

# Overlay Transport Virtualization Best Practices Guide

## Configuration Guide

November 2017

---

# Contents

<b>Introduction</b> .....	<b>3</b>
<b>Recommended Reading</b> .....	<b>3</b>
<b>What Hardware Supports OTV Encapsulation?</b> .....	<b>3</b>
<b>Should I Use Multicast or Unicast for My Transport?</b> .....	<b>3</b>
<b>Is Multihoming Recommended?</b> .....	<b>3</b>
<b>Can I Mix Cisco Nexus 7000 Series and Cisco ASR 1000 Series Devices for OTV?</b> .....	<b>4</b>
<b>Is the site-id Command a Required Configuration?</b> .....	<b>4</b>
<b>Can I Mix M and F Modules in My OTV VDC?</b> .....	<b>4</b>
<b>What Type of Scale Can OTV Accommodate?</b> .....	<b>4</b>
<b>What If My Link Between Data Centers Is Slower Than My Line-Card Ports?</b> .....	<b>4</b>
<b>Does OTV Support Packet Fragmentation?</b> .....	<b>4</b>
<b>Does OTV Require a Separate VDC?</b> .....	<b>5</b>
<b>Where Is My Spanning-Tree Root with OTV?</b> .....	<b>5</b>
<b>Are There Any Caveats with FabricPath?</b> .....	<b>5</b>
<b>Can I Flood Unknown Unicasts?</b> .....	<b>6</b>
<b>OTV in Enterprises</b> .....	<b>6</b>
<b>OTV in MSDCs</b> .....	<b>6</b>
<b>How Do I Isolate FHRP?</b> .....	<b>6</b>
<b>How Does OTV Interact with vPC</b> .....	<b>7</b>
<b>OTV and MTUs</b> .....	<b>7</b>
<b>OTV UDP Encapsulation</b> .....	<b>8</b>
<b>What Happens to My Ingress Routing with OTV?</b> .....	<b>8</b>
<b>Can I Translate VLAN X in Data Center 1 to VLAN Y in Data Center 2?</b> .....	<b>10</b>
<b>What If I Need to Run QoS Against My Multicast Core?</b> .....	<b>10</b>
<b>Microsoft Network Load Balancing and Its Implications for OTV</b> .....	<b>10</b>
NLB in Unicast Mode with OTV.....	11
NLB in Multicast Mode with OTV .....	11
<b>Conclusion</b> .....	<b>11</b>
<b>Configuration Example 1: Multicast and Unicast OTV Configuration</b> .....	<b>11</b>
<b>Configuration Example 2: FHRP Isolation</b> .....	<b>25</b>
<b>Configuration Example 3: Direct VLAN Translation</b> .....	<b>27</b>
<b>Configuration Example 4: VLAN Translation Through a Transit VLAN</b> .....	<b>28</b>

---

## Introduction

This document covers common best practices for deploying the Cisco® innovative LAN extension technology called Overlay Transport Virtualization (OTV). OTV provides Layer 2 LAN extension over Layer 2, Layer 3, or label switched (Multiprotocol Label Switching [MPLS])-based networks. This document expects that you have a basic understanding of OTV as a technology and provides a Q&A style format.

## Recommended Reading

Before you can understand the best practices of an OTV deployment, you need to understand the basic technology, terminology, and supported hardware. Following are two links that can help you:

- <https://www.cisco.com/en/US/netsol/ns1153/index.html>
- [https://www.cisco.com/en/US/docs/solutions/Enterprise/Data\\_Center/DCI/whitepaper/DCI\\_1.html](https://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DCI/whitepaper/DCI_1.html)

## What Hardware Supports OTV Encapsulation?

Cisco has two main products that currently support OTV: Cisco Nexus® 7000 Series Switches and Cisco ASR 1000 Series Aggregation Services Routers. Although the features, scale, and convergence times continue to improve, following are the minimum requirements to run OTV on each platform:

- Cisco Nexus 7000 Series and Cisco Nexus 7700 platform:
  - Any M-Series (Cisco Nexus 7000 Series) or F3 (Cisco Nexus 7000 Series or 7700 platform) line card for encapsulation
  - Cisco NX-OS Release 5.0(3) or later (Cisco Nexus 7000 Series) or Cisco NX-OS Release 6.2(2) or later (Cisco Nexus 7700 platform)
  - Transport Services license
- Cisco ASR 1000 Series:
  - Cisco IOS® XE Software Release 3.5 or later
  - Advanced IP Services or Advanced Enterprise Services license

## Should I Use Multicast or Unicast for My Transport?

OTV offers multicast and unicast as transports between sites. Multicast is the preferred transport because of its flexibility and smaller overhead when communicating with multiple sites. If you are planning only two or three sites, then unicast works just as well without losing any features or functions.

The Cisco Nexus 7000 Series and Cisco ASR 1000 Series both support multicast and unicast cores. For unicast cores the Nexus 7000 Series requires Cisco NX-OS Release 5.2(1) or later and the Cisco ASR 1000 Series requires Cisco IOS-XE 3.9 or later.

## Is Multihoming Recommended?

If possible, multihoming is always recommended because it adds another layer of redundancy and scalability. When in a multihomed site, all odd VLANs are on one Cisco Nexus 7000 Series Switch, and all even VLANs are on the other Cisco Nexus 7000 Series Switch. Please note that the multihoming of Cisco ASR 1000 Series and Cisco Nexus 7000 Series devices within a single site is not supported.

