White Paper

Layer 2 Everywhere: Overcoming Overlay Transport Virtualization (OTV) Site Limitations Within and Between Data Centers

Design Document

March 29, 2012

The specifications and information regarding the products in this manual are subject to change without notice. All statements, information, and recommendations in this manual are believed to be accurate but are presented without warranty of any kind, express or implied. Users must take full responsibility for their application of any products.

The software license and limited warranty for the accompanying product are set forth in the information packet that shipped with the product and are incorporated herein by this reference. If you are unable to locate the software license or limited warranty, contact your Cisco representative for a copy.

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

Notwithstanding any other warranty herein, all document files and software of these suppliers are provided "as is" with all faults. Cisco and the above-named suppliers disclaim all warranties, expressed or implied, including, without limitation, those of merchantability, fitness for a particular purpose and non-infringement or arising from a course of dealing, usage, or trade practice.

In no event shall Cisco or its suppliers be liable for any indirect, special, consequential, or incidental damages, including, without limitation, lost profits or loss or damage to data arising out of the use or inability to use this manual, even if Cisco or its suppliers have been advised of the possibility of such damages.

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 http://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.

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

Contents

List of Figures and Tables

About Authors

Xiaojie Zhang: Double CCIE #25063, is a Network Consulting Engineer with Cisco Systems in DC&V Network Architecture team. She has held a wide variety of network design, consulting, technical marketing, and solution validation positions. Her technical expertise includes extensive experience in designing and supporting complex networks, with a special focus on data center interconnect.

Azeem Suleman, Double CCIE #23427, is a Solutions Architect with Cisco Systems in DC&V Network Architecture team. He has written several documents, white papers and technical tips related to Data Center and Security on Cisco Connection Online, CCO (http://www.cisco.com). He has represented in Architecture Board and has been designing large scale Enterprise network for over a decade now.

Azeem has BS degree in Computer Engineering from SSUET, Pakistan and MS degree from UTA, TX. He has published IEEE Conference publication and has been with Cisco Advanced Services since 2009.

There are other contributors who helped in this whitepaper:

- Balaji Senthamil Selvan: Consultant Engineer with Cisco Systems in DC&V Network Architecture team.
- Hassan Durrani: Network Consulting Engineer with Cisco Systems in DC&V Network Architecture team.
- Faraz Siddiqui: Network Consulting Engineer with Cisco Systems in DC&V Network Architecture team.
- Vinay Suvarna: Consultant Engineer with Cisco Systems in DC&V Network Architecture team.

Technical review was done by industry experts including:

- Talha Hashmi: Manager with Cisco Systems in DC&V Network Architecture team.
- Victor Moreno: Distinguished Engineer with Cisco Systems in ECBU-NX-OS and Nexus 7000 Product Management team.
- Max Ardica: Technical Leader with Cisco Systems in SPB Cloud Computing Dev- SDU France team.
- Anoop Dawani: Technical Marketing Engineer with Cisco Systems in ECBU-NX-OS and Nexus 7000 Product Management team.

1. Introduction

Overlay Transport Virtualization (OTV) helps solve many of the challenges in today's increasingly virtualized and geographically dispersed data centers. OTV enables data centers to extend Layer 2 over the existing IP network.

This document introduces the following:

- New design that overcomes the site limitation inherent in OTV technology
- Flow within and between data centers
- Snapshot of the configuration
- End-to-end results of tests run on Cisco® NX-OS Software Release 5.2(3)a

The information in this document is for readers who have a basic understanding of OTV and IP network technology and the way that OTV works using the concept of Layer 2 over Layer 3.

1.1 Document Scope

This document discusses a data center interconnect (DCI) design option to demonstrate the use of OTV to interconnect separated Layer 2 domains within a data center as well as across data centers. The primary focus of the document is on a new design to overcome the site limitation inherent in OTV technology. This limitation is explained later.

The intent of this document is to highlight an OTV deployment and not to advise on the best practices or technology options for DCI designs. Some of the benefits of using OTV instead of other Layer 2 extension technologies include:

- Deployment of Ethernet over MPLS (EoMPLS) or Virtual Private LAN Services (VPLS) not required
- Layer 2 and Layer 3 connectivity provided using the same dark fiber connections
- Native Spanning Tree Protocol isolation; Bridge Protocol Data Unit (BPDU) filtering does not need to be explicitly configured
- Native unknown unicast flooding isolation; unknown unicast traffic is not sent to the overlay
- Native multihoming support
- Address Resolution Protocol (ARP) optimization with the OTV ARP cache
- Easy addition of sites

2. Business Requirements

2.1 Current OTV Design

Most large enterprises have data centers in different regions. Within each data center, multiple Layer 2 domains may be connected through Layer 3 routing. Extension of the same set of VLANs across all the Layer 2 domains within and between data centers can be a challenge. Site limitations currently exist and will be addressed in future software releases.

Figure 1 shows such a challenge. In current software release Cisco NX-OS 5.2, OTV supports a maximum of six sites. Thus, using current OTV design, you can connect up to six Layer 2 domains in Data Center West and Data Center East, leaving the rest of the Layer 2 domains in Data Center West and East and Data Center North isolated. To connect all the Layer 2 domains using OTV in the current situation, you have to use a hierarchical model, which is the proposed design until the enhanced release is available.

Figure 1. One-Tier OTV Design

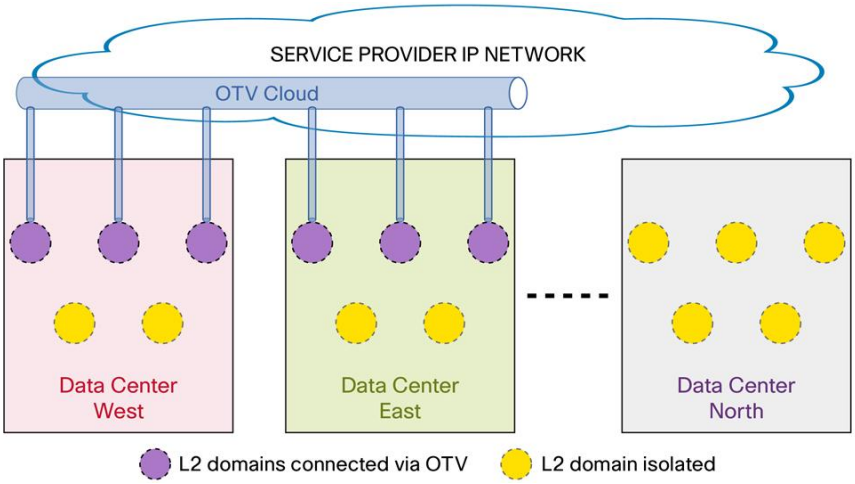

2.2 Two-Tier OTV Design

Cisco's two-tier OTV design can resolve this site limitation. With the current software release, Cisco NX-OS 5.2, which can support up to six sites in one OTV overlay, the two-tier design can extend to 30 the number of total supported sites. Using this design, the Layer 2 domains in the previous example are all connected through OTV, as shown in Figure 2.

