



Overlay Transport Virtualization

Using OTV to Extend Layer 2 between Two
Data Centers Connected Through
Dark Fiber Links

Introduction	3
Design Overview	3
Aggregation Tier and DCI Overview	4
Aggregation Tier: Layer 3 Overview	5
Aggregation Tier: Layer 2 Overview	5
Overlay Transport Virtualization (OTV) Overview	5
Traffic Flow Overview	6
Layer 3 Routing Design and Configuration	8
EIGRP Routing	8
Layer 2 Design and Configuration	10
vPC Design and Configuration	10
Server Connectivity	11
Spanning Tree Protocol	13
First Hop Redundancy Protocol (FHRP)	13
Overlay Transport Virtualization (OTV)	14
OTV Control Plane: Steady State Information	19
HW-St: Up Peer-ID: 3 If-Name: Tunnel16515	21
Multipathing and AED	21
Appendix	23

Introduction

This document provides guidance on configuring and designing a network with overlay transport virtualization (OTV) in the aggregation tier. A primary goal is to provide customers and field services with a deployment guide. We assume that the reader already has a good knowledge of the OTV functionality.

In this document we discuss a data center interconnect (DCI) design which uses OTV technology to interconnect two data centers. Our goal is to highlight OTV deployment, not to advise on the best practices or technology options for DCI designs.

The benefits of using OTV for Layer 2 extension include:

- No need for Ethernet over Multiprotocol Label Switching (EoMPLS) or Virtual Private LAN Services (VPLS) deployment for layer 2 extensions
- Provision of Layer 2 and Layer 3 connectivity using the same dark fiber connections
- Native Spanning Tree Protocol (STP) isolation: no need to explicitly configure bridge data protocol unit (BPDU) filtering
- Native unknown unicast flooding isolation: unknown unicast not sent to the overlay
- Address Resolution Protocol (ARP) optimization with the OTV ARP cache
- Simplified provisioning of First Hop Redundancy Protocol (FHRP) isolation
- Simplified addition of sites

The focus of this document is primarily on the aggregation tier in which the virtual device contexts (VDCs) such as the aggregation VDC and the OTV VDC reside. The aggregation tier design discussed here can be part of a two-tier (collapsed aggregation core and access) or three-tier design (core, aggregation, access) of a data center infrastructure.

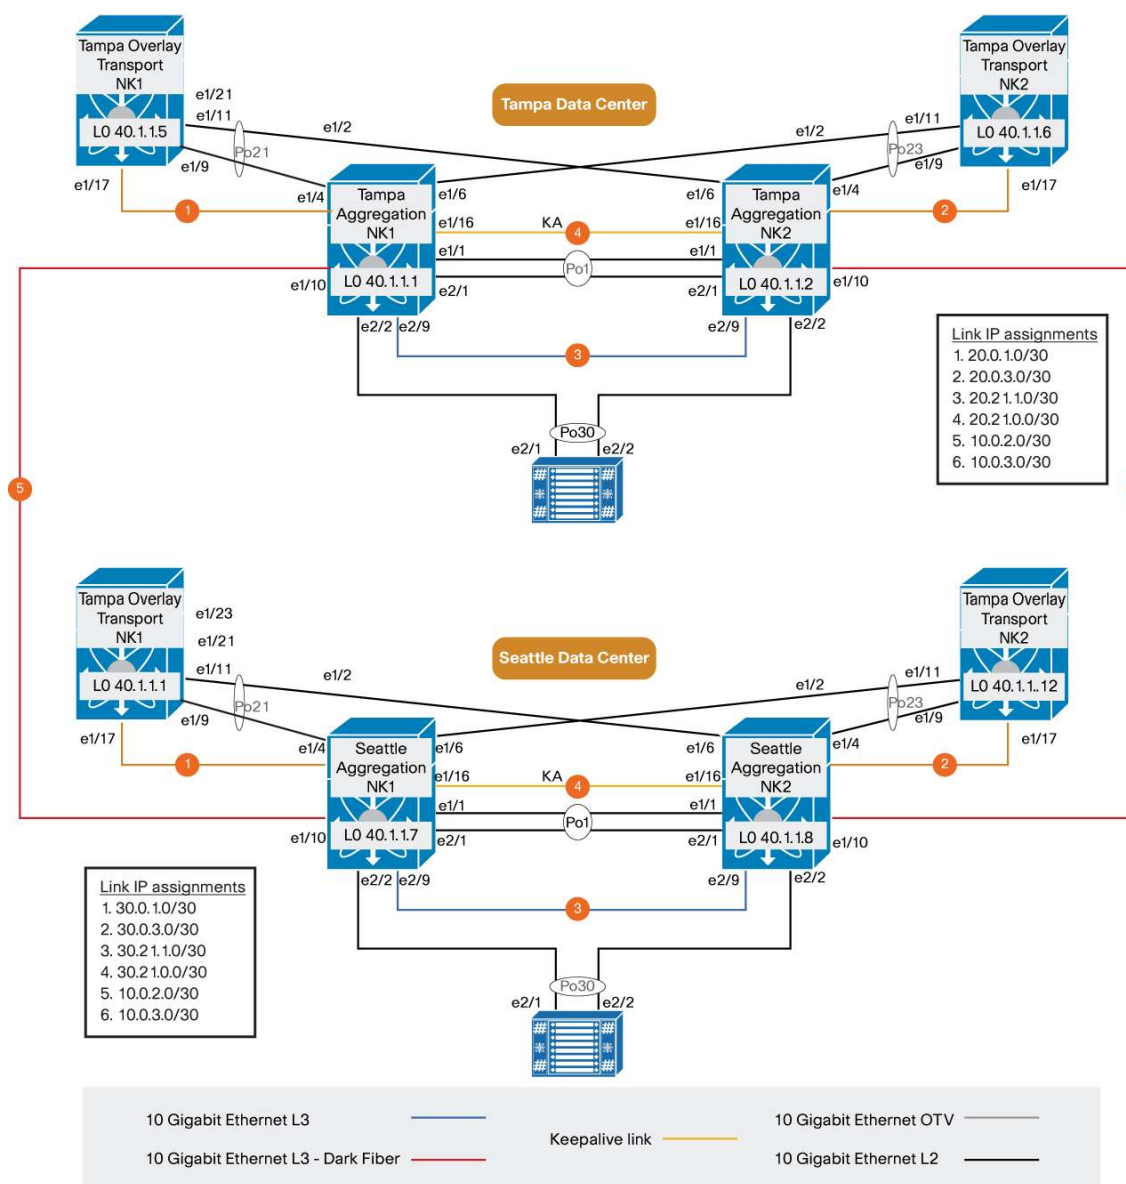
Design Overview

The logical design consists of three network tiers in the data center: core, aggregation and access. The intent of the design is to demonstrate Layer 2 extension between data centers over a Layer 2 or Layer 3 data center interconnect (DCI). The aggregation tier contains the functional elements to provide DCI.

The aggregation tier in each data center consists of two Cisco Nexus® 7000 Series Switches deployed in a redundant manner. The aggregation tier is configured to support multiple VDCs. Each of the Cisco Nexus 7000 Series Switches has three VDCs configured: the default VDC, the aggregation VDC, and the OTV VDC. The aggregation VDCs, as Figure 1 shows, act as a Layer 3 DCI border router in both the data centers. The OTV is configured in an “appliance” mode in its own VDC. The function of the OTV VDC is to serve as a Layer 2 switch over which the other VDCs or physical Layer 2 switches can extend their Layer 2 domains across the DCI to other data centers.

In this design, the aggregation VDC connects through Virtual PortChannel (vPC) technology to the OTV VDC. It should be noted that in a multitenancy environment, the same OTV VDC can be configured with multiple overlays to provide a segmented Layer 2 extension for different tenants or applications. It should also be noted that when multiple data center sites are interconnected, the OTV operations can benefit from the presence of multicast in the core.

The overall data center interconnect OTV design, based on the Tampa and Seattle data centers, is shown in the Figure 1.

Figure 1. Data Center Interconnect OTV design

Aggregation Tier and DCI Overview

The aggregation tiers of the Tampa and Seattle data centers are identical and therefore we describe only one of the data centers: Tampa. The aggregation layer consists of two Cisco Nexus 7000 Series Switches, N7K1 and N7K2, which provide high availability. Each of the Cisco Nexus 7000 Series Switches is configured to have three VDCs: the default VDC, the aggregation VDC, and the OTV VDC. The role of the default VDC is to provide resource management to the aggregation and to the OTV VDCs. The OTV VDC residing on the aggregation switches are deployed as an “OTV appliance on a stick” to provide Layer 2 extension between the data centers. This means that the OTV VDC is acting as the edge device for the Layer 2 extension, and it is in the direct path of the Layer 2 data flow between the two data centers. It should be further noted that the OTV VDC is one hop away from the Layer 3 border routers providing DCI connectivity.

The discussion of the best practices for placement options of the OTV VDC is outside the scope of this document and is not covered.

As noted before, the aggregation layer

The aggregation tier AGG VDCs provi

The design uses QTV in a virtual appliance mode, and

Author	Year	Country	Sample Size	Study Design	Findings
Smith et al.	2015	USA	1,200	Longitudinal	Increased risk of depression in children of parents with mental illness.
Johnson et al.	2016	UK	800	Cross-sectional	Higher levels of anxiety in children of parents with anxiety disorders.
Lee et al.	2017	Canada	950	Case-control	Association between parental mental illness and child behavioral problems.
Kim et al.	2018	South Korea	1,500	Longitudinal	Impact of parental mental health on child academic performance.
White et al.	2019	Australia	1,100	Cohort	Increased risk of substance use in children of parents with mental illness.
Green et al.	2020	Germany	1,300	Longitudinal	Association between parental mental health and child emotional well-being.
Brown et al.	2021	France	1,400	Cross-sectional	Impact of parental mental health on child social skills.
Miller et al.	2022	Italy	1,600	Longitudinal	Association between parental mental health and child cognitive development.
Wilson et al.	2023	Spain	1,700	Cohort	Impact of parental mental health on child physical health.
Anderson et al.	2024	Sweden	1,800	Longitudinal	Association between parental mental health and child self-esteem.
Clark et al.	2025	Japan	1,900	Cross-sectional	Impact of parental mental health on child life satisfaction.

achieve high availability, the OTV VDCs are configured on both N7K1 and N7K2 to provide multihoming capability. The OTV VDCs in T-N7k1 and T-N7K2 are connected through vPC to the AGG VDCs, as shown in the Figure 3. The use of vPC allows Layer 2 path redundancy from the OTV VDCs to the AGG VDCs.