---

## Can I Mix Cisco Nexus 7000 Series and Cisco ASR 1000 Series Devices for OTV?

Mixing the two types of devices is not supported at this time when the devices will be placed within the same site. However, using Cisco Nexus 7000 Series Switches in one site and Cisco ASR 1000 Series routers at another site for OTV is fully supported. For this scenario, please keep the separate scalability numbers in mind for the two different devices, because you will have to account for the lowest common denominator.

## Is the `site-id` Command a Required Configuration?

The **site-id** command was introduced as a way to harden multihoming for OTV. It is a configurable option that must be the same for devices within the same data center and different between any devices that are in different data centers. It specifies which site a particular OTV device is in so that two OTV devices in different sites cannot join each other as a multihomed site.

In earlier releases of OTV, the **site-id** command was not a required configuration, but starting in Cisco NX-OS Release 5.2 it is mandatory.

## Can I Mix M and F Modules in My OTV VDC?

Starting in the Cisco NX-OS Release 6.2(2) on the Cisco Nexus 7000 Series, F1 and F2e modules can be internal interfaces for OTV. These modules cannot perform OTV functions themselves and can be used only as internal interfaces.

## What Type of Scale Can OTV Accommodate?

The scale numbers for OTV are always improving as we evolve this technology. The latest numbers are available here:

[https://www.cisco.com/en/US/docs/switches/datacenter/sw/verified\\_scalability/b\\_Cisco\\_Nexus\\_7000\\_Series\\_NX-OS\\_Verified\\_Scalability\\_Guide.html#reference\\_18192F87114B45D9A40A41A0DEF3F74D](https://www.cisco.com/en/US/docs/switches/datacenter/sw/verified_scalability/b_Cisco_Nexus_7000_Series_NX-OS_Verified_Scalability_Guide.html#reference_18192F87114B45D9A40A41A0DEF3F74D).

Please note that while the number of VLANs supported is greater than 64 VLANs, the configuration line only supports up to 64 physical numbers listed at a time. For instance `1-64` uses 2 total numbers in the configuration line, while `1, 2, 3, 4, 5` uses 5 numbers.

## What If My Link Between Data Centers Is Slower Than My Line-Card Ports?

Some interconnects between data centers might have a slower link speed than is supported on the M-series cards being used for the OTV encapsulation within the data center. In this case, you need to make sure that you do not overrun the transport between the data centers and drop packets. We recommend quality of service (QoS) to resolve this problem so that the larger link speeds (10 or 40 Gbps) do not go above the link speed between data centers. If your link between data centers is only 1 Gbps and your M-series line card has 10-Gbps ports, then you should use QoS to ensure that the join interface in your OTV virtual device context (VDC) is limited to 1 Gbps of throughput.

A simple single-rate policing configuration can solve this problem, and example configurations are available on Cisco.com.

## Does OTV Support Packet Fragmentation?

At this time, OTV on the Cisco Nexus 7000 Series does not allow packet fragmentation. However, the Cisco ASR 1000 Series does support this feature. It is important to ensure that if any site running OTV is using a Cisco Nexus 7000 Series Switch for encapsulation, the encapsulated packets are not fragmented on any device.

## Does OTV Require a Separate VDC?

OTV currently enforces switch-virtual-interface (SVI) separation for the VLANs being extended across the OTV link, meaning that OTV is usually in its own VDC. With the VDC license on the Cisco Nexus 7000 Series you have the flexibility to have SVIs in other VDCs and have a dedicated VDC for OTV functions.

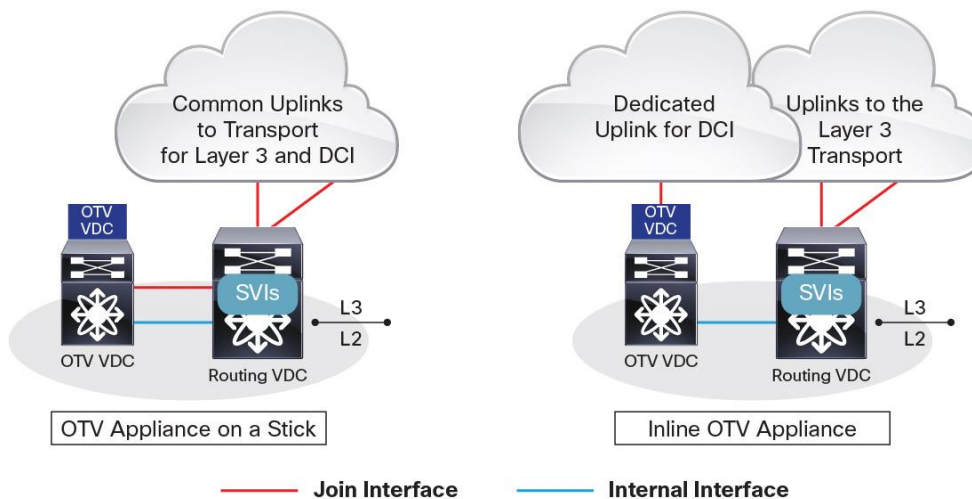
When using a separate VDC for OTV functions, two deployment models are possible ([Figure 1](#)):

- OTV Appliance on a Stick
- Inline OTV Appliance

In both scenarios, OTV functions remain the same. The only difference is where the join interface is placed. In the OTV Appliance on a Stick model, the join interface connects back through the VDCs that have SVIs on them. The inline mode requires that the join interface have a dedicated link out to the DCI transport.

In most cases, the OTV Appliance on a Stick model is used, because it requires no network redesigns or re-cabling if OTV is activated or deactivated for any reason.

**Figure 1.** VDC Deployment Options for OTV



## Where Is My Spanning-Tree Root with OTV?

One of the main advantages of OTV is the fault-domain isolation feature. With OTV, fault domains are actually isolated and separate from each other without the requirement of any additional configuration. With this isolation in place, there is no change to the spanning-tree root or configuration. Each site has its own root and remains separate from other sites with no intervention or planning required by the operators.