Figure 2. Two-Tier OTV Design Proposal

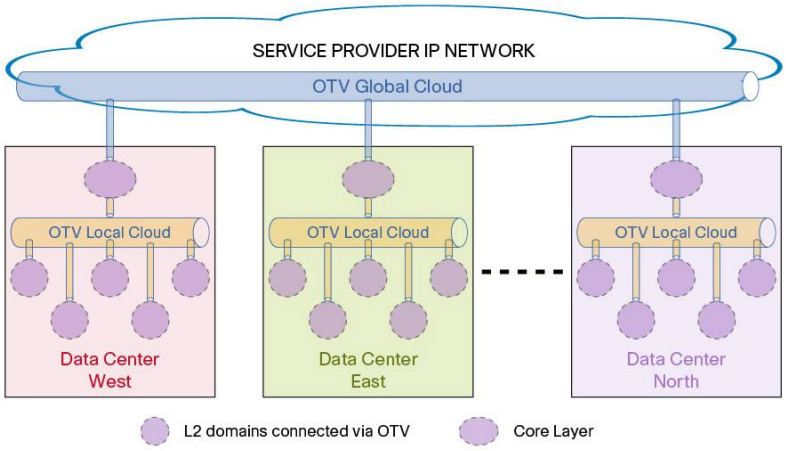

3. Design Proposal

3.1 Overview of Two-Tier OTV

Cisco proposes a two-tier OTV architecture based on a customer's current network to overcome the OTV site limitation and extend Layer 2 connections among all Layer 2 domains in all the distributed data centers, as shown in Figure 3.

Figure 3. Two-Tier OTV Design

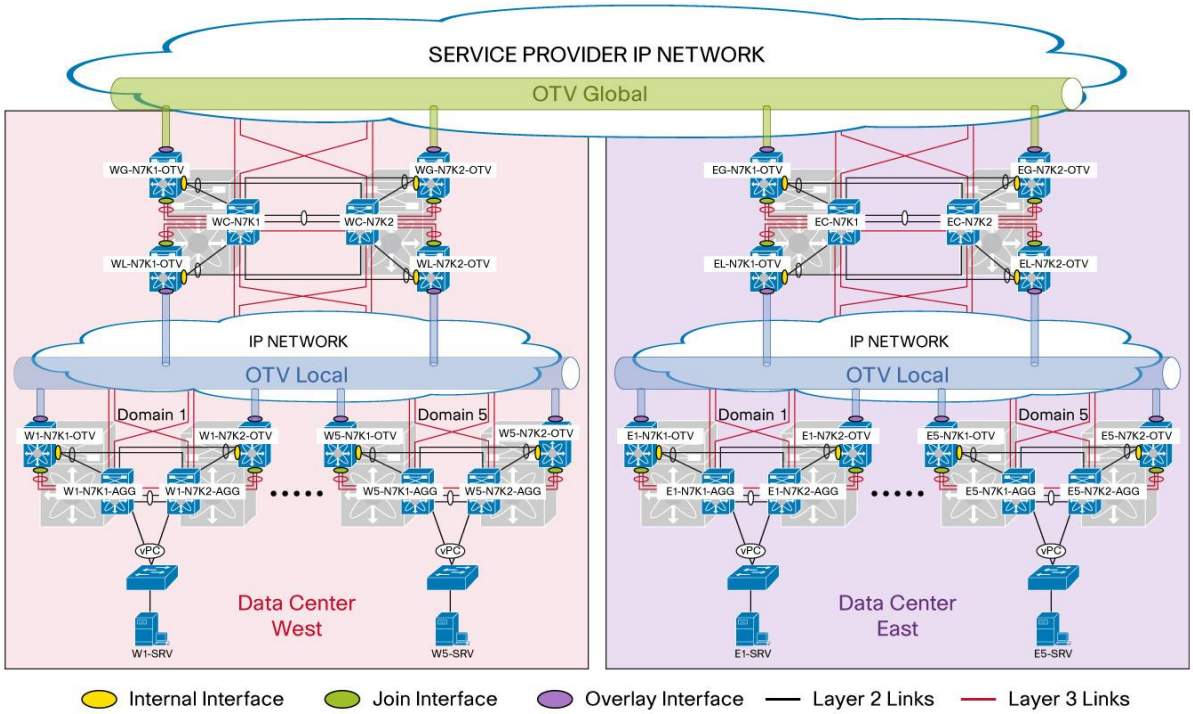

In this design, you create two OTV clouds, summarized in Table 1.

Table 1. Two-Tier OTV Cloud Proposal

Two-Tier OTV Cloud Proposal	
Cloud	**Description**
OTV Local	Extend the Layer 2 connection among all the Layer 2 domains and the core layer within the same data center.
OTV Global	Extend the Layer 2 connections among the core layer for all data centers.

3.2 Core Layer Design

On a customer's existing core-layer device - which must be a Cisco Nexus® 7000 Series Switch - in each data center, you create two additional virtual device contexts (VDCs). You use one VDC to join the local OTV cloud and the other to join the Global OTV cloud. You then create additional Layer 2 links in between to allow the two core devices to form virtual PortChannel (vPC) pair, and you connect each OTV VDC through vPC. You also create additional Layer 3 links between OTV VDC and Core VDC to provide Layer 3 routing to each OTV VDC.
In addition, you can use a Layer 3 PortChannel to provide redundancy.

Figure 4 shows the core layer at Data Center West.