As shown Figure 3, OTV has three types of interfaces: the join, the overlay, and the internal interface.

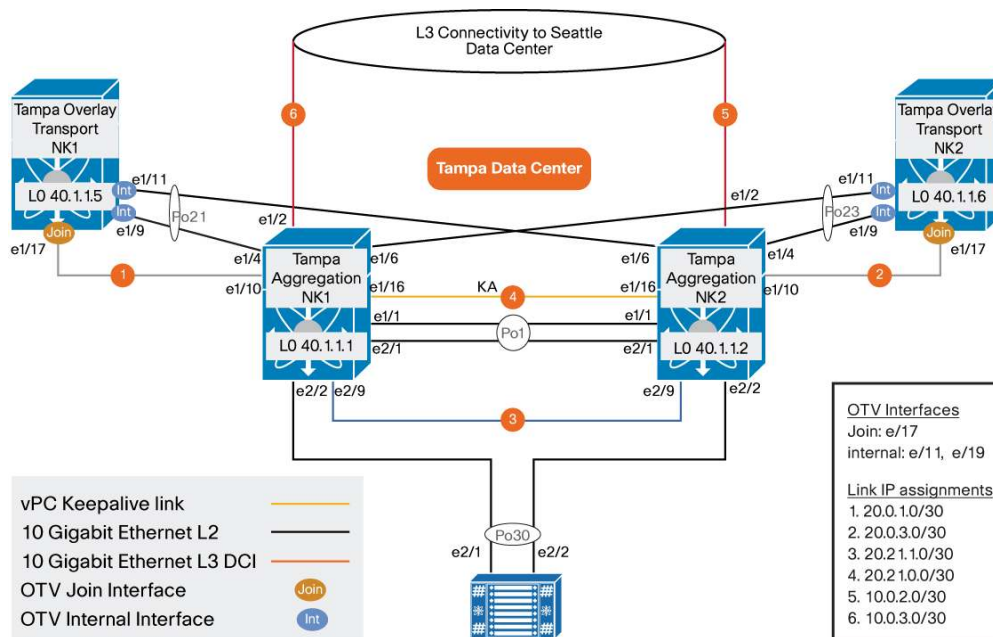
The OTV internal interfaces carry the VLANs to be extended and the OTV site VLAN (used within the data center to provide multihoming). They behave as regular Layer 2 switch port trunk interfaces; in fact, they send, receive, and process the Spanning Tree Protocol BPDUs as they would on a regular LAN bridge device.

The overlay interfaces encapsulate Layer 2 frames in IP unicast or multicast packets and are logical multiaccess, multicast-capable interfaces.

The join interfaces are point-to-point routed interfaces that are used by the OTV edge devices to join the overlay network. Currently, the OTV implementation uses the IP address of the physical join interface to advertise reachability of MAC-addresses present in the site. In a future release, the plan is to use the loopback address for this purpose.

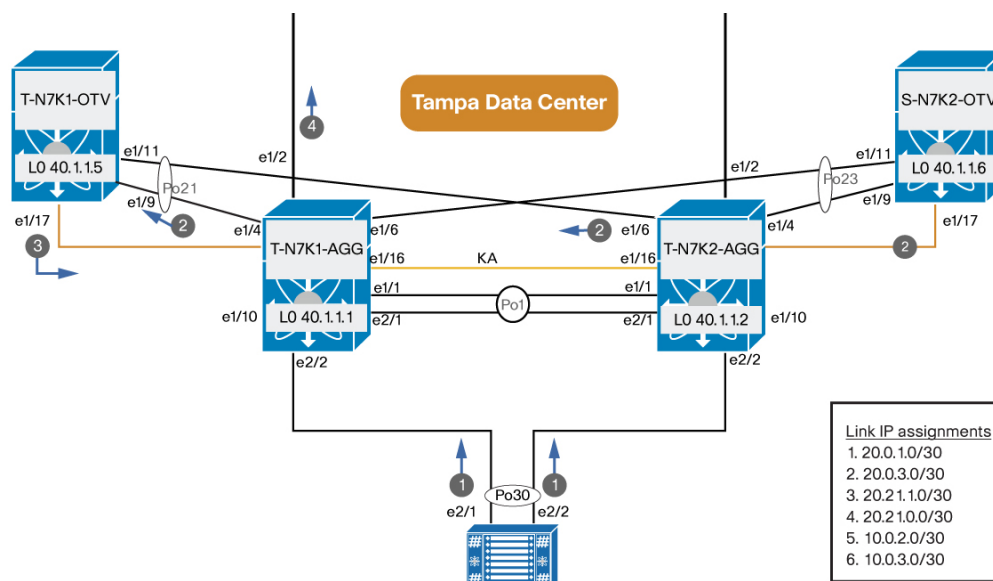
The same OTV VDCs can be used by multiple VDCs deployed at the aggregation tier, as well as by other Layer 2 switches connected to the OTV VDCs. This is done by configuring multiple OTV overlays. It's important to note that the extended VLANs within these multiple overlays should not overlap.

Figure 3. Overlay Transport Virtualization Overview

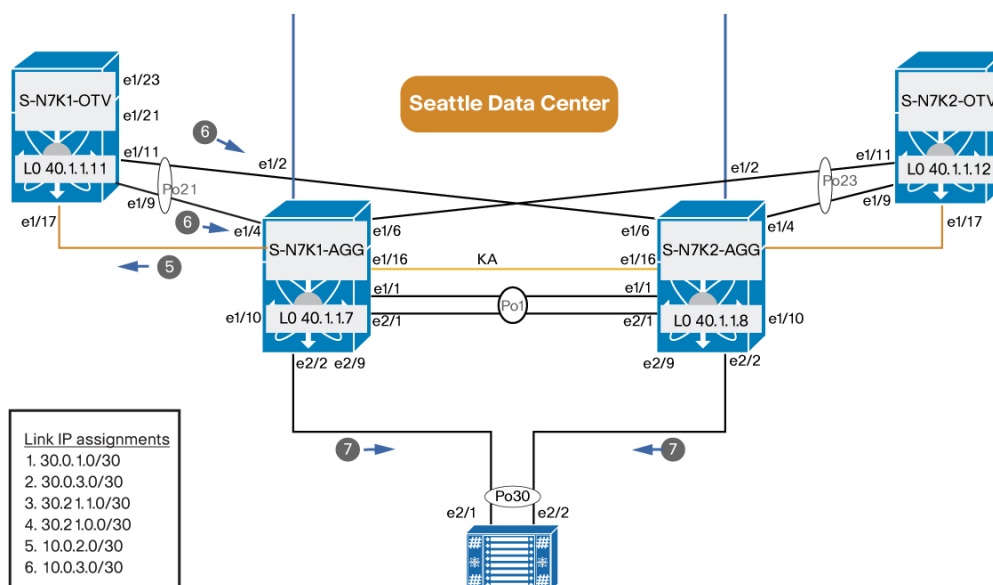


Traffic Flow Overview

In this section, we examine the flow of the traffic between two data centers to gain an understanding of the design functionality. In the steps that follow, we examine the path of the Layer 2 unicast traffic flow between two servers, one located in Tampa called **T-srv1** and the other located in Seattle called **S-srv1**. The server T-srv1 and S-srv1 are connected via vPC to their respective aggregation tier switches on VLAN 10. For this flow, assume that both servers reside in VLAN 10 and T-N7k1-OTV and S-N7k2-OTV are the authoritative edge devices (AEDs) for VLAN 10 in Tampa and Seattle. Figures 4 and 5 illustrate the steps in the traffic flow.

Figure 4. Tampa Data Center

1. The Layer 2 packet originating from T-srv1 travels over the vPC link Po30 and arrives at either T-N7K1-AGG or T-N7K2-AGG. The hashing algorithm determines which of the two links the packet will take.
2. From either the T-N7K1-AGG or T-N7K2-AGG the packet travels over the vPC link to the OTV AED for VLAN 10. Since AED for VLAN 10 in Tampa is T-N7K1-OTV, the packet would travel over vPC link Po21 to the OTV VDC T-N7K1-OTV.
3. The OTV VDC does a Layer 2 lookup on the destination MAC address of the frame. The MAC table of OTV VDC points to Seattle Data Center for this destination MAC. At this point, the OTV VDC performs a MAC-in-IP encapsulation where the destination IP address is the IP address of the OTV join interface of the Seattle's OTV VDC, which is the AED for VLAN 10. Based on its routing table, the OTV VDC forwards the packet over the OTV join interface e1/17 (which is a Layer 3 link) peering to interface e1/10 on T-N7K1-AGG.
4. The Tampa AGG VDC, T-N7K1-AGG, on receiving the encapsulated packet, does an IP destination lookup and forwards it via EIGRP over the e2/9 interface to the AGG VDC in Seattle S-N7K2-AGG.

Figure 5. Seattle Data Center

5. The Seattle AGG VDC S-N7K1-AGG, on receiving the packet on e1/2, forwards it over the interface e1/10 to the OTV join interface e/17 on S-N7k1-OTV.
6. Once in the OTV VDC, the IP packet gets deencapsulated to a Layer 2 packet. The packet is then forwarded toward its Layer 2 destination based on the MAC address table. In this case, the packet is sent over the internal interface, vPC port channel Po23, and could take either link in the port-channel
7. The Layer 2 packet arrives either at S-N7k1-AGG or S-N7K2-AGG VDCs, depending on the vPC hash. Either of the VDCs then directly forwards the packet down over Po30 to the destination server S-Srv1.

Layer 3 Routing Design and Configuration

EIGRP Routing

In a well-designed network, EIGRP scales well and provides extremely quick convergence times. For this testing, EIGRP is used as the routing protocol of choice to provide Layer 3 connectivity and fast convergence both within and between the two data centers (Tampa and Seattle).

In both data centers, all Cisco Nexus 7000 Series Switches are configured such that EIGRP process 100 is running on the interswitch link (Eth 2/9), interdata center link (Eth 1/2), OTV join interface (Eth 1/17), and VLAN 10 interface

Bidirectional Forwarding Detection (BFD) protocol is also used to provide fast forwarding path failure detection times for EIGRP.