## Are There Any Caveats with FabricPath?

Because OTV encapsulation is done on M-series modules, OTV cannot read FabricPath packets. Because of this restriction, terminating FabricPath and reverting to Classical Ethernet where the OTV VDC resides is necessary.

In addition, when running FabricPath in your network, we highly recommend that you use the **spanning-tree domain** command on all devices that are participating in these VLANs. This command speeds up convergence times greatly.

## Can I Flood Unknown Unicasts?

Yes, with the Cisco NX-OS Release 6.2(2), we now support selective unknown unicast flooding based on the MAC address. This feature is especially important for networks that use Microsoft's Network Load Balancer. In order to use this feature, the following one-line configuration must be added:

```
OTV-VDC(config)# otv flood mac 0000.2101.1111 vlan 72
```

Please note that both the MAC address for flooding and the VLAN in which it resides are needed in order to configure this feature.

## OTV in Enterprises

This model is what most OTV deployments use to extend Layer 2 between two or more physically separate sites. The vast majority of this guide covers this deployment scenario, so more details are not required.

## OTV in MSDCs


OTV can not only extend Layer 2 between physically separate data centers, it also can split large data centers to help mitigate any failures or storms that might occur. Some Massively Scalable Data Centers (MSDCs) use OTV to logically separate their networks. With OTV, this design isolates failure domains so any loops or failures do not propagate to the whole data center. As with the traditional deployment model used in enterprises, the configuration is the same when using OTV to logically split a larger data center.

## How Do I Isolate FHRP?

Another capability introduced by OTV is to filter First Hop Redundancy Protocol (FHRP - Hot Standby Router Protocol [HSRP], Virtual Router Redundancy Protocol [VRRP], and so on) messages across the logical overlay. This filtering is required to allow for the existence of the same default gateway in different locations and optimize the outbound traffic flows (server to client direction). [Figure 2](#) highlights the root of the problem.


Given that the same VLAN/IP subnet is available in different sites, the free exchange of FHRP messages across the OTV connection would lead to the election of a single default gateway, forcing traffic to follow a suboptimal path to reach the default gateway (in the site where it is deployed) each time it is required to be routed outside the subnet and the server is located in a different site.

**Figure 2.** Suboptimal Outbound Routing



[Figure 3](#) shows the deployment of independent default gateways in each data center site to optimize and localize routing of outbound traffic flows.

**Figure 3.** FHRP Isolation with OTV



It is critical that you enable the filtering of FHRP messages across the overlay because it allows the use of the same FHRP configuration in different sites. The end result is that the same default gateway is available, and it is characterized by the same virtual IP and virtual MAC addresses in each data center. Thus the outbound traffic will be able to follow the optimal and shortest path, always using the local default gateway.

The appendix has a sample configuration that can be copied into your Nexus 7000 Series in order to utilize this feature.

### How Does OTV Interact with vPC

Using Virtual Port Channels (vPCs) and OTV together provides an extra layer of resiliency and is thus recommended as a best practice. Because OTV is usually run in its own VDC, a vPC between the OTV and aggregation VDCs in a dual-homed scenario is the most common application. There are no constraints or special requirements when running both together. For more information about vPCs, please refer to the vPC Best Practices Guide or configuration guides found on [cisco.com](http://cisco.com).

### OTV and MTUs

OTV adds 42 bytes in the IP header packets, thus requiring a larger maximum transmission unit (MTU) for traffic to pass. It is also worth noting that OTV on the Cisco Nexus 7000 Series does not support fragmentation, so the larger MTU must be considered. There are two ways to solve this problem:

1. Configure a larger MTU on all interfaces where traffic will be encapsulated, including the join interface and any links between the data centers that are in an OTV transport.
2. Lower the MTU on all servers so that the total packet size does not exceed the MTU of the interfaces where traffic is encapsulated.

---

## OTV UDP Encapsulation

The OTV User Datagram Protocol (UDP) header encapsulation mode was introduced in the Cisco Nexus 7000 Series and 7700 platform switches that use F3 or M3 line cards and NX-OS 7.2.0. In this software release, the forwarding engine for the control-plane and data-plane packets supports UDP encapsulation over IP over Ethernet. The control and data paths use UDP headers for multicast and unicast core routing. The Internet Assigned Numbers Authority (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 Virtual Extensible LAN (VXLAN) header defined in IETF RFC 7348.

UDP encapsulation helps use more links in the core network because the UDP source port is varied automatically.

By default, the encapsulation format is MPLS generic routing encapsulation (GRE). You can configure the OTV encapsulation format as UDP using the `otv encapsulation-format ip udp` command.

## What Happens to My Ingress Routing with OTV?

As virtual machines or workloads are moved to different data centers, routing outside clients to these machines can become tricky. OTV in itself is a Layer 2 extension and cannot solve routing problems.

Consider the following example (refer to [Figure 4](#)): We have a client coming into Data Center 1 from the Internet and talking to the virtual machine highlighted in the figure. The client is routed directly to Data Center 1 instead of Data Center 2 upon initial communication, so the routing between the client and workload is optimal.


Later in the communication cycle, the virtual machine is VMotioned (or moved) to Data Center 2 because maintenance is being performed at Data Center 1 ([Figure 5](#)). Now when the client tries to communicate with the virtual machine, the client still goes to Data Center 1, where the traffic is then directed through the OTV link to Data Center 2.