Figure 4. Data Center West Core Layer

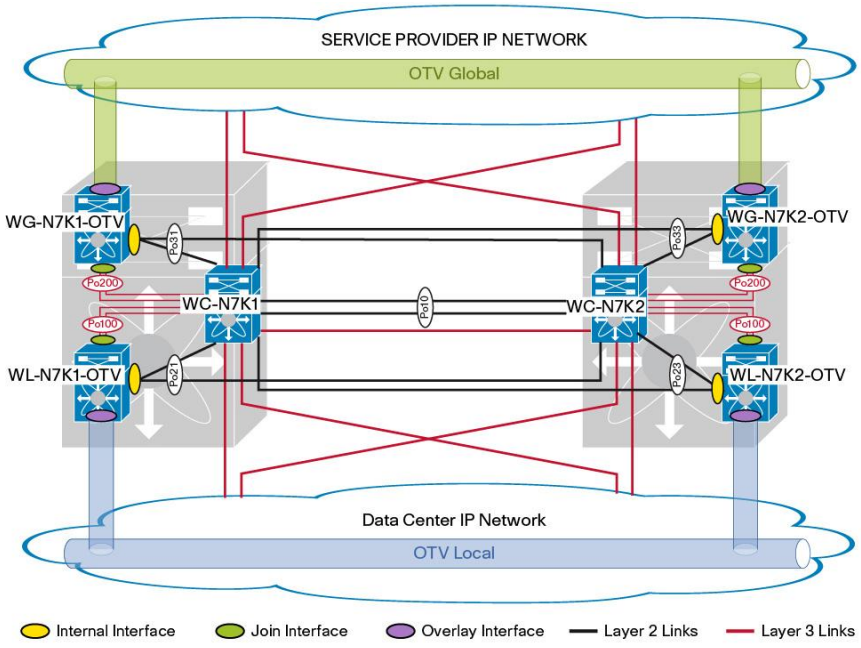

3.2.1 Layer 3 Design and Considerations

In the lab, Open Shortest Path First (OSPF) is used as the routing protocol of choice to provide Layer 3 connectivity and fast convergence both within and between the two data centers.

At the core layer, the two core VDCs (WC-N7K1 and WC-N7K2) are configured in such a way that OSPF process 1 is running on the following:

- Layer 3 link between WC-N7K1 and WC-N7K2
- Layer 3 links connecting to Data Center East

- Layer 3 PortChannels connecting to each OTV VDC
- Layer 3 links connecting to each Layer 2 domain
- Interface loopback0

On those OTV VDCs, OSPF process1 is running on the OTV join interfaces. To provide redundancy, you create two Layer 3 links between the OTV VDC and core VDC to form a Layer 3 PortChannel.

3.2.2 Layer 2 Design and Considerations

At the core layer, vPC technology is used to connect the OTV VDCs to two core VDCs, providing redundancy, higher bandwidth, and active-active bidirectional connectivity. The OTV VDCs see the vPC as a regular PortChannel and are unaware that it is connected to two separate core switches.

You create additional Layer 2 links between OTV VDCs and core VDCs to provide a Layer 2 path between core VDCs and OTV VDCs. On the two core routers, you form the vPC peer and let all the OTV VDCs have a vPC connection to this vPC pair. The vPC configuration on the core layer is a typical vPC design.

The Layer 2 links form PortChannel 10 between WC-N7K1 and WC-N7K2, the 2 VDCs act as a single virtual switch with separate control planes. Po10 is configured as a vPC peer link between the two core VDCs, providing vPC services to OTV VDCs (that is, WG-N7K1-OTV and WL-N7K1-OTV). Po10 carries Layer 2 vPC traffic between the two core VDCs. You use mgmt0 for the vPC keepalive link.

3.2.3 OTV Design and Considerations

OTV is a MAC-in-IP technique for supporting Layer 2 VPNs and extending LANs over any transport so long as the transport can carry IP packets.

OTV can be thought of as a MAC address routing technology in which destinations are MAC addresses and next hops are IP addresses. OTV simply maps MAC address destinations to IP next hops that can be reached 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 its next-hop address. Because traffic is IP forwarded, OTV is as efficient as the transport network and delivers optimal traffic load balancing, multicast traffic replication, and fast failover identical to the transport.

The core principles on which OTV operates include the following:

- Use of a control protocol to advertise MAC address reachability information (instead of using data-plane learning based on flooding)
- Packet switching of IP encapsulated Layer 2 traffic (instead of using circuit switching) for data forwarding

At the core layer of each data center, you create two additional VDCs on each existing core device: one for OTV VDC joining OTV Local, and one for OTV VDC joining OTV Global. OTV Local forms adjacencies with multiple Layer 2 domains within the same data center. OTV Global forms adjacencies with remote data centers (refer to Figure 4).

3.2.3.1 Internal Interfaces

The OTV VDCs connect back to the core layer switches through Layer 2 and Layer 3. The Layer 2 interfaces 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.

Each OTV VDC has two Internal Interfaces. These are configured as a vPC trunk, with one link going to each core VDC.

3.2.3.2 Join Interfaces

The Layer 3 link on OTV VDC is known as a join interface. OTV uses this interface to perform IP-based virtualization to send and receive overlay traffic between sites. The IP address of this interface is used to advertise the reachability of 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 be 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.

3.2.3.3 Overlay Interfaces

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 address–based forwarding on the site, the overlay interface is simply another bridged interface. However, no spanning-tree packets or unknown unicast packets are forwarded over the overlay interface. From the perspective of the IP transport, the overlay interface is not visible.

OTV encapsulates packets into an IP header, in which it sets the Do Not Fragment (DF) bit for all OTV control and data packets crossing the transport network. The encapsulation adds 42 bytes to the original IP maximum transmission unit (MTU) size. It is a best practice to configure the join interface and all Layer 3 interfaces facing the IP core between the OTV edge devices with the maximum MTU size supported by the transport.

OTV uses site VLANs to allow multiple OTV edge devices within the site to talk to each other and determine the authoritative edge device (AED) for the OTV-extended VLANs. In this design, at the core layer you use VLAN 200 as the site VLAN for OTV Global, and VLAN 100 as the site VLAN for OTV Local. It is a best practice to use a dedicated VLAN as a site VLAN. The site VLAN should not be extended and should be carried down to the core VDCs across the VPC peer link. Any change to the site VLAN configuration **must** happen with the overlay in shutdown mode.

Starting with Cisco NX-OS 5.2(1), the configuration of the OTV site identifier is mandatory to make OTV work. This value should be identical for all OTV edge devices that belong to the same site, and a different site identifier needs to be used for a different site.

On the core layer, since WG-N7K-OTV and WL-N7k-OTV belong to two different overlays - OTV Global and OTV Local - you can use the same site identifier for both OTVs, or you can use different site identifiers. For the purposes of this document, the same site identifier is used for both OTV Global and OTV Local in the core layer within the same data center.

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

The multicast control group identifies the overlay. Two different overlays must have two different multicast control groups. In this design, you use 239.1.1.2 as the multicast control group for OTV Global, 239.1.1.1 as the multicast control group for OTV Local in Data Center West, and 239.1.1.3 as the multicast control group for OTV Local in Data Center East. The control group is used for neighbor discovery and to exchange MAC address reachability information. However, the data group is a source-specific multicast (SSM) group range, which is used to carry multicast data traffic generated by the sites.