Tampa Data Center

T_N7K1-AGG#	T_N7K2-AGG#
<pre> interface Vlan10 description UCS ip address 172.16.10.2/24 ip router eigrp 100 ip passive-interface eigrp 100 ! interface Ethernet1/2 description S_N7K1-AGG e1/2 L3 ip address 10.0.2.1/30 ip router eigrp 100 ! interface Ethernet1/10 description T_N7K1-OTV e1/17 L3 ip address 20.0.1.2/30 ip router eigrp 100 ! interface Ethernet2/9 description T_N7K2-AGG e2/9 L3 ip address 20.21.1.1/30 ip router eigrp 100 ! interface loopback0 ip address 40.1.1.1/32 ip router eigrp 100 ! router eigrp 100 autonomous-system 100 </pre>	<pre> interface Vlan10 description UCS ip address 172.16.10.3/24 ip router eigrp 100 ip passive-interface eigrp 100 ! interface Ethernet1/2 description S_N7K2-AGG e1/2 L3 ip address 10.0.3.1/30 ip router eigrp 100 ! interface Ethernet1/10 description T_N7K2-OTV e1/17 L3 ip address 20.0.3.2/30 ip router eigrp 100 ! interface Ethernet2/9 description T_N7K1-AGG e2/9 L3 ip address 20.21.1.2/30 ip router eigrp 100 ! interface loopback0 ip address 40.1.1.1/32 ip router eigrp 100 ! router eigrp 100 autonomous-system 100 </pre>

<pre> bfd T_N7K1-OTV# interface Ethernet1/17 description T_N7K1-AGG e1/10 L3 ip address 20.0.1.1/30 ip router eigrp 100 ! router eigrp 100 bfd </pre>	<pre> bfd T_N7K2-OTV# interface Ethernet1/17 description T_N7K2-AGG e1/10 L3 ip address 20.0.2.1/30 ip router eigrp 100 ! router eigrp 100 bfd </pre>
---	---

Seattle Data Center

<pre> S_N7K1-AGG# interface Vlan10 description UAGG ip address 172.16.10.4/24 ip router eigrp 100 ip passive-interface eigrp 100 ! interface Ethernet1/2 description T_N7K1-AGG e1/2 L3 ip address 10.0.2.2/30 ip router eigrp 100 ! interface Ethernet1/10 description S_N7K1-OTV e1/17 L3 ip address 30.0.1.2/30 ip router eigrp 100 ! interface Ethernet2/9 description S_N7K2-AGG e2/9 L3 ip address 30.21.1.1/30 ip router eigrp 100 ! interface loopback0 ip address 40.1.1.7/32 ip router eigrp 100 ! router eigrp 100 autonomous-system 100 bfd T_N7K1-OTV# interface Ethernet1/17 description S_N7K1-AGG e1/10 L3 ip address 30.0.1.1/30 ip router eigrp 100 ! router eigrp 100 </pre>	<pre> S_N7K2-AGG# interface Vlan10 description UCS ip address 172.16.10.5/24 ip router eigrp 100 ip passive-interface eigrp 100 ! interface Ethernet1/2 description T_N7K2-AGG e1/2 L3 ip address 10.0.3.2/30 ip router eigrp 100 ! interface Ethernet1/10 description S_N7K2-OTV e1/17 L3 ip address 30.0.3.2/30 ip router eigrp 100 ! interface Ethernet2/9 description S_N7K1-AGG e2/9 L3 ip address 30.21.1.2/30 ip router eigrp 100 ! interface loopback0 ip address 40.1.1.8/32 ip router eigrp 100 ! router eigrp 100 autonomous-system 100 bfd T_N7K2-OTV# interface Ethernet1/17 description S_N7K2-AGG e1/10 L3 ip address 30.0.3.1/30 ip router eigrp 100 ! router eigrp 100 </pre>
--	---

bfd	bfd
-----	-----

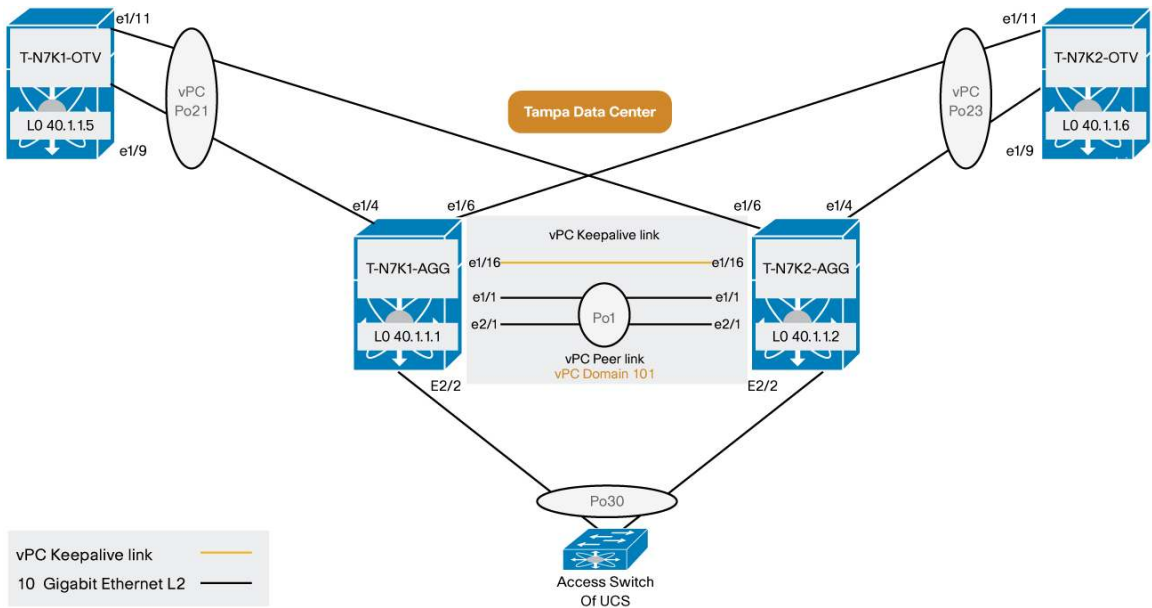
Layer 2 Design and Configuration

In this section, we discuss the Layer 2 connectivity between the OTV VDCs and the AGG VDC on the aggregation tier switches. We also discuss the server connectivity to the aggregation tier.

vPC Design and Configuration

The vPC configuration on the aggregation tier is a typical vPC design. The vPC peer-link is configured as a port channel between the two Cisco Nexus 7000 Series Switches in AGG VDC that act as a single virtual switch with separate control planes. As shown in the Figure 6, Po1 is configured as a vPC peer-link between the two switches, providing vPC services to access layer devices and OTV VDCs (that is, T-N7K1-OTV and T-N7K2-OTV). The Po1 carries Layer 2 vPC traffic between the two switch VDCs. A vPC keepalive link (e1/16 to e1/16) is also configured between T-N7K1-AGG and T-N7K2-AGG for dual active detection. Note that a separate Layer 3 link (refer to Figure 1) between T-N7K1-AGG (e2/9) and T-N7K2-AGG (e2/9) is configured as per best practices to carry any Layer 3 traffic between the two aggregation VDCs.

Figure 6. Virtual Port Channel Connectivity



The configuration for vPC in T-N7K1-AGG and T-N7K2-AGG is as shown in the following configuration snippet:

Tampa Data Center

Peer-link Configuration	
T_N7K1-AGG#	T_N7K2-AGG#
vpc domain 101	vpc domain 101
interface port-channell	interface port-channell
description peer link to T-N7K2-AGG2	description peer link T-N7K2-AGG1
switchport	switchport
switchport mode trunk	switchport mode trunk
vpc peer-link	vpc peer-link
spanning-tree port type network	spanning-tree port type network

<pre> interface Ethernet1/1 description VPC peer e1/1 e2/1 T-AGG2 switchport switchport mode trunk channel-group 1 mode active interface Ethernet2/1 description VPC peer e1/1 e2/1 T-AGG2 switchport switchport mode trunk channel-group 1 mode active </pre>	<pre> interface Ethernet1/1 description VPC peer e1/1 e2/2 T-AGG1 switchport switchport mode trunk channel-group 1 mode active interface Ethernet2/1 description vPC peer e1/1 e2/2 T-AGG1 switchport switchport mode trunk channel-group 1 mode active </pre>
Keep-alive link configuration	
<pre> feature vpc vrf context keepalive vpc domain 101 peer-keepalive destination 20.21.0.2 source 20.21.0.1 vrf keepalive interface Ethernet1/16 description vpc peer vrf member keepalive ip address 20.21.0.1/30 </pre>	<pre> feature vpc vpc context keepalive vpc domain 101 peer-keepalive destination 20.21.0.1 source 20.21.0.2 vrf keepalive interface Ethernet1/16 description vpc peer vrf member keepalive ip address 20.21.0.2/30 </pre>

Server Connectivity

At the aggregation layer, vPC technology is used to connect the access layer server device (the Cisco UCS 6120XP 20-Port Fabric Interconnect) to two upstream switches, providing redundancy, higher bandwidth, and active-active bidirectional connectivity. The server sees vPC as a regular port-channel and does not know that it is connected to two separate upstream switches. As noted earlier, Po30 is used to connect to the server through vPC.

Tampa Data Center

T_N7K1-AGG#	T_N7K2-AGG#
<pre> interface port-channel30 description P030 to UCS switchport switchport mode trunk vpc 30 switchport trunk allowed vlan 10 interface Ethernet2/2 description port-channel to UCS switchport switchport mode trunk switchport trunk allowed vlan 10 channel-group 30 mode active </pre>	<pre> interface port-channel30 description P030 to UCS switchport switchport mode trunk vpc 30 switchport trunk allowed vlan 10 interface Ethernet2/2 description port-channel to UCS switchport switchport mode trunk switchport trunk allowed vlan 10 channel-group 30 mode active </pre>

Since traffic from the attached server can get hashed to any uplink in the bundle, vPC is also used between aggregation and OTV edge device as a means to provide direct path for Layer 2 server traffic sent to and received from the overlay.