This example shows how the routing is now un-optimized and requires traversal through the OTV link. As mentioned previously, OTV cannot solve this problem by itself. Some sort of routing protocol or the Cisco Locator/ID Separation Protocol (LISP) needs to be integrated into this design to resolve this suboptimal routing if it is a problem for your data centers.


For more information about LISP, please visit [https://lisp.cisco.com/lisp\\_over.html](https://lisp.cisco.com/lisp_over.html).



**Figure 4.** Optimal Routing



**Figure 5.** Suboptimal Routing



---

## Can I Translate VLAN X in Data Center 1 to VLAN Y in Data Center 2?

Yes, you can translate VLAN X in Data Center 1 to VLAN Y in Data Center 2. Cisco NX-OS 6.2(2) introduces a VLAN translation feature. Two types of VLAN translation are supported:

- Direct VLAN translation
- VLAN translation through a transit VLAN

Because VLAN translation is performed on both ingress and egress traffic, the configuration for direct translation requires configuration at only one site. In this scenario, if you want VLAN 100 in Data Center 1 to be translated to VLAN 200 in Data Center 2, you need to specify the configuration at only one of the sites. Configuration Example 3 shows this scenario.

For scenarios in which the VLAN number may differ among three or more sites, translation through a transit VLAN is necessary. In this scenario, you use a VLAN that is only in the core, and each data center will translate that unique VLAN number to its local VLAN. Configuration Example 4 shows this scenario, using VLAN 400 as the transit VLAN and VLANs 100, 200, and 300 as the local VLANs for Data Centers 1, 2, and 3. The local VLAN number will always be the first one listed in the configuration line.

## What If I Need to Run QoS Against My Multicast Core?

Cisco NX-OS 6.2(2) introduces a new optional feature: dedicated broadcast group. This feature allows the administrator to configure a dedicated broadcast group to be used by OTV in the multicast core network. By separating the two multicast addresses, you can now enforce different QoS treatments for data traffic and control traffic.

## Microsoft Network Load Balancing and Its Implications for OTV

Network load balancing (NLB) is a clustering technology included in the Microsoft Windows Advanced Server and Datacenter Server operating systems. It enhances the scalability and availability of mission-critical, TCP/IP-based services, such as web servers, terminal services, virtual private networking, and streaming media servers. NLB technology is used to distribute client requests across a set of servers. This component runs within cluster hosts as part of the Windows operating system. NLB has three primary configuration modes: unicast, multicast, and Internet Group Management Protocol (IGMP) multicast. The widely deployed modes for NLB are Unicast and Multicast which are discussed below

- **Unicast mode** assigns the cluster a virtual IP address and virtual MAC address. This method relies on unknown unicast flooding. Because the virtual MAC address is not learned on any switch ports, traffic destined to the virtual MAC address is flooded within the VLAN. As a result, all clustered servers receive traffic destined for the virtual MAC address. One downside to this method is that all devices in the VLAN receive this traffic. The only way to mitigate this behavior is to limit the NLB VLAN to only the NLB server interfaces to avoid flooding to interfaces that should receive the traffic.
- **Multicast mode** assigns a unicast IP address to a non-IANA multicast MAC address (03xx.xxxx.xxxx). IGMP snooping does not dynamically program this address, which results in flooding of the NLB traffic in the VLAN.

---

## NLB in Unicast Mode with OTV

When NLB is used in the unicast mode and deployed concurrently with OTV, it relies on unknown unicast flooding. OTV as a technology does not allow unknown unicast flooding across sites. 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. To use NLB in unicast mode across OTV sites, you must enable selective unicast flooding in OTV. This feature enables flooding for the specified destination MAC address to all other edge devices in the OTV overlay network.

See the configuration guide at the following URL for details about how to enable the selective unicast flooding feature in OTV: [https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/OTV/config\\_guide/b\\_Cisco\\_Nexus\\_7000\\_Series\\_NX-OS\\_OTV\\_Configuration\\_Guide/adv-otv.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/OTV/config_guide/b_Cisco_Nexus_7000_Series_NX-OS_OTV_Configuration_Guide/adv-otv.html).

## NLB in Multicast Mode with OTV

When NLB is used in multicast mode, the NLB sources the Address Resolution Protocol (ARP) reply packets with an IP address, but the MAC address is a multicast MAC address. The source MAC address of the outer (Layer 2) header is the MAC address of the network interface card (NIC) of the server. However, the sender MAC address (inner ARP header) is the multicast MAC address for the cluster. OTV **arp nd-cache** (which is enabled by default) performs a form of ARP inspection and treats the ARP packets with intelligence. Consequently, it identifies this mismatch and drops these packets. Hence, clients at the different sites in OTV cannot resolve the ARP of the NLB server. To use NLB in multicast mode across sites in OTV, you have to disable the default behavior of ARP local caching on OTV. To achieve this, disable **arp nd-cache**.

See the configuration guide at the following URL for details about how to disable the **arp nd-cache** feature in OTV: [https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/OTV/config\\_guide/b\\_Cisco\\_Nexus\\_7000\\_Series\\_NX-OS\\_OTV\\_Configuration\\_Guide/adv-otv.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/OTV/config_guide/b_Cisco_Nexus_7000_Series_NX-OS_OTV_Configuration_Guide/adv-otv.html).


## Conclusion

This guide discusses the deployment scenarios, benefits, and drawbacks of the different possible designs for OTV. It also covers some of the newer features and questions that users have about the Cisco Nexus 7000 Series. This document will be updated as needed.

## Configuration Example 1: Multicast and Unicast OTV Configuration

The following is an example of a multicast-enabled transport running multihomed OTV across two sites. The design uses a separate OTV VDC in the inline mode previously shown. Unicast configuration differences are also shown to highlight the minor differences in actual configuration ([Figure 6](#)).