After the overlay interface is configured with the control group and the **no shutdown** command is entered (make sure that the site VLAN configuration and site identifier are already in place), the OTV edge device sends an Internet Group Management Protocol (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. After the adjacencies are established, the OTV control-plane packets communicate MAC-to-IP address mappings (MAC address reachability information) 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.

3.3 Layer 2 Domain Aggregation Design

On the aggregation devices in each Layer 2 domain, you create one more VDC and dedicate it to the OTV edge device. These OTV VDCs join the OTV Local overlay as shown in Figure 5.

Figure 5. Data Center West Domain 1

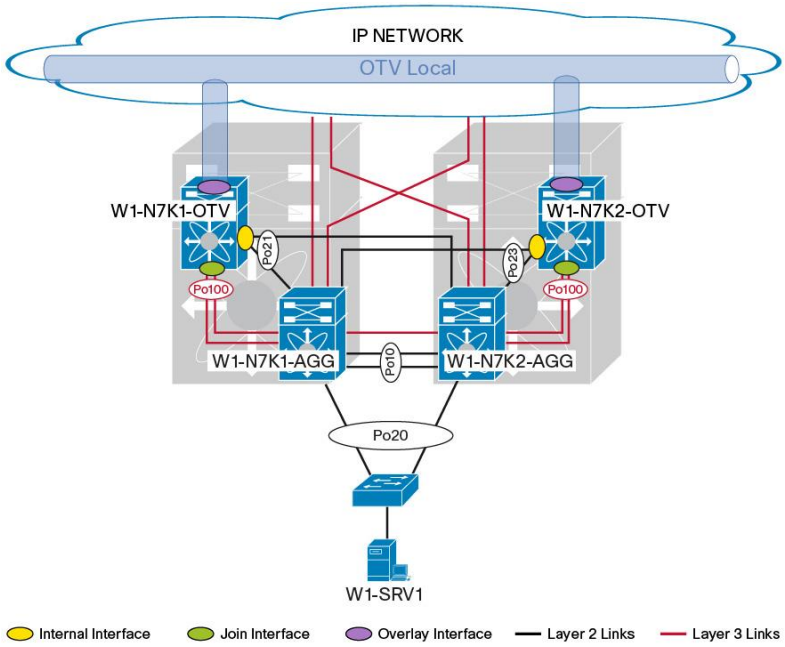

3.3.1 Layer 3 Design and Considerations

Within a data center, OSPF is used as the routing protocol of choice to provide Layer 3 connectivity and fast convergence among the Layer 2 domains. In each Layer 2 domain, the Cisco Nexus 7000 Series Switches are configured in such a way that OSPF Process 1 is running on all Layer 3 links, loopback interfaces, and switch virtual interfaces (SVIs) for the extended VLANs.

3.3.1.1 Hot Standby Router Protocol

First Hop Redundancy Protocol (FHRP) provides default gateway redundancy for connected hosts such as servers. This laboratory example uses Hot Standby Router Protocol (HSRP) as the FHRP.

Each VLAN in the aggregation block, which requires Layer 3 connectivity to the rest of the network, is configured with an HSRP gateway, including those VLANs extended on the overlay between the data centers. To allow the extended VLANs to use their local HSRP gateway, an IP gateway localization technique is used to prevent HSRP protocol data units (PDUs) from being forwarded on the overlay network.

This technique uses MAC address access control list (ACL) filtering of HSRP packets to block the propagation of HSRP packets between different Layer 2 domains. This approach prevents suboptimal outbound routing.

Inbound traffic optimization solutions are beyond the scope of this document. For more information, refer to the virtualized workload mobility design guide available at http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DCI/4.0/EMC/EMC.pdf.

3.3.2 Layer 2 Design and Considerations

The vPC configuration on the aggregation tier in each domain is a typical vPC design. The two aggregation-layer VDC formed a vPC pair, that acts as a single virtual switch with separate control planes.

Po10 is configured as a vPC peer link between the two VDCs, providing vPC services to access-layer devices and OTV VDCs. Po10 carries Layer 2 vPC traffic between the two switch VDCs. A vPC keepalive link is also configured between W1-N7K1-AGG and W1-N7K2-AGG for dual active detection.

Because traffic from the attached server can be hashed to any uplink in the bundle, vPC is also used between the aggregation layer and OTV VDCs as a means of providing a direct path for Layer 2 server traffic sent to, and received from, the overlay.

3.3.3 OTV Design and Considerations

This design use OTV edge device as a virtual appliance on a stick, OTV is configured under a separate VDC. This approach is used because the current implementation of OTV requires separation between SVIs 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 and Layer 3 links. The Layer 2 links are 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 vPC trunk with one link going to each aggregation-layer VDC.

The Layer 3 interface is known as a join interface. OTV uses this interface to perform IP-based virtualization to send and receive overlay traffic between Layer 2 domains and the core layer. The IP address of this interface is used to advertise the reachability of MAC addresses present in the local Layer 2 domain.

OTV uses the site VLAN to allow multiple OTV edge devices within the site to talk to each other and form site adjacency. Starting from Cisco NX-OS 5.2(1), both site adjacency and overlay adjacency are used to determine the AED for the OTV-extended VLANs. The site VLAN used in local Layer 2 domains is VLAN 100.

In each Layer 2 domain, you need to define a unique site identifier, and the OTV edge devices within the same Layer 2 domain should have the same site identifier.

The OTV edge device is also configured with the overlay interface 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. In Domain 1, interface PortChannel 100 is used as a join interface for OTV Local, interface PortChannel 21 is used as internal interface on W1-N7k1-OTV, and interface PortChannel 23 is used as internal interface on W1-N7K2-OTV.

In the Layer 2 domain, each OTV edge device is configured with common control group 239.1.1.1 and data group 239.3.2.0/28 on Overlay Interface 1 to join OTV Local.