<p>T_N7K1-AGG#</p> <pre> interface port-channel21 description VPC to T_N7K1-OTV switchport switchport mode trunk vpc 21 interface Ethernet1/4 description PO21 to N7K1-OTV e1/9 switchport switchport mode trunk channel-group 21 mode active </pre>	<p>T_N7K2-AGG#</p> <pre> interface port-channel21 description VPC to T_N7K1-OTV switchport switchport mode trunk vpc 21 interface Ethernet1/6 description PO21 to N7K1-OTV e1/11 switchport switchport mode trunk channel-group 21 mode active </pre>
<p>T_N7K1-AGG#</p> <pre> interface port-channel23 description VPC to T_N7K2-OTV switchport switchport mode trunk vpc 23 interface Ethernet1/6 description PO23 to N7K1-OTV e1/9 switchport switchport mode trunk channel-group 23 mode active </pre>	<p>T_N7K2-AGG#</p> <pre> interface port-channel23 description VPC to T_N7K2-OTV switchport switchport mode trunk vpc 23 interface Ethernet1/4 description PO23 to N7K1-OTV e1/11 switchport switchport mode trunk channel-group 23 mode active </pre>
<p>T_N7K1-OTV#</p> <pre> interface port-channel21 description T-N7K-AGG switchport switchport mode trunk interface Ethernet1/9 description PO21 to N7K1-AGG e1/4 switchport switchport mode trunk channel-group 21 mode active interface Ethernet1/11 description PO21 to N7K2-AGG e1/6 switchport </pre>	<p>T_N7K2-OTV#</p> <pre> interface port-channel23 description T-N7K-AGG switchport switchport mode trunk interface Ethernet1/9 description PO21 to N7K2-AGG e1/4 switchport switchport mode trunk channel-group 23 mode active interface Ethernet1/11 description PO21 to N7K-AGG e1/6 switchport </pre>

```
switchport mode trunk
channel-group 21 mode active
```

```
switchport mode trunk
channel-group 23 mode active
```

Spanning Tree Protocol

The Spanning Tree Protocol used is Rapid-PVST+. The design does not have any blocked ports because vPC is used when dual-homing the server and the OTV edge device to the two aggregation VDCs. T_N7K1-AGG in Tampa and S_N7K1-AGG in Seattle are configured to be the root for all VLANs in their local sites.

First Hop Redundancy Protocol (FHRP)

The First Hop Redundancy Protocol used in this test topology is Hot Standby Routing Protocol (HSRP). HSRP provides default gateway redundancy for connected hosts such as servers.

Each VLAN in the aggregation block, which requires Layer 3 connectivity to rest of the network, is configured with an HSRP gateway, including ones extended on the overlay between the data centers.

Tampa Data Center

T_N7K1-AGG#	T_N7K2-AGG#
<pre>interface Vlan10 description UCS ip address 172.16.10.2/24 hsrp 10 preempt ip 172.16.10.1</pre>	<pre>interface Vlan10 description UCS ip address 172.16.10.3/24 hsrp 10 ip 172.16.10.1</pre>

To allow the extended VLANs to use their local HSRP gateway, an IP gateway localization technique is used to keep HSRP protocol data units (PDUs) from getting forwarded on the overlay network. This technique uses a combination of VLAN access control lists (VACLs) and OTV MAC route filters in the OTV VDC to block the propagation of HSRP packets between the Tampa and Seattle data centers. Alternatively, **port access control lists** PACL-based filtering of HSRP packets on the inside interface can also be used for similar results; however, this is outside the scope of this document.

Tampa Data Center

T_N7K1-OTV#
<pre>**VACL Filter** ip access-list ALL_IPs 10 permit ip 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 vlan access-map HSRP_Localization 10 match ip address HSRP_IP action drop vlan access-map HSRP_Localization 20 match ip address ALL_IPs action forward</pre>

```
vlan filter HSRP_Localization vlan-list 10

**OTV MAC route filter**

mac-list OTV_HSRP_VMAC_deny seq 10 deny 0000.0c07.ac00 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 20 deny 0000.0c9f.f000 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 30 permit 0000.0000.0000 0000.0000.0000

route-map OTV_HSRP_filter permit 10
  match mac-list OTV_HSRP_VMAC_deny

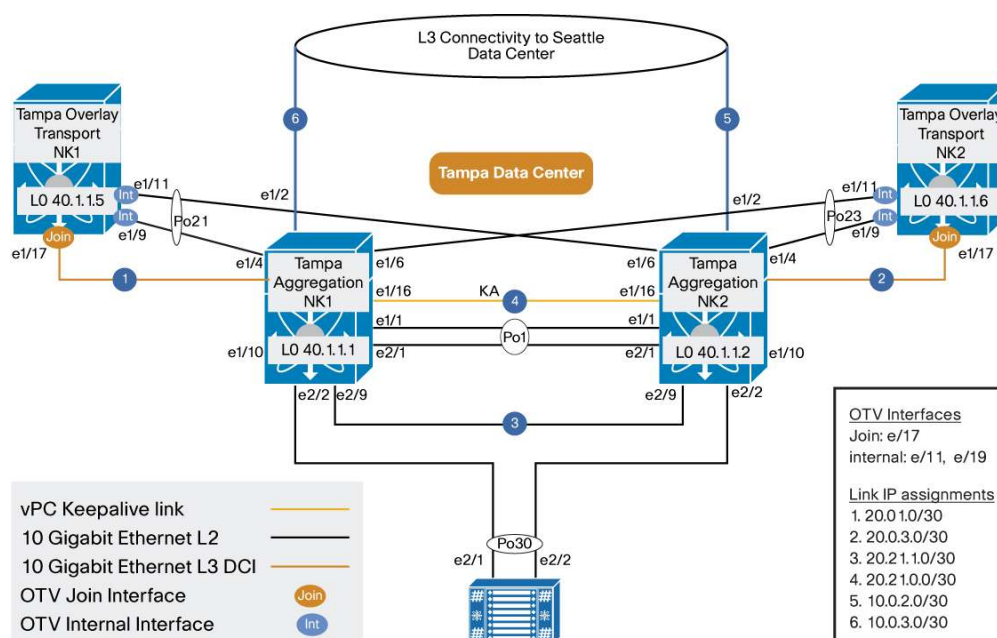
otv-isis default
  vpn Overlay0
    redistribute filter route-map OTV_HSRP_filter
```

Overlay Transport Virtualization (OTV)

OTV connectivity, as illustrated in Figure 7, is a "MAC in IP" technique for supporting Layer 2 VPNs and extending LANs over any transport, as long as this transport can carry IP packets, which is always the case.

OTV can be thought of as MAC-address routing in which destinations are MAC addresses and next hops are IP addresses. OTV simply maps MAC address destinations to IP next hops that are reachable through the transport cloud. The traffic destined for a remote MAC address is encapsulated in IP and carried through the IP cloud as regular IP traffic to their next hop address. Since traffic is IP forwarded, OTV is as efficient as the transport network and delivers optimal traffic load balancing, multicast traffic replication, and fast failover just as the transport is able to do.

The core principles on which OTV operates are the use of a control protocol to advertise MAC address reachability information (instead of using data plane learning, which is based on flooding) and packet switching of IP encapsulated Layer 2 traffic (instead of using circuit switching) for data forwarding.

Figure 7. Overlay Transport Virtualization Connectivity

The design used to perform this testing uses the OTV edge device as a virtual appliance on a stick, requiring OTV to be configured under a separate VDC. This is because the current implementation of OTV requires separation between the Layer 3 routing performed for the VLANs by the SVIs at the aggregation VDC and the Layer 2 extension, which is now performed by the OTV VDC.

The OTV VDC connects back to the aggregation layer switches through Layer 2 as well as Layer 3 links as shown in Figure 7. The Layer 2 links are also known as internal interfaces and are used by the OTV edge device to learn the MAC addresses of the site and forward Layer 2 traffic across the sites for the extended VLANs.

The two OTV internal interfaces are configured as a virtual port-channel trunk (Po21 and Po23) with one link going to each aggregation VDC.

Tampa Data Center

T_N7K1-OTV#	T_N7K2-OTV#
<pre> interface port-channel21 description T-N7K-AGG switchport switchport mode trunk switchport trunk allowed vlan 10,987 ! interface Ethernet1/9 switchport switchport mode trunk switchport trunk allowed vlan 10,987 channel-group 21 mode active ! interface Ethernet1/11 switchport switchport mode trunk </pre>	<pre> interface port-channel23 description T-N7K-AGG switchport switchport mode trunk switchport trunk allowed vlan 10,987 ! interface Ethernet1/9 switchport switchport mode trunk switchport trunk allowed vlan 10,987 channel-group 23 mode active ! interface Ethernet1/11 switchport switchport mode trunk </pre>

```
switchport trunk allowed vlan 10,987
channel-group 21 mode active
```

```
switchport trunk allowed vlan 10,987
channel-group 23 mode active
```

The Layer 3 link is known as the join interface, which OTV uses to perform IP-based virtualization to send and receive overlay traffic between sites. The IP address of this interface is used to advertise reachability of a MAC addresses present in the site.

There is one Join interface per OTV Overlay; however, if multiple Layer 3 interfaces are present on the OTV edge device, the unicast extended traffic can get routed over any of these links. Link bundling can be used to present a single aggregated Layer 3 link to OTV, providing link redundancy and resiliency. However, for this testing a single interface is used.

Tampa Data Center

T_N7K1-AGG#

```
interface Ethernet1/17
description T_N7K1-AGG e1/10 L3
mtu 9216
ip address 20.0.1.1/30
ip router eigrp 100
```

T_N7K2-AGG#

```
interface Ethernet1/17
description T_N7K2-AGG e1/10 L3
mtu 9216
ip address 20.0.3.1/30
ip router eigrp 100
```

OTV encapsulates packets into an IP header and where it sets the Don't Fragment (DF) bit for all OTV control and data packets crossing the transport network. The encapsulation adds 42 bytes to the original IP maximum transition unit (MTU) size. So it is a best practice to configure the join interface and all Layer 3 interfaces that face the IP core between the OTV edge devices with the max possible MTU size supported by the transport.

OTV uses site VLAN to allow multiple OTV edge devices within the site to talk to each other and determine the AED for the OTV-extended VLANs. The site VLAN used in this design is 987. It is a best practice to use a dedicated VLAN as site VLAN. The site VLAN should not be extended and should be carried down to the aggregation layer across the VPC peer link. Any change to the site VLAN configuration **must** happen with the overlay in shutdown mode. For more information about AED functionality, please see the discussion of multipath and AED later in this document.

T_N7K1-OTV#

```
feature otv
!
otv site-vlan 987
```

T_N7K2-OTV#

```
feature otv
!
otv site-vlan 987
```

The OTV edge device is also configured with the overlay interface, which is associated with the join interface to provide connectivity to the physical transport network. The overlay interface is used by OTV to send and receive Layer 2 frames encapsulated in IP packets. From the perspective of MAC-based forwarding on the site, the overlay interface is simply another bridged interface. However, no Spanning Tree Protocol packets or unknown unicast packets are forwarded over the overlay interface. From the perspective of the IP transport, the overlay interface is not visible.

T_N7K1-OTV#

```
interface Overlay1
otv join-interface Ethernet1/17
```

T_N7K2-OTV#

```
interface Overlay1
otv join-interface Ethernet1/17
```


description Overlay for AGG	description Overlay for AGG
-----------------------------	-----------------------------

Note: The overlay interface does not come up until you configure a multicast group address **and** the site-VLAN has at least an active port on the device.