**Figure 6.** Testbed Configuration Example



### NR-n7k-a

```

version 5.1(1a)
license grace-period
hostname NR-n7k-a
!Following configuration set is not terminated by a newline
no vdc combined-hostname
vdc NR-n7k-a id 1
    limit-resource vlan minimum 16 maximum 4094
    limit-resource monitor-session minimum 0 maximum 2
    limit-resource monitor-session-erspan-dst minimum 0 maximum 23
    limit-resource vrf minimum 2 maximum 1000
    limit-resource port-channel minimum 0 maximum 768
    limit-resource u4route-mem minimum 32 maximum 32
    limit-resource u6route-mem minimum 16 maximum 16
    limit-resource m4route-mem minimum 58 maximum 58
    limit-resource m6route-mem minimum 8 maximum 8
vdc East-a id 2
    allocate interface Ethernet1/26, Ethernet1/28, Ethernet1/30, Ethernet1/32
    allocate interface Ethernet2/6, Ethernet2/21, Ethernet2/23, Ethernet2/25,
Ethernet2/27
    boot-order 3
    limit-resource vlan minimum 16 maximum 4094
    limit-resource monitor-session minimum 0 maximum 2
    limit-resource monitor-session-erspan-dst minimum 0 maximum 23
    limit-resource vrf minimum 2 maximum 1000
    limit-resource port-channel minimum 0 maximum 768
    limit-resource u4route-mem minimum 8 maximum 8
    limit-resource u6route-mem minimum 4 maximum 4
    limit-resource m4route-mem minimum 8 maximum 8

```

```
limit-resource m6route-mem minimum 2 maximum 2
feature telnet
cfs eth distribute
feature ospf
feature pim
feature udld
feature interface-vlan
feature hsrp
feature lacp
feature vpc
username adminbackup password 5 ! role network-operator
username admin password 5 $1$zOYf1VLm$Wh2/WnLdQDN894obifIpZ1 role network-admin
username adminbackup password 5 ! role network-operator
no password strength-check
ip domain-lookup
snmp-server user admin network-admin auth md5
0x88018226be1701759b4301a3c0519193 priv 0x88018226be1701759b4301a3c0
519193 localizedkey
vrf context management
    ip route 0.0.0.0/0 172.26.245.1
vlan 1-4,99-199
spanning-tree vlan 99-199 priority 4096
vpc domain 1
    role priority 4086
    peer-keepalive destination 172.26.245.10 source 172.26.245.20
interface Vlan100
    no shutdown
    no ip redirects
    ip address 10.100.1.4/24
    ip ospf network broadcast
    ip ospf passive-interface
    ip router ospf 1 area 0.0.0.0
    hsrp 100
        preempt delay minimum 60
        priority 40
        timers 1 3
        ip 10.100.1.1
<SNIP>
interface Vlan199
    no shutdown
    no ip redirects
    ip address 10.199.1.4/24
    ip ospf network broadcast
    ip ospf passive-interface
    ip router ospf 1 area 0.0.0.0
```

```
hsrp 199
  preempt delay minimum 60
  priority 40
  timers 1 3
  ip 10.199.1.1
interface port-channel1
  description [ To N5K Access ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  spanning-tree port type network
  vpc 1
interface port-channel10
  description [ To N7K-b ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  spanning-tree port type network
  vpc peer-link
interface port-channel20
  description [ To this OTV VDC ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  vpc 20
interface port-channel30
  description [ To this OTV VDC ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  vpc 30
interface Ethernet1/1
  description [ To N7K-b ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  channel-group 10 mode active
  no shutdown
interface Ethernet1/9
  description [ To N5K Access ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  channel-group 1 mode active
  no shutdown
```

```
interface Ethernet2/1
  description [ To N7K-b ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  channel-group 10 mode active
  no shutdown
interface Ethernet2/5
  description [ To the OTV Join-Interface ]
  uddl aggressive
  ip address 172.26.255.93/30
  ip ospf network point-to-point
  ip router ospf 2 area 0.0.0.0
  ip pim sparse-mode
  ip igmp version 3
  no shutdown
interface Ethernet2/17
  description [ To Core-A ]
  ip address 172.26.255.70/30
  ip ospf network point-to-point
  ip router ospf 2 area 0.0.0.0
  ip pim sparse-mode
  no shutdown
interface Ethernet2/18
  description [ To Core B ]
  ip address 172.26.255.78/30
  ip ospf network point-to-point
  ip router ospf 2 area 0.0.0.0
  ip pim sparse-mode
  no shutdown
interface Ethernet2/20
  description [ To this OTV VDC ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  channel-group 20 mode active
  no shutdown
interface Ethernet2/22
  description [ To other OTV VDC ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  channel-group 30 mode active
  no shutdown
interface mgmt0
```

```

    ip address 172.26.245.20/24
interface loopback2
    ip address 172.26.255.153/32
    ip router ospf 2 area 0.0.0.0
cli alias name sw switchto vdc east-a
line console
    exec-timeout 0
    speed 115200
line vty
    exec-timeout 0
boot kickstart bootflash:/n7000-s1-kickstart.5.1.1a.gbin sup-1
boot system bootflash:/n7000-s1-dk9.5.1.1a.gbin sup-1
boot kickstart bootflash:/n7000-s1-kickstart.5.1.1a.gbin sup-2
boot system bootflash:/n7000-s1-dk9.5.1.1a.gbin sup-2
router ospf 1
    auto-cost reference-bandwidth 1000000
router ospf 2
    router-id 172.26.255.153
ip pim rp-address 172.26.255.101 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8
*****
*****