4. Packet Walk Analysis

This section presents a Packet Walk for the following two traffic flows (Figure 6):

- Traffic flow between different Layer 2 domains within the same data center
- Traffic flow between different Layer 2 domains between different data centers

Figure 6. Intra- and Inter-Data Center Traffic Flows

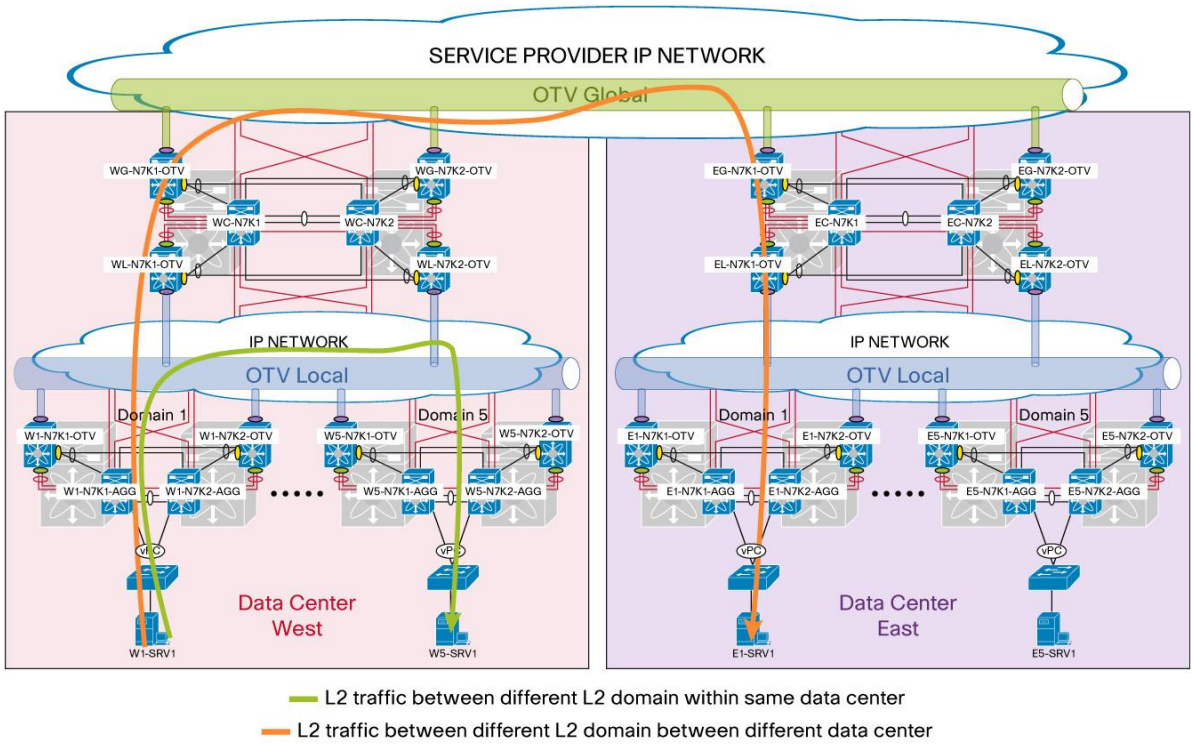

4.1 Intra-Data Center Interdomain Traffic Flow

This section examines the flow of the traffic between two Layer 2 domains within the same data center. The steps in Table 2 examine the path of the Layer 2 unicast traffic flow between two servers: one in Data Center West Domain 1 (W1-SRV1) and one in Data Center West Domain 5 (W5-SVR1).

The servers W1-SRV1 and W5-SRV1 are connected through vPC to their respective aggregation-tier switches on VLAN 10. For this flow, assume the following:

- Both servers reside in VLAN 10
- W1-N7K1-OTV and W5-N7K1-OTV are the AEDs for VLAN 10 in Data Center West Domain 1 (Figure 7) and Domain 5 (Figure 8)

Figure 7. Intra-Data Center Flow W1-SRV1 to W5-SRV1 in Domain 1

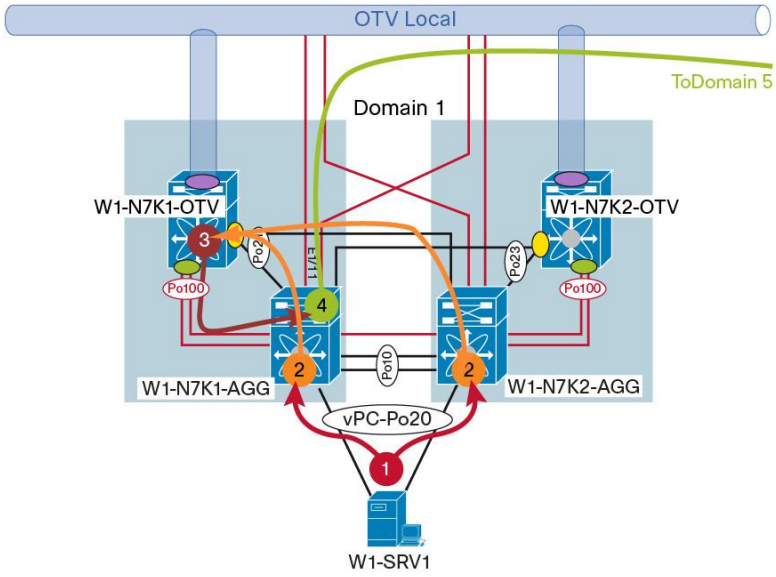

Table 2. Intra-Data Center Flow Path: W1-SRV1 to W5-SRV1

Step	Description
1	Layer 2 traffic travels to aggregation layer. The Layer 2 packet originating from W1-SRV1 travels over vPC link Po20 and arrives at either W1-N7K1-AGG or W1-N7K2-AGG. The hashing algorithm determines which of the two links the packet uses.
2	Layer 2 traffic travels to OTV VDC. From either W1-N7K1-AGG or W1-N7K2-AGG, the packet travels over the vPC link to the OTV AED for VLAN 10. Since AED for VLAN 10 in Domain 1 is W1-N7K1-OTV, the packet travels over vPC link Po21 to OTV VDC W1-N7K1-OTV.
3	Encapsulated Layer 3 packet is forwarded to aggregation layer. W1-N7K1-OTV performs a Layer 2 lookup on the destination MAC address of the frame. The MAC address table of W1-N7K1-OTV points to the IP address of the W5-N7k1-OTV join interface for this destination MAC address. At this point, W1-N7k1-OTV performs a MAC-in-IP encapsulation, where the destination IP address is the IP address of the OTV join interface of the Domain 5 OTV VDC, which is the AED for VLAN 10 in Domain 5. On the basis of its routing table, W1-N7k1-OTV forwards the packet over the OTV join interface Po100 connected to W1-N7k1-AGG.
4	Layer 3 packet leaves Domain 1. The Domain 1 aggregation-layer VDC, W1-N7K1-AGG, upon receiving the encapsulated packet, performs an IP destination lookup and forwards the packet using OSPF over the e1/11 interface to the aggregation-layer VDC in Domain 5, W5-N7K1-AGG.