By default, a VLAN is not advertised on the overlay network; therefore, forwarding cannot occur over the overlay network unless the VLANs are explicitly extended. Once the VLAN is extended, the OTV edge device will begin advertising locally learned MAC addresses on the overlay network.

T_N7K1-OTV# <pre>interface Overlay1 otv extend-vlan 10</pre>	T_N7K2-OTV# <pre>interface Overlay1 otv extend-vlan 10</pre>
--	--

Note: Currently, in order to modify the extended VLANs, the overlay interface needs to be in shutdown mode.

OTV builds control plane relationships to enable edge devices to discover each other, form adjacencies and exchange MAC reachability across the overlay network. These control plane relationships can be built using either a multicast- or unicast-enabled transport.

The unicast-enabled transport requires OTV to support the adjacency server functionality, and at the time of this writing this feature is not available in Cisco **NX-OS** Software for Cisco Nexus 7000 Series Switches.

This document focuses on deploying OTV over a multicast enable transport. One of the key advantages of using multicast is that it allows optimal multicast traffic replication to multiple sites and avoids head-end replication that leads to suboptimal bandwidth utilization.

In this testing environment each OTV edge device is configured with a common control group 239.1.1.2 and data group 239.192.1.0/24 on the overlay interface. The multicast control group identifies the overlay; two different overlays must have two different multicast control groups. The control group is used for neighbor discovery and to exchange MAC address reachability. The data group however is an SSM (Source Specific Group) group range, which is used to carry multicast data traffic generated by the sites

Once the overlay interface is configured with the control group and the “no shutdown” command is issued (make sure that the site VLAN configuration is already in place), the OTV edge device sends an IGMP report message to join the control group in the transport network. The OTV edge devices then exchange OTV control-plane hellos to build adjacencies with each other. Once the adjacencies are established, the OTV control-plane packets communicate MAC to IP mappings (MAC address reachability) to the OTV-adjacent devices. These update packets contain the Join Interface IP address, the VLAN IDs, and the learned MAC addresses that are reachable through the sending OTV edge device.

Tampa Data Center

T_N7K1-OTV# <pre>interface Overlay1 otv control-group 239.1.1.2 otv data-group 239.192.1.0/24 interface Ethernet1/17 description Join Interface to T_N7K1-</pre>	T_N7K2-OTV# <pre>interface Overlay1 otv control-group 239.1.1.2 otv data-group 239.192.1.0/24 interface Ethernet1/17 description Join Interface to T_N7K2-</pre>
--	--

AGG e1/10 L3 ip igmp version 3	AGG e1/10 L3 ip igmp version 3
-----------------------------------	-----------------------------------

Note: The control-plane protocol used by OTV is IS-IS. However, IS-IS does not need to be explicitly configured. It runs in the background once OTV is enabled.

In the aggregation layer, Protocol Independent Multicast (PIM) is configured on all intra- and inter-data-center Layer 3 links to allow multicast states to be built in the core network.

Tampa Data Center

T_N7K1-AGG# <pre>interface Ethernet1/2 description S_N7K1-AGG e1/2 L3 ip pim sparse-mode ! interface Ethernet1/10 description T_N7K1-OTV e1/17 L3 ip pim sparse-mode ip igmp version 3 ! interface Ethernet2/9 description T_N7K2-AGG e2/9 L3 ip pim sparse-mode</pre>	T_N7K2-AGG# <pre>interface Ethernet1/2 description S_N7K2-AGG e1/2 L3 ip pim sparse-mode ! interface Ethernet1/10 description T_N7K2-OTV e1/17 L3 ip pim sparse-mode ip igmp version 3 ! interface Ethernet2/9 description T_N7K1-AGG e2/9 L3 ip pim sparse-mode</pre>
--	--

Since PIM sparse mode requires rendezvous point (RP) to build a multicast tree, one of the aggregation switches in each data center is used as an RP. Local RP allows both local sources and receivers to join local RP rather than having to go to different data center to reach an RP in order to build a shared tree. Anycast-RP and Multicast Source Discovery Protocol (MSDP) features are used for local RP availability and provide redundancy in case of an RP failure. For more information about MSDP and Anycast features of multicast, visit:

http://www.cisco.com/en/US/docs/ios/solutions_docs/ip_multicast/White_papers/anycast.html

Tampa Data Center	Seattle Data Center
T_N7K1-AGG# <pre>ip pim rp-address 40.1.1.50 group-list 224.0.0.0/4 ip pim ssm range 232.0.0.0/8</pre>	S_N7K1-AGG# <pre>ip pim rp-address 40.1.1.50 group-list 224.0.0.0/4 ip pim ssm range 232.0.0.0/8</pre>
T_N7K2-AGG# <pre>interface loopback1 description MSDP Loopback ip address 40.1.1.50/32 ip pim sparse-mode ! ip pim rp-address 40.1.1.50 group-list 224.0.0.0/4 ip pim ssm range 232.0.0.0/8</pre>	S_N7K2-AGG# <pre>interface loopback1 description MSDP Loopback ip address 40.1.1.50/32 ip pim sparse-mode ! ip pim rp-address 40.1.1.50 group-list 224.0.0.0/4 ip pim ssm range 232.0.0.0/8</pre>

```
ip msdp originator-id loopback0
ip msdp peer 40.1.1.8 connect-source
loopback0
```

```
ip msdp originator-id loopback0
ip msdp peer 40.1.1.2 connect-source
loopback0
```

OTV Control Plane: Steady State Information

The OTV edge devices become part of the overlay when they join the multicast control group defined in the core; members of the same group are members of the same overlay. For example the control group used in this test environment is 239.1.1.2. This allows the hellos and updates between the overlay edge devices to travel over this control group defined in the core. Thus, edge devices peer with each other as if they were directly connected at Layer 2. This peering is possible as all the traffic for the OTV control plane is encapsulated with the control group address and sent into the core. Thus, all edge devices in a given overlay receive the control plane multicast traffic as if they were all on the same segment.

The mroute outputs shown in the following code, taken on S_N7K2-AGG and T_N7K1-AGG switches, show the multicast state information for the control group. As the output indicates, the edge devices are acting both as sources and receivers on control group 239.1.1.2, which allows them to send and receive hellos and share Intermediate System-to-Intermediate System (IS-IS) control plane information with all other edge devices.

```
S_N7K2-AGG# show ip mroute 239.1.1.2 summary
IP Multicast Routing Table for VRF "default"

(*, 239.1.1.2/32), uptime: 1d12h, igmp pim ip
  Incoming interface: loopback1, RPF nbr: 40.1.1.50
  Outgoing interface list: (count: 2)
    Ethernet2/9, uptime: 1d04h, pim
    Ethernet1/10, uptime: 1d12h, igmp

(20.0.1.1/32, 239.1.1.2/32), uptime: 1d04h, pim ip
  Incoming interface: Ethernet2/9, RPF nbr: 30.21.1.1, internal
  Outgoing interface list: (count: 1)
    Ethernet1/10, uptime: 1d04h, mrrib

(20.0.3.1/32, 239.1.1.2/32), uptime: 1d04h, pim ip msdp
  Incoming interface: Ethernet1/2, RPF nbr: 10.0.3.1
  Outgoing interface list: (count: 2)
    Ethernet2/9, uptime: 12:44:20, pim
    Ethernet1/10, uptime: 1d04h, mrrib

(30.0.1.1/32, 239.1.1.2/32), uptime: 1d08h, pim ip
  Incoming interface: Ethernet2/9, RPF nbr: 30.21.1.1, internal
  Outgoing interface list: (count: 2)
    Ethernet1/2, uptime: 12:44:20, pim
    Ethernet1/10, uptime: 1d08h, mrrib

(30.0.3.1/32, 239.1.1.2/32), uptime: 2d01h, ip pim mrrib
  Incoming interface: Ethernet1/10, RPF nbr: 30.0.3.1, internal
  Outgoing interface list: (count: 3)
    Ethernet1/2, uptime: 12:44:21, pim
    Ethernet2/9, uptime: 1d04h, pim
```

T_N7K1-AGG# show ip mroute 239.1.1.2

IP Multicast Routing Table for VRF "default"

```
(*, 239.1.1.2/32), uptime: 1d11h, pim ip igmp
  Incoming interface: Ethernet1/2, RPF nbr: 10.0.2.2
  Outgoing interface list: (count: 1)
    Ethernet1/10, uptime: 1d11h, igmp

(20.0.1.1/32, 239.1.1.2/32), uptime: 1d11h, ip pim mrib
  Incoming interface: Ethernet1/10, RPF nbr: 20.0.1.1
  Outgoing interface list: (count: 3)
    Ethernet2/9, uptime: 12:47:32, pim
    Ethernet1/2, uptime: 1d04h, pim

(20.0.3.1/32, 239.1.1.2/32), uptime: 1d04h, ip pim
  Incoming interface: Ethernet2/9, RPF nbr: 20.21.1.2
  Outgoing interface list: (count: 1)
    Ethernet1/10, uptime: 1d04h, mrib

(30.0.1.1/32, 239.1.1.2/32), uptime: 1d04h, ip pim
  Incoming interface: Ethernet1/2, RPF nbr: 10.0.2.2
  Outgoing interface list: (count: 1)
    Ethernet1/10, uptime: 1d04h, mrib

(30.0.3.1/32, 239.1.1.2/32), uptime: 1d04h, ip pim
  Incoming interface: Ethernet1/2, RPF nbr: 10.0.2.2
  Outgoing interface list: (count: 1)
    Ethernet1/10, uptime: 1d04h, mrib
```

After an edge device in the OTV overlay exchanges control plane information with devices at the other sites, it forms an adjacency with these devices. These adjacencies need to be up before the edge device at a site can send or receive Layer 2 information on the overlay. The following output shows the adjacencies formed by OTV edge devices in Tampa and Seattle with all other devices participating in the overlay, both within and across different sites.