```

### East-a (Multicast Mode)

```

version 5.1(1a)
hostname East-a
feature telnet
feature ospf
feature otv
feature lacp
feature dhcp
username admin password 5 $1$36p3G1AA$Pq09DfOCaBaSvfVj1U1ld. role vdc-admin
no password strength-check
ip domain-lookup
ip access-list ALL_IPs
    10 permit ip any any
mac access-list ALL_MACs
    10 permit any any
ip access-list HSRPv1_IP
    10 permit udp any 224.0.0.2/32 eq 1985
mac access-list HSRP_VMAC
    10 permit 0000.0c07.ac00 0000.0000.00ff any
arp access-list HSRP_VMAC_ARP
    10 deny ip any mac 0000.0c07.ac00 ffff.ffff.ff00
    20 permit ip any mac any

```



```

vlan access-map HSRPv1_Loc 10
  match mac address HSRP_VMAC
  match ip address HSRPv1_IP
  action drop
vlan access-map HSRPv1_Loc 20
  match mac address ALL_MACs
  match action forward ip address ALL_IPs
vlan filter HSRPv1_Loc vlan-list 100-199
ip arp inspection filter HSRP_VMAC_ARP <100-199>
snmp-server user admin vdc-admin auth md5 0x88018226be1701759b4301a3c0519193 pri
v 0x88018226be1701759b4301a3c0519193 localizedkey
vrf context management
  ip route 0.0.0.0/0 172.26.245.1
vlan 1,99-199
otv site-vlan 99
otv site-identifier 0x1
mac-list HSRP_VMAC_Deny seq 5 deny 0000.0c07.ac00 ffff.ffff.ff00
mac-list HSRP_VMAC_Deny seq 10 permit 0000.0000.0000 0000.0000.0000
route-map stop-HSRP permit 10
  match mac-list HSRP_VMAC_Deny
interface port-channel2
  description [ To N7K-a - Internal Interface ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
interface Overlay0
  otv join-interface Ethernet2/6
  otv control-group 239.1.1.1
  otv data-group 232.1.1.0/28
  otv extend-vlan 100-199
  no shutdown
interface Ethernet2/6
  description [ OTV Join-Interface ]
  ip address 172.26.255.94/30
  ip ospf network point-to-point
  ip router ospf 2 area 0.0.0.0
  ip igmp version 3
  no shutdown
interface Ethernet2/21
  description [ To N7K-a - Internal Interface ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-227
  channel-group 2 mode active
  no shutdown

```

```

interface Ethernet2/23
  description [ To N7K-b - Internal Interface ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-227
  channel-group 2 mode active
  no shutdown
interface mgmt0
  ip address 172.26.245.21/24
interface loopback0
  ip address 172.26.255.151/32
  ip router ospf 2 area 0.0.0.0
cli alias name his show cli hist unfo 20
line console
  exec-timeout 0
line vty
  exec-timeout 0
router ospf 2
  router-id 172.26.255.151
  timers throttle spf 10 100 500
otv-isis default
  vpn Overlay0
  redistribute filter route-map stop-HSRP

```

### East-a (Unicast-Only Mode)

The configuration is mostly identical to the one shown above. The only difference is in the Overlay interface configuration, as shown below.

```

interface Overlay0
  otv join-interface Ethernet2/6
  otv adjacency-server unicast-only
  otv use-adjacency-server 172.26.255.94 172.27.255.94 unicast-only
  otv extend-vlan 100-199
  no shutdown
*****
*****

```

### NR-n7k-b

```

version 5.1(1a)
license grace-period
hostname NR-n7k-b
!Following configuration set is not terminated by a newline
no vdc combined-hostname
vdc NR-n7k-b id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2

```

```

limit-resource monitor-session-erspan-dst minimum 0 maximum 23
limit-resource vrf minimum 16 maximum 1000
limit-resource port-channel minimum 0 maximum 768
limit-resource u4route-mem minimum 32 maximum 32
limit-resource u6route-mem minimum 16 maximum 16
limit-resource m4route-mem minimum 58 maximum 58
limit-resource m6route-mem minimum 8 maximum 8
vdc East-b id 2
allocate interface Ethernet1/26, Ethernet1/28, Ethernet1/30, Ethernet1/32
allocate interface Ethernet2/6, Ethernet2/21, Ethernet2/23
boot-order 3
limit-resource vlan minimum 16 maximum 4094
limit-resource monitor-session minimum 0 maximum 2
limit-resource monitor-session-erspan-dst minimum 0 maximum 23
limit-resource vrf minimum 16 maximum 1000
limit-resource port-channel minimum 0 maximum 768
limit-resource u4route-mem minimum 8 maximum 8
limit-resource u6route-mem minimum 4 maximum 4
limit-resource m4route-mem minimum 8 maximum 8
limit-resource m6route-mem minimum 2 maximum 2
feature telnet
cfs eth distribute
feature ospf
feature pim
feature interface-vlan
feature hsrp
feature lacp
feature vpc
logging level monitor 7
username adminbackup password 5 $1$Oip/C5Ci$oOdx7oJS1BCFpNRmQK4na. role
network-operator
username admin password 5 $1$T1wpkssO$4U6JRuGrh5M8WvbYXTsnV0 role network-admin
username adminbackup password 5 $1$Oip/C5Ci$oOdx7oJS1BCFpNRmQK4na. role
network-operator
no password strength-check
ip domain-lookup
snmp-server user admin network-admin auth md5
0x1ef34a157db87c5884230ac8e89f4663 priv 0x1ef34a157db87c5884230ac8e89f4663
localizedkey
ntp server 171.68.10.80 use-vrf management
ntp server 171.68.10.150 use-vrf management
ntp source-interface mgmt0
vrf context management
ip route 0.0.0.0/0 172.26.245.1
vlan 1-4, 99-199

