



Cisco Nexus 7000 Series OTV Quick Start Guide

Cisco Nexus 7000 Series OTV Quick Start Guide	2
Configuring an OTV Network	2
Configuration Examples for OTV	5
Configuration Example for Load Balancing	5
Verifying the OTV Configuration	6
Default Settings for OTV	7
Prerequisites for OTV	8
Guidelines and Limitations for OTV	8
Communications, Services, and Additional Information	10
Full Cisco Trademarks with Software License	?

Revised: August 8, 2022

Cisco Nexus 7000 Series OTV Quick Start Guide

This guide describes the basic configuration for Overlay Transport Virtualization (OTV) on Cisco NX-OS devices. OTV is a MAC-in-IP method that extends Layer 2 connectivity across a transport network infrastructure. OTV uses MAC address-based routing and IP-encapsulated forwarding across a transport network to provide support for applications that require Layer 2 adjacency, such as clusters and virtualization. You deploy OTV on the edge devices in each site.

Configuring an OTV Network

This procedure provides a step-by-step walkthrough of how to configure an OTV edge device. This procedure includes the configuration for a physical interface that acts as the join interface, and a VLAN that is extended over the overlay network. This procedure should be applied to each OTV Edge Device.

A basic OTV configuration requires IP connectivity of the OTV Edge Device to the transport network connecting all OTV edge devices. Also, VLANs that will be extended using OTV have to be configured on the OTV edge devices. Once these prerequisites are fulfilled, you can configure OTV.

Configuring a Physical Interface

This procedure shows how to configure a physical interface for IP connectivity to the IP transport network.

Procedure

Step 1 `interface ethernet interface`

Example:

```
switch(config)# interface ethernet 2/1
switch(config-if)#
```

Enters interface configuration mode for the physical interface that will become the join interface for the OTV network. OTV uses this interface to reach the transport network. This interface can be a Layer 3 interface, Layer 3 port channel or subinterface on a Layer 3 interface or Layer 3 port channel.

Step 2 Required: `ip address ipaddress/mask`

Example:

```
switch(config-if)# ip address 192.0.2.1/24
```

Configures the IP address and network mask length in dotted decimal notation on the physical interface.

Step 3 `ip igmp version 3`

Example:

```
switch (config-if)# ip igmp version 3
```

Enables IGMPv3 on this physical interface. This is a requirement for this interface when it becomes the OTV join interface.

What to do next

You must configure static or dynamic routing and enable the interface with the **no shutdown** command.

Configuring a VLAN

This procedure shows how to configure a VLAN.

Procedure

vlan *vlan-id*

Example:

```
switch (config-if)# vlan 5
switch (config-vlan)#
```

Configures a VLAN.

The range is from 1 to 3967 and from 4048 to 4093. The default is 1.

Configuring OTV

This procedure shows how to configure OTV.

Before you begin

Ensure that you have configured a physical interface that provides connectivity to the IP core and that you have configured the VLANs that will be extended over the OTV network.

Procedure

Step 1 **feature otv**

Example:

```
switch (config)# feature otv
```

This command enables the OTV feature on this device. You may try this feature for a limited time without the license by using the license grace period.

Step 2 **otv site-vlan** *vlan-id*

Example:

```
switch (config)# otv site-vlan 10
```

Note VLAN 1 is the default site VLAN. We recommend that you use a dedicated VLAN as site VLAN. Ensure that the site VLAN is active on at least one of the edge device ports. The site VLAN should not be extended across the overlay.

This command configures a VLAN that all local edge devices in a site communicate on. OTV uses this site VLAN to send hello messages that other configured edge devices in the site respond to. OTV uses a VLAN hashing algorithm to select the authoritative edge device from one of these local site edge devices. OTV can load balance traffic over different edge devices for the VLANs that each edge device is authoritative for.

You should configure the site 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 Required: **otv site-identifier** *id*

Example:

```
switch(config)# otv site-identifier 0018.g957.6rk0
```

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 0x0 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.

Step 4 **interface overlay** *interface*

Example:

```
switch(config)# interface overlay 1  
switch(config-if-overlay)#
```

Creates an OTV overlay interface and enters interface configuration mode. The overlay interface is a logical multi-access multicast-capable interface that encapsulates Layer 2 frames in IP unicast or multicast headers.

Step 5 Required: **otv control-group** *mcast-address*

Example:

```
switch(config-if-overlay)# otv control-group 239.1.1.1
```

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 and must be an ASM or Bidir group.

Step 6 Required: **otv data-group** *mcast-range1 [mcast-range2...]*

Example:

```
switch(config-if-overlay)# otv data-group 232.1.1.0/28
```

Configures one or more ranges of local IPv4 multicast group prefixes used for multicast data traffic. The multicast group address is an IPv4 address in dotted decimal notation. These prefixes are SSM groups. A subnet mask is used to indicate ranges of addresses. You can define up to eight data-group ranges.

Step 7 Required: **otv join-interface** *interface*

Example:

```
switch(config-if-overlay)# otv join-interface ethernet 2/1
```

Joins the OTV overlay interface with a Layer 3 interface, Layer 3 port channel or sub-interface on a Layer 3 interface or Layer 3 port channel. You must configure an IP address and enable IGMPv3 on this interface.

You can specify only one join interface per Overlay. You can decide to use one of the following methods:

- A single join interface, which is shared across multiple Overlays.
- A different join interface for each Overlay thus increasing the OTV reliability.

For a higher resiliency, you can use a port-channel but it is not mandatory. There are no requirements either in terms of 1GE vs 10GE or in terms of Dedicated vs Shared mode.

Note The join-interface must belong to the default-VRF.

Step 8 Required: **otv extend-vlan** *vlan-range*

Example:

```
switch(config-if-overlay)# otv extend-vlan 2,5-34
```

Extends a range of VLANs over this overlay interface and enables OTV advertisements for these VLANs. OTV will not forward Layer 2 packets for VLANs not in the extended VLAN range for the overlay interface. You can assign a VLAN to only one overlay interface.

The *vlan-range* is from 1 to 3967 and from 4048 to 4093.

What to do next

You need to enable the interface with the **no shutdown** command.

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
```

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

```

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.
<code>show otv [overlay interface] [vlan [vlan-range] [authoritative detail]]</code>	Displays information about VLANs that are associated with an overlay interface.
<code>show otv isis site [database statistics]</code>	Displays the BFD configuration state on both local and neighboring edge devices.
<code>show otv site [all]</code>	Displays information about the local site.

Command	Purpose
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.

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

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.

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).
- 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.

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2016 Cisco Systems, Inc. All rights reserved.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA 95134-1706
USA

Asia Pacific Headquarters
CiscoSystems(USA)Pte.Ltd.
Singapore

Europe Headquarters
CiscoSystemsInternationalBV
Amsterdam,TheNetherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.