Figure 8. Intra-Data Center Flow W1-SRV1 to W5-SRV1 in Domain 5

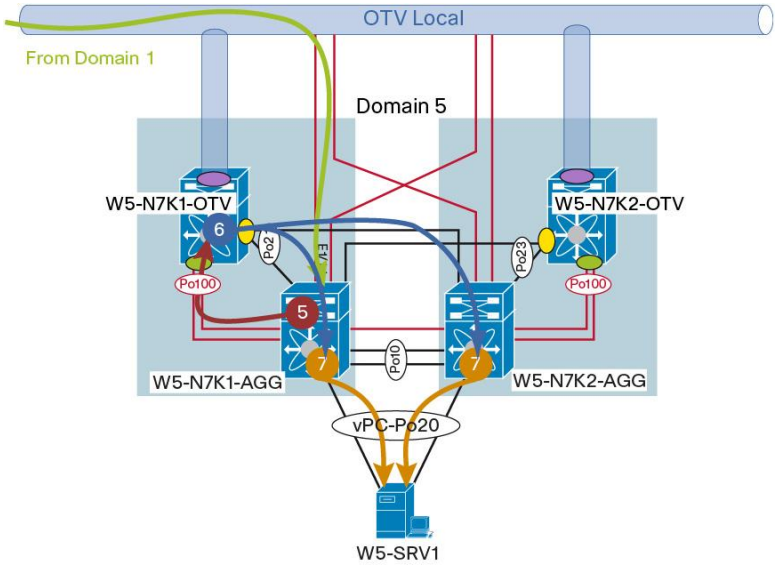

Step	Description
5	Layer 3 traffic forwarded to OTV vDC in Domain 5. Domain 5 aggregation-layer VDC W5-N7K1-AGG upon receiving the packet on e1/15 forwards it over the Layer 3 PortChannel to OTV join interface Po100 on W5-N7k1-OTV.
6	Layer 3 packet is decapsulated and sent to aggregation layer. After it arrives in the OTV VDC, the IP packet is decapsulated to a Layer 2 packet. The packet is then forwarded toward its Layer 2 destination based on the MAC address table on W5-N7K1-OTV. In this case, the packet is sent over the internal interface, vPC Po21, and can use either link in the PortChannel.
7	Layer 2 packet is forwarded to destination. The Layer 2 packet arrives either at the W5-N7k1-AGG or W5-N7K2-AGG VDC depending on the vPC hash. Either VDC then forwards the packet over Po20 to the destination server, W5-SRV1.

4.2 Inter-Data Center Traffic Flow

This section examines the flow of traffic between two data centers to gain an understanding of the design functions. The steps in Table 3 examine the path of the Layer 2 unicast traffic flow between two servers: one in Data Center West Domain 1 (W1-SRV1) and one in Data Center East Domain 1 (E1-SRV1).

The servers W1-SRV1 and E1-SRV1 are connected through vPC to their respective aggregation-tier switches on VLAN 10. For this flow, assume the following:

- Both servers reside in VLAN 10
- W1-N7K1-OTV and E1-N7K1-OTV are the AEDs for VLAN 10 in Data Center West Domain 1 and Data Center East Domain 1
- WL-N7K1-OTV and WG-N7K1-OTV are the AEDs for VLAN10 in the Data Center West Core OTV Local overlay and OTV Global overlay (Figure 9)
- EL-N7K1-OTV and EG-N7K1-OTV are the AEDs for VLAN10 in the Data Center East Core OTV Local overlay and OTV Global overlay (Figure 10)

Figure 9. Inter-Data Center Flow: W1-SRV1 to E1-SRV1 in Data Center West

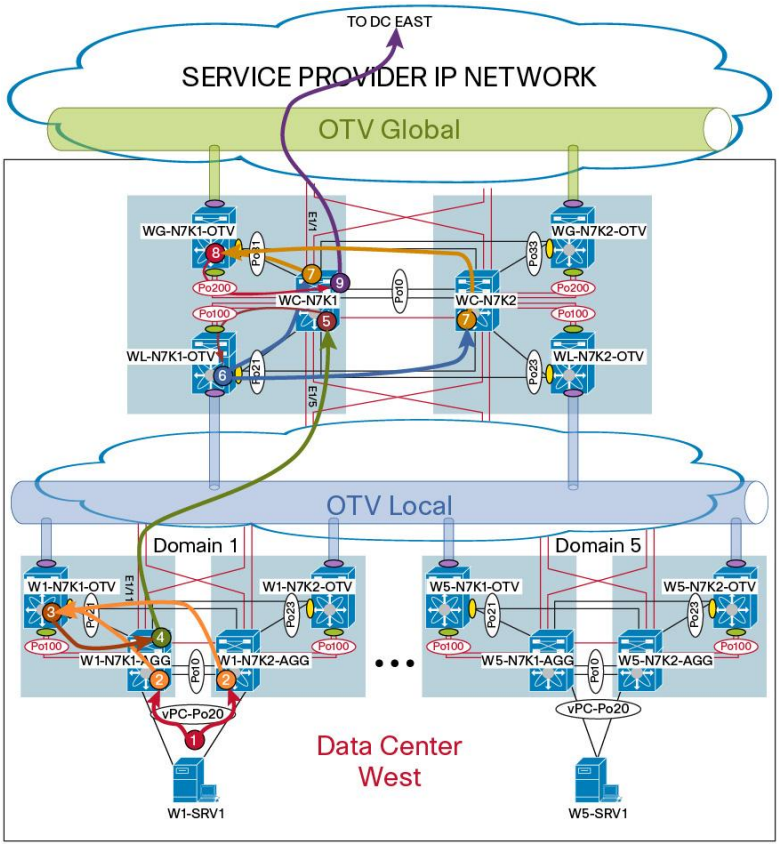

Table 3. Inter-Data Center Flow Path: W1-SRV1 to E1-SRV1

Step	Description
1	Layer 2 traffic travels to aggregation layer. The Layer 2 packet originating from W1-srv1 travels over vPC link Po20 and arrives at either W1-N7K1-AGG or W1-N7K2-AGG. The hashing algorithm determines which of the two links the packet uses.
2	Layer 2 traffic travels to OTV VDC. From either W1-N7K1-AGG or W1-N7K2-AGG, the packet travels over the vPC link to the OTV AED for VLAN 10. Since the AED for VLAN 10 in Domain 1 is W1-N7K1-OTV, the packet travels over vPC link Po21 to OTV VDC W1-N7K1-OTV.
3	Encapsulated Layer 3 packet is forwarded to aggregation layer. W1-N7K1-OTV performs a Layer 2 lookup on the destination MAC address of the frame. The MAC address table of W1-N7K1-OTV points to the IP address of the WL-N7k1-OTV join interface for this destination MAC address. At this point, W1-N7k1-OTV performs a MAC-in-IP encapsulation, where the destination IP address is the IP address of the OTV join interface of WL-N7K1-OTV, which is the AED for VLAN 10 on the core OTV edge device that joins the OTV Local overlay. On the basis of its routing table, W1-N7k1-OTV forwards the packet over OTV join interface Po100 connect to W1-N7k1-AGG.
4	Layer 3 packet leaves Domain 1. The Domain 1 aggregation-layer VDC, W1-N7K1-AGG, upon receiving the encapsulated packet, performs an IP destination lookup and forwards the packet using e1/11 interface to the core VDC, WC-N7K1.
5	Layer 3 traffic is forwarded to OTV vDC in the core. The core VDC WC-N7K1 upon receiving the packet on e1/5 forwards it over the Layer 3 PortChannel to OTV join interface po100 on WL-N7k1-OTV.
6	Layer 3 packet is decapsulated and sent to the core. After it arrives in the OTV VDC, the IP packet is decapsulated to a Layer 2 packet. The packet is then forwarded based on the MAC address table on WL-N7K1-OTV. In this case, the packet is sent over the internal interface, vPC Po21, and can use either link in the PortChannel.