Tampa DC (OTV VDC):**T_N7K1-OTV# show otv adjacency detail**

Overlay-Interface Overlay1 :

System-ID	HostName	Dest Addr	Up Time	Adj-State
0024.f718.7542	T_N7K2-OTV	20.0.3.1	4d02h	UP
HW-St: Up Peer-ID: 3 If-Name: Tunnel16454				
0024.98e9.0842	S_N7K1-OTV	30.0.1.1	4d02h	UP
HW-St: Up Peer-ID: 2 If-Name: Tunnel16453				
0024.98e9.05c3	S_N7K2-OTV	30.0.3.1	3d21h	UP
HW-St: Up Peer-ID: 1 If-Name: Tunnel16468				

T_N7K2-OTV# show otv adjacency detail

Overlay-Interface Overlay1 :

System-ID	HostName	Dest Addr	Up Time	Adj-State
0024.98e9.01c2	T_N7K1-OTV	20.0.1.1	4d02h	UP
HW-St: Up Peer-ID: 3 If-Name: Tunnel16428				
0024.98e9.0842	S_N7K1-OTV	30.0.1.1	6d22h	UP
HW-St: Up Peer-ID: 2 If-Name: Tunnel16412				
0024.98e9.05c3	S_N7K2-OTV	30.0.3.1	3d21h	UP
HW-St: Up Peer-ID: 1 If-Name: Tunnel16442				

Seattle DC (OTV VDC):

S_N7K1-OTV# show otv adjacency detail

Overlay-Interface Overlay1 :

System-ID	HostName	Dest Addr	Up Time	Adj-State
0024.98e9.01c2	T_N7K1-OTV	20.0.1.1	4d02h	UP
HW-St: Up Peer-ID: 2 If-Name: Tunnel16440				
0024.f718.7542	T_N7K2-OTV	20.0.3.1	6d22h	UP
HW-St: Up Peer-ID: 3 If-Name: Tunnel16424				
0024.98e9.05c3	S_N7K2-OTV	30.0.3.1	3d21h	UP
HW-St: Up Peer-ID: 1 If-Name: Tunnel16454				

S_N7K2-OTV# show otv adjacency detail

Overlay-Interface Overlay1 :

System-ID	HostName	Dest Addr	Up Time	Adj-State
0024.98e9.01c2		20.0.1.1	3d21h	UP
HW-St: Up Peer-ID: 2 If-Name: Tunnel16514				
0024.f718.7542		20.0.3.1	3d21h	UP
HW-St: Up Peer-ID: 1 If-Name: Tunnel16513				
0024.98e9.0842	S_N7K1-OTV	30.0.1.1	3d21h	UP

HW-St: Up Peer-ID: 3 If-Name: Tunnel16515

Multipathing and AED

When sites are multihomed, as in the case here, there could be potential for loops to be created between the OTV overlay and the Layer 2 domains at different sites. Unlike other Layer 2 VPN techniques OTV does not use Spanning Tree Protocol on the overlay to break any loops between sites. It has its own loop prevention mechanism, which allows each site to maintain its own Spanning Tree Protocol domain separate from all other sites.

OTV achieves this separation by electing an authoritative edge device (AED) for each VLAN, which is the only device that can forward the traffic for the extended VLAN inside and outside the data center. The AED election is based on an internal hash of VLAN ID and number of edge devices in a site. Today the extended VLANs are split in odd and even and automatically assigned to the site's edge devices. This results in an edge device being the AED for the odd VLANs and the other edge device being the AED for the even VLANs. This assignment is not configurable at this time and is done automatically by NX-OS.

In our test topology, each site (Tampa and Seattle) has an AED elected by OTV in each VLAN. For Example, as shown in the output below, **T_N7K1-OTV** was the AED for VLAN 10 in Tampa and **S_N7K2-OTV** was the AED for VLAN 10 in Seattle.

Tampa DC (OTV VDC):

T_N7K1-OTV#show otv vlan 10 authoritativeOTV VLAN Configuration Information (* - AED)

VLAN-ID	VlanState	Overlay
*10	active	Overlay1

Seattle DC(OTV VDC):

S_N7K2-OTV# show otv vlan 10 authoritative
OTV VLAN Configuration Information (* - AED)

VLAN-ID	VlanState	Overlay
*10	active	Overlay1

Once the overlay is up and AEDs elected at each site, the MAC addresses can be learnt and advertised across the overlay. The interface on which these addresses are learnt is based on the site's Layer 2 topology.

For example, in Tampa, VLAN 10 MAC address (**0024.98e9.01c3**) is learnt over the port channel "Po23" on T_N7K2 in OTV VDC and over the peer link on T_N7K2 in the "AGG" VDC, whereas the same MAC address is learnt over the overlay on S_N7K2 in OTV VDC.

Note: 0024.98e9.01c3 is the MAC address of VLAN 10 interface in **T_N7K1-AGG**.

Tampa DC (OTV VDC):

T_N7K1-OTV# show mac address-table vlan 10 address 0024.98e9.01c3

VLAN	MAC Address	Type	age	Secure	NTFY	Ports
* 10	0024.98e9.01c3	dynamic	0	F	F	Po21

T_N7K2-OTV# show mac address-table vlan 10 address 0024.98e9.01c3

VLAN	MAC Address	Type	age	Secure	NTFY	Ports
* 10	0024.98e9.01c3	dynamic	240	F	F	Po23

Tampa DC(Default "AGG" VDC):

T_N7K1-AGG# show mac address-table vlan 10 address 0024.98e9.01c3

VLAN	MAC Address	Type	age	Secure	NTFY	Ports
G 10	0024.98e9.01c3	static	-	F	F	sup-eth1(R)

T_N7K2-AGG# show mac address-table vlan 10 address 0024.98e9.01c3

VLAN	MAC Address	Type	age	Secure	NTFY	Ports
* 10	0024.98e9.01c3	static	-	F	F	vPC Peer-Link

Seattle DATA CENTER(OTV VDC):

S_N7K1-OTV# show mac address-table vlan 10 address 0024.98e9.01c3

VLAN	MAC Address	Type	age	Secure	NTFY	Ports
* 10	0024.98e9.01c3	dynamic	210	F	F	Po21

S_N7K2-OTV# show mac address-table vlan 10 address 0024.98e9.01c3

VLAN	MAC Address	Type	age	Secure	NTFY	Ports
O 10	0024.98e9.01c3	dynamic	0	F	F	Overlay1

Seattle DATA CENTER(Default "AGG" VDC):

S_N7K1-AGG# show mac address-table vlan 10 address 0024.98e9.01c3

VLAN	MAC Address	Type	age	Secure	NTFY	Ports
* 10	0024.98e9.01c3	dynamic	0	F	F	Po23

S_N7K2-AGG# show mac address-table vlan 10 address 0024.98e9.01c3

VLAN	MAC Address	Type	age	Secure	NTFY	Ports
* 10	0024.98e9.01c3	dynamic	0	F	F	Po23

Appendix

This appendix shows relevant extracts of the various Cisco Nexus VDC configurations.

T_N7K1 Configuration

```
T_N7K1#

version 5.0(2)
feature telnet
feature tacacs+
cfs eth distribute
feature pim
feature msdp
feature eigrp
feature udld
```

```

feature interface-vlan
feature hsrp
feature vpc

hostname T_N7K1

vrf context management
  ip route 0.0.0.0/0 14.1.1.1
vlan 1
vdc T_N7K1 id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  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 OTV id 2
  allocate interface Ethernet1/9,Ethernet1/11,Ethernet1/13,
Ethernet1/15,Ethernet1/17, Ethernet1/19, Ethernet1/21,Ethernet1/23
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  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
vdc AGG id 3
  allocate interface Ethernet1/1-
8,Ethernet1/10,Ethernet1/12,Ethernet1/14,Ethernet1/16
  allocate interface Ethernet2/1-16
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  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

```

T_N7K1-OTV Configuration

T_N7K1-OTV#

```

version 5.0(2)
feature telnet

```



```
cfs eth distribute
feature pim
feature eigrp
feature otv
feature udld
feature lacp
feature bfd

hostname OTV
ip access-list ALL_IPs
  10 permit ip 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

vlan access-map HSRP_Localization 10
  match ip address HSRP_IP
  action drop
vlan access-map HSRP_Localization 20
  match ip address ALL_IPs
  action forward

vlan filter HSRP_Localization vlan-list 10
vrf context management
vlan 1,10,987
spanning-tree vlan 1-3967,4048-4093 priority 49152
mac-list OTV_HSRP_VMAC_deny seq 10 deny 0000.0c07.ac00 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 20 deny 0000.0c9f.f000 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 30 permit 0000.0000.0000 0000.0000.0000

route-map OTV_HSRP_filter permit 10
  match mac-list OTV_HSRP_VMAC_deny

interface port-channel21
  description T_N7K-AGG
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,987

interface Overlay1
  otv join-interface Ethernet1/17
  otv control-group 239.1.1.2
  otv data-group 239.192.1.0/24
  otv extend-vlan 10
  no otv suppress-arp-nd
  description Overlay for AGG
  no shutdown

interface Ethernet1/9
```

```
switchport
switchport mode trunk
channel-group 21 mode active
switchport trunk allowed vlan 10,987
no shutdown

interface Ethernet1/11
switchport
switchport mode trunk
channel-group 21 mode active
switchport trunk allowed vlan 10,987
no shutdown

interface Ethernet1/17
description T_N7K1-AGG e1/10 L3
mtu 9216
no ip redirects
ip address 20.0.1.1/30
ip router eigrp 100
ip igmp version 3
no shutdown

interface loopback0
ip address 40.1.1.5/32
ip router eigrp 100

router eigrp 100
bfd

otv-isis default

vpn Overlay1
redistribute filter route-map OTV_HSRP_filter

otv site-vlan 987
```