```

```
spanning-tree vlan 99-199 priority 8192
vpc domain 1
  role priority 8192
  peer-keepalive destination 172.26.245.20 source 172.26.245.10
interface Vlan100
  no shutdown
  management
  no ip redirects
  ip address 10.100.1.5/24
  ip ospf network broadcast
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 100
    preempt delay minimum 60
    priority 20
    timers 1 3
    ip 10.100.1.1
<SNIP>
interface Vlan199
  no shutdown
  no ip redirects
  ip address 10.199.1.5/24
  ip ospf network broadcast
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 199
    preempt delay minimum 60
    priority 20
    timers 1 3
    ip 10.199.1.1
interface port-channel1
  description [ To N5K Access ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  spanning-tree port type network
  vpc 1
interface port-channel10
  description [ To N7K-a ]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 99-199
  spanning-tree port type network
  vpc peer-link
interface port-channel20
```

```
description [ To this OTV VDC ]
switchport
switchport mode trunk
switchport trunk allowed vlan 99-199
vpc 20
interface port-channel30
description [ To this OTV VDC ]
switchport
switchport mode trunk
switchport trunk allowed vlan 99-199
vpc 30
interface Ethernet1/1
description [ To N7K-a ]
switchport
switchport mode trunk
switchport trunk allowed vlan 99-199
channel-group 10 mode active
no shutdown
interface Ethernet1/9
description [ To N5K Access ]
switchport
switchport mode trunk
switchport trunk allowed vlan 99-199
channel-group 1 mode active
no shutdown
interface Ethernet2/1
description [ To N7K-a ]
switchport
switchport mode trunk
switchport trunk allowed vlan 99-199
channel-group 10 mode active
no shutdown
interface Ethernet2/5
description [ To the OTV Join-Interface ]
ip address 172.26.255.97/30
ip ospf network point-to-point
ip router ospf 2 area 0.0.0.0
ip pim sparse-mode
ip igmp version 3
no shutdown
interface Ethernet2/17
description [ To Core A ]
ip address 172.26.255.74/30
ip ospf network point-to-point
ip router ospf 2 area 0.0.0.0
```

```

ip pim sparse-mode
ip igmp version 3
no shutdown
interface Ethernet2/18
description [ To Core B ]
ip address 172.26.255.82/30
ip ospf network point-to-point
ip router ospf 2 area 0.0.0.0
ip pim sparse-mode
ip igmp version 3
no shutdown
interface Ethernet2/20
description [ To this OTV VDC ]
switchport
switchport mode trunk
switchport trunk allowed vlan 99-199
channel-group 30 mode active
no shutdown
interface Ethernet2/22
description [ To other OTV VDC ]
switchport
switchport mode trunk
switchport trunk allowed vlan 99-199
channel-group 20 mode active
no shutdown
interface mgmt0
ip address 172.26.245.10/24
interface loopback2
ip address 172.26.255.154/32
ip router ospf 2 area 0.0.0.0
cli alias name sw switchto vdc east-b
line console
exec-timeout 0
speed 115200
line vty
exec-timeout 0
boot kickstart bootflash:/n7000-s1-kickstart.5.1.1a.gbin sup-1
boot system bootflash:/n7000-s1-dk9.5.1.1a.gbin sup-1
boot kickstart bootflash:/n7000-s1-kickstart.5.1.1a.gbin sup-2
boot system bootflash:/n7000-s1-dk9.5.1.1a.gbin sup-2
router ospf 1
auto-cost reference-bandwidth 1000000
router ospf 2
router-id 172.26.255.154
timers throttle spf 10 100 5000

```

```

ip pim rp-address 172.26.255.101 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8
ip routing multicast holddown 0
logging monitor 7
logging console 7
*****
*****

```

### East-b (Multicast Mode)

```

version 5.1(1a)
hostname East-b
feature telnet
feature ospf
feature otv
feature lacp
feature dhcp
logging level otv 7
username admin password 5 $1$mDXdlrBj$3UtOG.HD2w.PI41n2apYe/ role vdc-admin
no password strength-check
ip domain-lookup
ip access-list ALL_IPs
  10 permit ip any any
mac access-list ALL_MACs
  10 permit any any
ip access-list HSRPv1_IP
  10 permit udp any 224.0.0.2/32 eq 1985
mac access-list HSRP_VMAC
  10 permit 0000.0c07.ac00 0000.0000.00ff any
arp access-list HSRP_VMAC_ARP
  10 deny ip any mac 0000.0c07.ac00 ffff.ffff.ff00
  20 permit ip any mac any
vlan access-map HSRPv1_Loc 10
  match mac address HSRP_VMAC
  match ip address HSRP_IP
  action drop
vlan access-map HSRPv1_Loc 20
  match mac address ALL_MACs
  match ip address ALL_IPs
  action forward
vlan filter HSRPv1_Loc vlan-list 100-199
ip arp inspection filter HSRP_VMAC_ARP <100-199>
snmp-server user admin vdc-admin auth md5 0x1ef34a157db87c5884230ac8e89f4663 pri
v 0x1ef34a157db87c5884230ac8e89f4663 localizedkey
vrf context management
ip route 0.0.0.0/0 172.26.245.1