Step	Description
7	Layer 2 packet is forwarded to the OTV edge device for the OTV Global overlay. The Layer 2 packet arrives at either the WC-N7k1 or WC-N7K2 VDC, depending on the vPC hash. Either VDC then forwards the packet to WG-N7K1-OTV over Po31.
8	Encapsulated Layer 3 packet is forwarded to the core. WG-N7K1-OTV performs a Layer 2 lookup on the destination MAC address of the frame. The MAC address table of WG-N7K1-OTV points to the IP address of the EG-N7k1-OTV join interface for this destination MAC address. At this point, WG-N7k1-OTV performs a MAC-in-IP encapsulation, where the destination IP address is the IP address of the OTV join interface of EG-N7K1-OTV, which is the AED for VLAN 10 on the Data Center East core OTV VDC that joins the OTV Global overlay. On the basis of its routing table, WG-N7k1-OTV forwards the packet over OTV join interface po200 connecting to WC-N7k1.
9	Layer 3 packet leaves Data Center West. WC-N7K1, upon receiving the encapsulated packet, performs an IP destination lookup and forwards the packet using OSPF over the e1/1 interface to Data Center East Core VDC EC-N7K1.

Figure 10. Inter-Data Center Flow: W1-SRV1 to E1-SRV1 in Data Center East

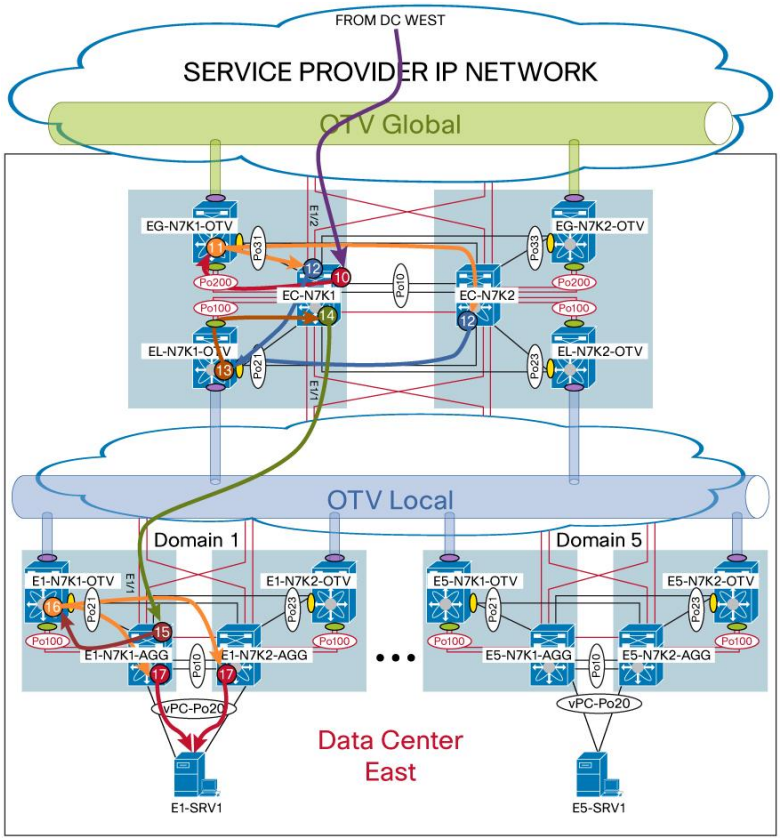

Step	Description
10	Layer 3 traffic is forwarded to OTV vDC in the Data Center East core. In Data Center East, core VDC EC-N7K1, upon receiving the packet on e1/2, forwards it over the Layer 3 PortChannel to OTV Global join interface po200 on EG-N7k1-OTV.
11	Layer 3 packet is decapsulated and sent to the core. After arriving in Global OTV VDC EG-N1K1-OTV, the IP packet is decapsulated to a Layer 2 packet. The packet is then forwarded based on the MAC address table on EG-N7K1-OTV. In this case, the packet is sent over the internal interface, vPC Po31, and can use either link in the PortChannel.
12	Layer 2 packet is forwarded to the OTV edge device for the OTV Local overlay. The Layer 2 packet arrives at either EC-N7k1 or EC-N7K2 VDC, depending on the vPC hash. Either VDC then directly forwards the packet to EL-N7K1-OTV over Po21.

Step	Description
13	Encapsulated Layer 3 packet is forwarded to the core. EL-N7K1-OTV performs a Layer 2 lookup on the destination MAC address of the frame. The MAC address table of EL-N7K1-OTV points to the IP address of the E1-N7k1-OTV join interface for this destination MAC address. At this point, EL-N7k1-OTV performs a MAC-in-IP encapsulation, where the destination IP address is the IP address of the OTV join interface of E1-N7K1-OTV, which is the AED for VLAN 10 on Data Center East Domain 1. On the basis of its routing table, EL-N7k1-OTV forwards the packet over OTV join interface po100 connecting to EC-N7k1.
14	Layer 3 packet leaves the core. EC-N7K1, upon receiving the encapsulated packet, performs an IP destination lookup and forwards the packet using OSPF over the e1/1 interface to the Data Center East Domain 1 aggregation-layer VDC, E1-N7K1-AGG.
15	Layer 3 traffic is forwarded to OTV vDC in DC East Domain 1. Domain 1 aggregation-layer VDC E1-N7K1-AGG, upon receiving the packet on e1/1, forwards the packet over the Layer 3 PortChannel to OTV join interface po100 on E1-N7k1-OTV.
16	Layer 3 packet is decapsulated and sent to aggregation layer. After it arrives in the OTV VDC, the IP packet is decapsulated to a Layer 2 packet. The packet is then forwarded to its Layer 2 destination based on the MAC address table. In this case, the packet is sent over the internal interface, vPC Po21, and can use either link in the PortChannel.
17	Layer 2 packet is forwarded to destination. The Layer 2 packet arrives at either the E1-N7k1-AGG or E1-N7K2-AGG VDC, depending on the vPC hash. Either VDC then directly forwards the packet over Po20 to the destination server, E1-SRV1.