T_N7K1-AGG Configuration

```
T_N7K1-AGG#

version 5.0(2)
feature telnet
cfs eth distribute
feature pim
feature eigrp
feature udld
feature interface-vlan
feature hsrp
```

```
feature lacp
feature vpc
feature bfd

hostname T_N7K1-AGG

vrf context keepalive
vrf context management
vlan 1
vlan 10
    name UCS-Hosts

spanning-tree vlan 1-3967,4048-4093 priority 24576
vpc domain 101
    peer-keepalive destination 20.21.0.2 source 20.21.0.1 vrf keepalive

interface Vlan1

interface Vlan10
    no shutdown
    description UCS
    no ip redirects
    ip address 172.16.10.2/24
    ip router eigrp 100
    ip passive-interface eigrp 100
    hsrp 10
        preempt
        ip 172.16.10.1

interface port-channel1
    switchport
    switchport mode trunk
    vpc peer-link
    spanning-tree port type network

interface port-channel21
    description VPC to T_N7K1-OTV
    switchport
    switchport mode trunk
    vpc 21
    switchport trunk allowed vlan 10,987

interface port-channel23
    description VPC to T_N7K2-OTV
    switchport
    switchport mode trunk
    vpc 23
    switchport trunk allowed vlan 10,987
```

```
interface port-channel30
  description Port-channel to UCS
  switchport
  switchport mode trunk
  vpc 30
  switchport trunk allowed vlan 10

interface Ethernet1/1
  description VPC peer e1/1 e2/1 T_NK2-AGG
  switchport
  switchport mode trunk
  channel-group 1 mode active
  no shutdown

interface Ethernet1/2
  description S_N7K1-AGG e1/2 L3
  mtu 9216
  no ip redirects
  ip address 10.0.2.1/30
  ip router eigrp 100
  ip pim sparse-mode
  no shutdown

interface Ethernet1/4
  description PO21 to T_N7K1-OTV e1/9
  switchport
  switchport mode trunk
  channel-group 21 mode active
  switchport trunk allowed vlan 10,987
  no shutdown

interface Ethernet1/6
  description PO23 to T_N7K2_OTV 1/11
  switchport
  switchport mode trunk
  channel-group 23 mode active
  switchport trunk allowed vlan 10,987
  no shutdown

interface Ethernet1/10
  description T_N7K1-OTV e1/17 L3
  mtu 9216
  no ip redirects
  ip address 20.0.1.2/30
  ip router eigrp 100
  ip pim sparse-mode
  ip igmp version 3
  no shutdown
```

```
interface Ethernet1/16
  description vpc peer
  vrf member keepalive
  ip address 20.21.0.1/30
  no shutdown

interface Ethernet2/1
  description VPC peer e1/1 e2/1 T_N7K2-AGG
  switchport
  switchport mode trunk
  channel-group 1 mode active
  no shutdown

interface Ethernet2/2
  description port-channel to UCS
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10
  spanning-tree port type edge trunk
  channel-group 30 mode active
  no shutdown

interface Ethernet2/9
  description T_N7K2-AGG e2/9 L3
  mtu 9216
  no ip redirects
  ip address 20.21.1.1/30
  ip router eigrp 100
  ip pim sparse-mode
  no shutdown

interface loopback0
  ip address 40.1.1.1/32
  ip router eigrp 100
  ip pim sparse-mode

router eigrp 100
  autonomous-system 100
  bfd
  ip pim rp-address 40.1.1.50 group-list 224.0.0.0/4
  ip pim ssm range 232.0.0.0/8
```

T_N7K2 Configuration

```
T_N7K2#

version 5.0(2)
feature telnet
```

```

feature tacacs+
cfs eth distribute
feature pim
feature msdp
feature eigrp
feature udld
feature interface-vlan
feature hsrp
feature vpc

hostname T_N7K2
vrf context management
    ip route 0.0.0.0/0 14.1.83.1
vlan 1
spanning-tree vlan 1-3967,4048-4093 priority 28672

vdc T_N7K2 id 1
    limit-resource vlan minimum 16 maximum 4094
    limit-resource monitor-session minimum 0 maximum 2
    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 OTV id 2
    allocate interface Ethernet1/9,Ethernet1/11,Ethernet1/13,Ethernet1/15,
Ethernet1/17, Ethernet1/19, Ethernet1/21,Ethernet1/23
    boot-order 1
    limit-resource vlan minimum 16 maximum 4094
    limit-resource monitor-session minimum 0 maximum 2
    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
vdc AGG id 3
    allocate interface Ethernet1/1-
8,Ethernet1/10,Ethernet1/12,Ethernet1/14,Ethernet1/16
    allocate interface Ethernet2/1-16
    boot-order 1
    limit-resource vlan minimum 16 maximum 4094
    limit-resource monitor-session minimum 0 maximum 2
    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

```

T_N7K2-OTV Configuration

```
T_N7K2-OTV#

version 5.0(2)
feature telnet
cfs eth distribute
feature pim
feature eigrp
feature otv
feature udld
feature lacp
feature bfd

hostname OTV
ip access-list ALL_IPs
  10 permit ip 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

vlan access-map HSRP_Localization 10
  match ip address HSRP_IP
  action drop
vlan access-map HSRP_Localization 20
  match ip address ALL_IPs
  action forward

vlan filter HSRP_Localization vlan-list 10
vrf context management
vlan 1,10,987
spanning-tree vlan 1-3967,4048-4093 priority 49152
mac-list OTV_HSRP_VMAC_deny seq 10 deny 0000.0c07.ac00 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 20 deny 0000.0c9f.f000 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 30 permit 0000.0000.0000 0000.0000.0000

route-map OTV_HSRP_filter permit 10
  match mac-list OTV_HSRP_VMAC_deny

interface port-channel23
  description T-N7K-AGG
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,987

interface Overlay1
  otv join-interface Ethernet1/17
  otv control-group 239.1.1.2
```

```

otv data-group 239.192.1.0/24
otv extend-vlan 10
no otv suppress-arp-nd
description Overlay for AGG
no shutdown

interface Ethernet1/9
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,987
  channel-group 23 mode active
  no shutdown

interface Ethernet1/11
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,987
  channel-group 23 mode active
  no shutdown

interface Ethernet1/17
  description T_N7K2-AGG e1/10 L3
  mtu 9216
  no ip redirects
  ip address 20.0.3.1/30
  ip router eigrp 100
  ip igmp version 3
  no shutdown

interface loopback0
  ip address 40.1.1.6/32
  ip router eigrp 100

router eigrp 100
  bfd
  otv-isis default

  vpn Overlay1
    redistribute filter route-map OTV_HSRP_filter
ip pim ssm range 232.0.0.0/8
otv site-vlan 987

```

T_N7K2-AGG Configuration

```

T_N7K2-AGG#

version 5.0(2)
feature telnet
cfs eth distribute
feature pim

```



```
feature msdp
feature eigrp
feature udld
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature bfd

hostname T_N7K2-AGG

vrf context keepalive
vrf context management
vlan 1
vlan 10
    name UCS-Hosts

spanning-tree vlan 1-3967,4048-4093 priority 28672
vpc domain 101
    peer-keepalive destination 20.21.0.1 source 20.21.0.2 vrf keepalive

interface Vlan1

interface Vlan10
    no shutdown
    description UCS
    no ip redirects
    ip address 172.16.10.3/24
    ip router eigrp 100
    ip passive-interface eigrp 100
    hsrp 10
        ip 172.16.10.1

interface port-channel1
    description peer link T-CS1
    switchport
    switchport mode trunk
    vpc peer-link
    spanning-tree port type network

interface port-channel21
    description VPC to T_N7K1-OTV
    switchport
    switchport mode trunk
    vpc 21
    switchport trunk allowed vlan 10,987

interface port-channel23
```

```
description VPC to T_N7K2-OTV
switchport
switchport mode trunk
vpc 23
switchport trunk allowed vlan 10,987

interface port-channel30
description PO30 to UCS
switchport
switchport mode trunk
vpc 30
switchport trunk allowed vlan 10

interface Ethernet1/1
description peer link po1
switchport
switchport mode trunk
channel-group 1 mode active
no shutdown

interface Ethernet1/2
description S_N7K2-AGG e1/2 L3
mtu 9216
no ip redirects
ip address 10.0.3.1/30
ip router eigrp 100
ip pim sparse-mode
no shutdown

interface Ethernet1/4
description T_N7K2-OTV e1/9
switchport
switchport mode trunk
switchport trunk allowed vlan 10,987
channel-group 23 mode active
no shutdown

interface Ethernet1/6
description PO21 to T_N7K1-OTV e1/11
switchport
switchport mode trunk
switchport trunk allowed vlan 10,987
channel-group 21 mode active
no shutdown

interface Ethernet1/10
description T_N7K2-OTV e1/17 L3
mtu 9216
no ip redirects
```

```
ip address 20.0.3.2/30
ip router eigrp 100
ip pim sparse-mode
ip igmp version 3
no shutdown

interface Ethernet1/16
description vpc peer
vrf member keepalive
ip address 20.21.0.2/30
no shutdown

interface Ethernet2/1
description peer link po1
switchport
switchport mode trunk
channel-group 1 mode active
no shutdown

interface Ethernet2/2
description port-channel to UCS
switchport
switchport mode trunk
switchport trunk allowed vlan 10
spanning-tree port type edge trunk
channel-group 30 mode active
no shutdown

interface Ethernet2/9
description T_N7K1-AGG e2/9 L3
mtu 9216
no ip redirects
ip address 20.21.1.2/30
ip router eigrp 100
ip pim sparse-mode
no shutdown

interface loopback0
ip address 40.1.1.2/32
ip router eigrp 100
ip pim sparse-mode

interface loopback1
description MSDP Loopback
ip address 40.1.1.50/32
ip router eigrp 100
```

```

ip pim sparse-mode

router eigrp 100
  bfd
ip pim rp-address 40.1.1.50 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8
ip msdp originator-id loopback0
ip msdp peer 40.1.1.8 connect-source loopback0

```

S_N7K1 Configuration

```

S_N7K1#

version 5.0(2)
feature telnet
feature tacacs+
cfs eth distribute
feature pim
feature msdp
feature eigrp
feature udld
feature interface-vlan
feature hsrp
feature vpc

hostname S_N7K1

vrf context management
  ip route 0.0.0.0/0 14.1.81.1
vlan 1
vdc S_N7K1 id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  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 OTV id 2
  allocate interface
  Ethernet1/9,Ethernet1/11,Ethernet1/13,Ethernet1/15,Ethernet1/17,Ethernet1/19,
  Ethernet1/21,Ethernet1/23
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  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
vdc AGG id 3
    allocate interface Ethernet1/1-
8,Ethernet1/10,Ethernet1/12,Ethernet1/14,Ethernet1/16
    allocate interface Ethernet2/1-16
    boot-order 1
    limit-resource vlan minimum 16 maximum 4094
    limit-resource monitor-session minimum 0 maximum 2
    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

```

S_N7K1-OTV Configuration

```

S_N7K1-OTV#

version 5.0(2)
feature telnet
cfs eth distribute
feature pim
feature eigrp
feature otv
feature udld
feature lacp
feature bfd

hostname OTV
ip access-list ALL_IPs
    10 permit ip 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

vlan access-map HSRP_Localization 10
    match ip address HSRP_IP
    action drop
vlan access-map HSRP_Localization 20
    match ip address ALL_IPs
    action forward

vlan filter HSRP_Localization vlan-list 10
vrf context management
vlan 1,10,987
spanning-tree vlan 1-3967,4048-4093 priority 49152
mac-list OTV_HSRP_VMAC_deny seq 10 deny 0000.0c07.ac00 ffff.ffff.ff00

