01:15:21  Inactive
  Overlay3   0.0.4.0         ffff.ffff.ffff.ff-ff 01:15:21  Inactive
  Overlay4   0.0.5.0         ffff.ffff.ffff.ff-ff 01:15:21  Inactive

  Neighbor SystemID: 0022.557a.3043
  IPv4 site groups:
    0.0.4.0
    0.0.5.0
    239.1.1.1
    239.1.1.2
```

Configuration Example for Disabling Tunnel Depolarization with IP Pools

The following examples show the how to disable and verify tunnel depolarization on an overlay network:

```
switch# configure terminal
switch(config)# feature otv
switch(config)# otv depolarization disable
switch(config)# exit


switch# show otv

OTV Overlay Information
Site Identifier 0000.0000.0001
```

```
Overlay interface Overlay1

 VPN name            : Overlay1
 VPN state           : UP
 Extended vlans      : 10-11 101-102 (Total:4)
 Control group       : 239.1.1.1
 Data group range(s) : 232.10.10.0/28
 Broadcast group     : 239.1.1.1
 Join interface(s)   : Eth1/13 (20.0.0.100)
   Secondary IP Addresses:  20.0.0.101
 Site vlan           : 10 (up)
 AED-Capable         : No (ISIS Ctrl Group Sync Pending)
 Capability          : Multicast-Reachable


switch# show otv adjacency detail

Overlay Adjacency database

Overlay interface Overlay1
Hostname                         System-ID      Dest Addr       Up Time   State
meN7K-1-N7K-B1                   64a0.e741.84c2 20.0.0.2        00:10:24  UP
```

Monitoring OTV

To monitor OTV, perform one of the following tasks:

Command	Purpose
show otv orib clients	Displays information about the ORIB clients.
show otv route [**overlay** *interface* | **vlan** *vlan-id* | **vpn** *vpn-name*]	Shows unicast MAC routes.
show otv mroute [**overlay** *interface* | **vlan** *vlan-id* | **vpn** *vpn-name*]	Displays information about multicast MAC routes.
show otv statistics multicast vlan *vlan-id*	Shows OTV statistics.
show otv isis statistics {***** | **overlay** *interface*}	Shows statistics for the OTV control-plane protocol.
show otv isis track-adjacency-nexthop	Displays the OTV IS-IS next-hop adjacencies.

To clear OTV information, perform the following task:

Command	Purpose
clear otv isis statistics {***** | **overlay** *interface*}	Clears OTV statistics.

Additional References

This section includes additional information related to implementing OTV.

Related Documents

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

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.
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.

Feature Name	Releases	Feature Information
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, on page 2
Configuring the Site VLAN and Site Identifier
Assigning the Extended VLAN Range