```

```

vlan 1,99-199
otv site-vlan 99
otv site-identifier 0x1
mac-list HSRP_VMAC_Deny seq 5 deny 0000.0c07.ac00 ffff.ffff.ff00
mac-list HSRP_VMAC_Deny seq 10 permit 0000.0000.0000 0000.0000.0000
route-map stop-HSRP permit 10
    match mac-list HSRP_VMAC_Deny
interface port-channel3
    description [ OTV Internal Interface ]
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 99-199
interface Overlay0
    otv join-interface Ethernet2/6
    otv control-group 239.1.1.1
    otv data-group 232.1.1.0/28
    otv extend-vlan 100-199
    no shutdown
interface Ethernet2/6
    description [ OTV Join-Interface ]
    ip address 172.26.255.98/30
    ip ospf network point-to-point
    ip router ospf 2 area 0.0.0.0
    ip igmp version 3
    no shutdown
interface Ethernet2/21
    description [ To N7K-a - Internal Interface ]
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 99-199
    channel-group 3 mode active
    no shutdown
interface Ethernet2/23
    description [ To N7K-b - Internal Interface ]
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 99-199
    channel-group 3 mode active
    no shutdown
interface mgmt0
    ip address 172.26.245.11/24
interface loopback0
    ip address 172.26.255.152/32
    ip router ospf 2 area 0.0.0.0
logging monitor 7

```



```
logging console 7
line console
  exec-timeout 0
line vty
  exec-timeout 0
router ospf 2
  router-id 172.26.255.152
otv-isis default
  vpn Overlay0
    redistribute filter route-map stop-HSRP
```

### **East-b (Unicast-Only Mode)**

The configuration is mostly identical to the one shown previously. The only difference is in the overlay interface configuration, as follows:

```
interface Overlay0
  otv join-interface Ethernet2/6
  otv use-adjacency-server 172.26.255.94 172.27.255.94 unicast-only
  otv extend-vlan 100-199
  no shutdown
```

## **Configuration Example 2: FHRP Isolation**

Step 1. Configure a VLAN ACL (VACL) on the OTV VDC.

```
ip access-list ALL_IPs

 10 permit ip any any

!

mac access-list ALL_MACs

 10 permit any any

!

ipv6 access-list ALL_IPV6s

 10 permit ipv6 any any

!

ip access-list HSRP_IP

 10 permit udp any 224.0.0.2/32 eq 1985
```

```
20 permit udp any 224.0.0.102/32 eq 1985

!

ipv6 access-list HSRP_IPV6
  10 permit udp any ff02::66/128

!

mac access-list HSRP_VMAC

  10 permit 0000.0c07.ac00 0000.0000.00ff any

  20 permit 0000.0c9f.f000 0000.0000.0fff any

  30 permit 0005.73a0.0000 0000.0000.0fff any

!

arp access-list HSRP_VMAC_ARP

  10 deny ip any mac 0000.0c07.ac00 ffff.ffff.ff00

  20 deny ip any mac 0000.0c9f.f000 ffff.ffff.f000
  30 deny ip any mac 0005.73a0.0000 ffff.ffff.f000

  40 permit ip any mac any

vlan access-map HSRP_Localization 10

  match mac address HSRP_VMAC

  match ip address HSRP_IP
  match ipv6 address HSRP_IPV6

action drop

vlan access-map HSRP_Localization 20

  match mac address ALL_MACs

  match ip address ALL_IPs

  match ipv6 address ALL_IPV6s
```

```
    action forward

    !

    feature dhcp

    ip arp inspection filter HSRP_VMAC_ARP <OTV_Extended_VLANs>

    vlan filter HSRP_Localization vlan-list <OTV_Extended_VLANs>
```

**Step 2. Apply a route-map to the OTV control protocol (Intermediate System-to-Intermediate System [IS-IS]).**

```
mac-list OTV_HSRP_VMAC_deny seq 10 deny 0000.0c07.ac00 ffff.ffff.ff00

mac-list OTV_HSRP_VMAC_deny seq 11 deny 0000.0c9f.f000 ffff.ffff.f000
mac-list OTV_HSRP_VMAC_deny seq 12 deny 0005.73a0.0000 ffff.ffff.f000

mac-list OTV_HSRP_VMAC_deny seq 20 permit 0000.0000.0000 0000.0000.0000

    !

    route-map OTV_HSRP_filter permit 10

    match mac-list OTV_HSRP_VMAC_deny

    !

    otv-isis default

    vpn Overlay0

    redistribute filter route-map OTV_HSRP_filter
```

### Configuration Example 3: Direct VLAN Translation

**Step 1. Enter configuration mode for the overlay.**

```
conf
interface overlay1
```

**Step 2. Apply the **vlan mapping** command with the local VLAN number listed first.**

```
otv vlan mapping 100 to 200
```

---

## Configuration Example 4: VLAN Translation Through a Transit VLAN

Step 3. Enter configuration mode for the overlay in Data Center 1.

```
conf
interface overlay1
```

Step 4. Apply the **vlan mapping** command with the local VLAN number listed first for Data Center 1.

```
otv vlan mapping 100 to 400
```

Step 5. Enter configuration mode for the overlay in Data Center 2.

```
conf
interface overlay1
```

Step 6. Apply the **vlan mapping** command with the local VLAN number listed first for Data Center 2.

```
otv vlan mapping 200 to 400
```

Step 7. Enter configuration mode for the overlay in Data Center 3.

```
conf
interface overlay1
```

Step 8. Apply the **vlan mapping** command with the local VLAN number listed first for Data Center 3.

```
otv vlan mapping 300 to 400
```




---

**Americas Headquarters**  
Cisco Systems, Inc.  
San Jose, CA

**Asia Pacific Headquarters**  
Cisco Systems (USA) Pte. Ltd.  
Singapore

**Europe Headquarters**  
Cisco Systems International BV Amsterdam,  
The Netherlands

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

 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/go/trademarks>. 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. (1110R)