```

```
mac-list OTV_HSRP_VMAC_deny seq 20 deny 0000.0c9f.f000 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 30 permit 0000.0000.0000 0000.0000.0000

route-map OTV_HSRP_filter permit 10
  match mac-list OTV_HSRP_VMAC_deny

interface port-channel21
  description T-N7K-AGG
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,987

interface Overlay1
  otv join-interface Ethernet1/17
  otv control-group 239.1.1.2
  otv data-group 239.192.1.0/24
  otv extend-vlan 10
  no otv suppress-arp-nd
  description Overlay for AGG
  no shutdown

interface Ethernet1/9
  switchport
  switchport mode trunk
  channel-group 21 mode active
  switchport trunk allowed vlan 10,987
  no shutdown

interface Ethernet1/11
  switchport
  switchport mode trunk
  channel-group 21 mode active
  switchport trunk allowed vlan 10,987
  no shutdown

interface Ethernet1/17
  description to S_N7K1-AGG e1/10
  mtu 9216
  no ip redirects
  ip address 30.0.1.1/30
  ip router eigrp 100
  ip igmp version 3
  no shutdown

interface loopback0
  ip address 40.1.1.11/32
  ip router eigrp 100
```

```
router eigrp 100
  router-id 91.1.1.1
  bfd

otv-isis default
  vpn Overlay1
  redistribute filter route-map OTV_HSRP_filter

ip pim ssm range 232.0.0.0/8
otv site-vlan 987
```

S_N7K1-AGG configuration

```
S_N7K1-AGG#

version 5.0(2)
feature telnet
cfs eth distribute
feature pim
feature eigrp
feature udld
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature bfd

hostname S_N7K1-AGG

vrf context keepalive
vrf context management
vlan 1
vlan 10
  name UCS-Hosts

spanning-tree vlan 1-3967,4048-4093 priority 24576
vpc domain 101
  peer-keepalive destination 30.21.0.2 source 30.21.0.1 vrf keepalive

interface Vlan1

interface Vlan10
  no shutdown
  no ip redirects
  ip address 172.16.10.4/24
  ip router eigrp 100
  ip passive-interface eigrp 100
  hsrp 10
```

```
ip 172.16.10.1

interface port-channel1
  switchport
  switchport mode trunk
  vpc peer-link
  spanning-tree port type network

interface port-channel21
  description vPC to S_N7k1-OTV
  switchport
  switchport mode trunk
  vpc 21
switchport trunk allowed vlan 10,987

interface port-channel23
  description vPC to S_N7k2-OTV
  switchport
  switchport mode trunk
  vpc 23
switchport trunk allowed vlan 10,987

interface port-channel30
  description port-channel to UCS Seattle
  switchport
  switchport mode trunk
  vpc 30
switchport trunk allowed vlan 10

interface Ethernet1/1
  description VPC peer e1/1 e2/1 S_N7K2-AGG
  switchport
  switchport mode trunk
  channel-group 1 mode active
  no shutdown

interface Ethernet1/2
  description L3 Connection to T_N7K1-AGG E1/2
  mtu 9216
  no ip redirects
  ip address 10.0.2.2/30
  ip router eigrp 100
  ip pim sparse-mode
  no shutdown

interface Ethernet1/4
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,987
```



```
channel-group 21 mode active
no shutdown

interface Ethernet1/6
switchport
switchport mode trunk
switchport trunk allowed vlan 10,987
channel-group 23 mode active
no shutdown
```

```
interface Ethernet1/10

description S_N7K2-OTV e1/17 L3
mtu 9216
no ip redirects
ip address 30.0.1.2/30
ip router eigrp 100
ip pim sparse-mode
ip igmp version 3
no shutdown

interface Ethernet1/16
description to S_7K2 e1/16
vrf member keepalive
ip address 30.21.0.1/30
no shutdown

interface Ethernet2/1
switchport
switchport mode trunk
channel-group 1 mode active
no shutdown

interface Ethernet2/2
description Connection to UCS
switchport
switchport mode trunk
switchport trunk allowed vlan 10
spanning-tree port type edge trunk
channel-group 30 mode active
no shutdown

interface Ethernet2/9
description to S_7K2 e2/9
mtu 9216
no ip redirects
ip address 30.21.1.1/30
ip router eigrp 100
ip pim sparse-mode
no shutdown
```

```

interface loopback0
  ip address 40.1.1.7/32
  ip router eigrp 100
  ip pim sparse-mode

router eigrp 100
  bfd
  ip pim rp-address 40.1.1.50 group-list 224.0.0.0/4
  ip pim ssm range 232.0.0.0/8

```

S_N7K2 Configuration

```

S_N7K2#
version 5.0(2)

feature telnet
feature tacacs+
cfs eth distribute
feature pim
feature msdp
feature eigrp
feature udld
feature interface-vlan
feature hsrp
feature vpc

hostname S_N7K2

vrf context management
  ip route 0.0.0.0/0 14.1.81.1
vlan 1
vdc S_N7K2 id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource vrf minimum 16 maximum 200
  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 CS id 2
  allocate interface Ethernet1/1-
8,Ethernet1/10,Ethernet1/12,Ethernet1/14,Ethernet1/16
  allocate interface Ethernet2/1-16
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2

```

```

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
vdc OTV id 3
  allocate interface Ethernet1/9,Ethernet1/11,Ethernet1/13,Ethernet1/15,
Ethernet1/17, Ethernet1/19, Ethernet1/21,Ethernet1/23
  boot-order 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  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

```

S_N7K2-OTV Configuration

```

S_N7K2-OTV#

version 5.0(2)
feature telnet
cfs eth distribute
feature pim
feature eigrp
feature otv
feature udld
feature lacp
feature bfd

hostname OTV
ip access-list ALL_IPs
  10 permit ip 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

vlan access-map HSRP_Localization 10
  match ip address HSRP_IP
  action drop
vlan access-map HSRP_Localization 20
  match ip address ALL_IPs
  action forward

vlan filter HSRP_Localization vlan-list 10

```

```
vrf context management
vlan 1,10,987
spanning-tree vlan 1-3967,4048-4093 priority 49152
mac-list OTV_HSRP_VMAC_deny seq 10 deny 0000.0c07.ac00 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 20 deny 0000.0c9f.f000 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 30 permit 0000.0000.0000 0000.0000.0000

route-map OTV_HSRP_filter permit 10
  match mac-list OTV_HSRP_VMAC_deny

interface port-channel23
  description S-N7K-AGG
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,987

interface Overlay1
  otv join-interface Ethernet1/17
  otv control-group 239.1.1.2
  otv data-group 239.192.1.0/24
  otv extend-vlan 10
  no otv suppress-arp-nd
  description Overlay for AGG
  no shutdown

interface Ethernet1/9
  switchport
  switchport mode trunk
  channel-group 23 mode active
  switchport trunk allowed vlan 10,987
  no shutdown

interface Ethernet1/11
  switchport
  switchport mode trunk
  channel-group 23 mode active
  switchport trunk allowed vlan 10,987
  no shutdown

interface Ethernet1/17
  description to S_N7K2-AGG e1/10
  mtu 9216
  no ip redirects
  ip address 30.0.3.1/30
  ip router eigrp 100
  ip igmp version 3
  no shutdown
```

```
interface loopback0
  ip address 40.1.1.12/32
  ip router eigrp 100

router eigrp 100
  bfd

otv-isis default
  vpn Overlay1
  redistribute filter route-map OTV_HSRP_filter

ip pim ssm range 232.0.0.0/8
otv site-vlan 987
```

S_N7K2-AGG Configuration

```
S_N7K2-AGG#

version 5.0(2)
feature telnet
cfs eth distribute
feature pim
feature msdp
feature eigrp
feature udld
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature bfd

hostname S_N7K2-AGG

vrf context keepalive
vrf context management
vlan 1
vlan 10
  name UCS-Hosts

spanning-tree vlan 1-3967,4048-4093 priority 28672
vpc domain 101
  peer-keepalive destination 30.21.0.1 source 30.21.0.2 vrf keepalive

interface Vlan1

interface Vlan10
  no shutdown
  no ip redirects
```

```
ip address 172.16.10.5/24
ip router eigrp 100
ip passive-interface eigrp 100
hsrp 10
    ip 172.16.10.1

interface port-channel1
    switchport
    switchport mode trunk
    vpc peer-link
    spanning-tree port type network

interface port-channel21
    switchport
    switchport mode trunk
    vpc 21
    switchport trunk allowed vlan 10,987

interface port-channel23
    switchport
    switchport mode trunk
    vpc 23
    switchport trunk allowed vlan 10,987

interface port-channel30
    description port-channel to UCS
    switchport
    switchport mode trunk
    vpc 30
    switchport trunk allowed vlan 10

interface Ethernet1/1
    switchport
    switchport mode trunk
    channel-group 1 mode active
    no shutdown

interface Ethernet1/2
    description to T_N7K2 e1/2
    mtu 9216
    no ip redirects
    ip address 10.0.3.2/30
    ip router eigrp 100
    ip pim sparse-mode
    no shutdown

interface Ethernet1/4
    switchport
    switchport mode trunk
```

```
switchport trunk allowed vlan 10,987
channel-group 23 mode active
no shutdown

interface Ethernet1/6
switchport
switchport mode trunk
switchport trunk allowed vlan 10,987
channel-group 21 mode active
no shutdown

interface Ethernet1/10
description S_N7K1-OTV e1/17 L3
mtu 9216
no ip redirects
ip address 30.0.3.2/30
ip router eigrp 100
ip pim sparse-mode
ip igmp version 3
no shutdown

interface Ethernet1/16
description Keepalive to S_N7K1-AGG e1/16
vrf member keepalive
ip address 30.21.0.2/30
no shutdown

interface Ethernet2/1
switchport
switchport mode trunk
channel-group 1 mode active
no shutdown

interface Ethernet2/2
description Connection to UCS
switchport
switchport mode trunk
switchport trunk allowed vlan 10
spanning-tree port type edge trunk
channel-group 30 mode active
no shutdown

interface Ethernet2/9
description to S_N7K1-AGG e2/9
mtu 9216
no ip redirects
ip address 30.21.1.2/30
ip router eigrp 100
ip pim sparse-mode
```

```
no shutdown

interface loopback0
  ip address 40.1.1.8/32
  ip router eigrp 100
  ip pim sparse-mode

interface loopback1
  description MSDP Loopback
  ip address 40.1.1.50/32
  ip router eigrp 100
  ip pim sparse-mode

router eigrp 100
  bfd
  ip pim rp-address 40.1.1.50 group-list 224.0.0.0/4
  ip pim ssm range 232.0.0.0/8
  ip msdp originator-id loopback0
  ip msdp peer 40.1.1.2 connect-source loopback0
```



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 www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at 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. (1005R)