5. Setup, Configuration, and Validation

This section discusses setup, configuration, and validation on the core and aggregation layers in each Layer 2 domain. Data Center West is used as an example.

The example in this document deploys OTV over a multicast-enabled transport. One of the main advantages of using multicast is that it allows optimal multicast traffic replication to multiple sites and avoids head-end replication, which leads to suboptimal bandwidth utilization.

For unicast-only transport configuration, please refer to the OTV Technology Introduction and Deployment Considerations Guide at the following link: http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DCI/whitepaper/DCI_1.html.

5.1 Core Layer Setup, Configuration, and Validation

On each of the core switches, WC-N7K1 and WC-N7K2, two additional VDCs are created, used as OTV VDCs:

- WG-N7K1-OTV and WG-N7K2-OTV are used to join OTV Global, forming adjacencies with remote data centers
- WL-N7K1-OTV and WL-N7K2-OTV are used to join OTV Local, forming adjacencies with Domains 1 through 5 within the same data center

Figure 11. Data Center West Core Layer

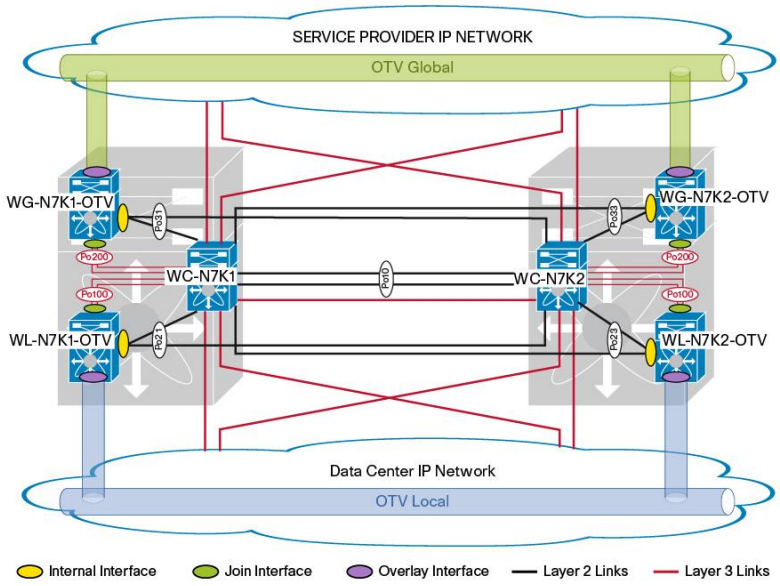

5.1.1 Core Layer Setup and Configuration

Layer 3 Configuration

WC-N7K1#	**WC-N7K2#**
feature ospf	feature ospf

router ospf 1
  log-adjacency-changes
	router ospf 1
	  log-adjacency-changes

interface port-channel100
  description L3 port-channel to WL-N7K1-OTV
  mtu 9216
  ip address 10.4.2.1/24
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode
	interface port-channel100
	  description L3 port-channel to WL-N7K2-OTV
	  mtu 9216
	  ip address 10.3.2.1/24
	  ip ospf network point-to-point
	  ip router ospf 1 area 0.0.0.0
	  ip pim sparse-mode

interface port-channel200
description L3 port-channel to WG-N7K1-OTV
  mtu 9216
  ip address 10.4.3.1/24
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode
	interface port-channel200
	  description L3 port-channel to WG-N7K2-OTV
	  mtu 9216
	  ip address 10.3.3.1/24
	  ip ospf network point-to-point
	  ip router ospf 1 area 0.0.0.0
	  ip pim sparse-mode

WL-N7K1-OTV#

```
feature ospf

router ospf 1
  log-adjacency-changes

interface port-channel100
  description L3 port-channel connect
to WC-N7K1
  mtu 9216
  ip address 10.4.2.2/24
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip igmp version 3
```

WG-N7K1-OTV#

```
feature ospf

router ospf 1
  log-adjacency-changes

interface port-channel200
description L3 port-channel to WC-N7K1
  mtu 9216
  ip address 10.4.3.2/24
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip igmp version 3
```

WL-N7K2-OTV#

```
feature ospf

router ospf 1
  log-adjacency-changes

interface port-channel100
  description L3 port-channel connect
to WC-N7K2
  mtu 9216
  ip address 10.3.2.2/24
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip igmp version 3
```

WG-N7K1-OTV#

```
feature ospf

router ospf 1
  log-adjacency-changes

interface port-channel200
description L3 port-channel to WC-N7K2
  mtu 9216
  ip address 10.3.3.2/24
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip igmp version 3
```

Layer 2 Configuration

WC-N7K1#

```
feature vpc

vpc domain 1
  role priority 100
  peer-keepalive destination
172.21.55.33 source 172.21.55.34

interface port-channel10
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100,200
  spanning-tree port type network
```

WC-N7K1#

```
feature vpc

vpc domain 1
  role priority 200
  peer-keepalive destination
172.21.55.34 source 172.21.55.33

interface port-channel10
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100,200
  spanning-tree port type network
```

```
  vpc peer-link

interface port-channel21
  description vpc to WL-N7K1-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
  vpc 21

interface port-channel23
  description vpc to WL-N7K2-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
  vpc 23

interface port-channel31
  description vpc to WG-N7K1-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,200
  vpc 31

interface port-channel33
  description vpc to WG-N7K2-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,200
  vpc 33
```

WL-N7K1-OTV#

```
interface port-channel21
  description vpc 21
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
```

WG-N7K1-OTV#

```
interface port-channel31
  description vpc 31
```

```
  vpc peer-link

interface port-channel21
  description vpc to WL-N7K1-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
  vpc 21

interface port-channel23
  description vpc to WL-N7K2-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
  vpc 23

interface port-channel31
  description vpc to WG-N7K1-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,200
  vpc 31

interface port-channel33
  description vpc to WG-N7K2-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,200
  vpc 33
```

WL-N7K2-OTV#

```
interface port-channel23
  description vpc 23
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
```

WG-N7K2-OTV#

```
interface port-channel33
  description vpc 33
```

```
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,200
```

```
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,200
```

OTV Configuration

```
WL-N7K1-OTV#
feature otv

otv site-vlan 100

interface Overlay1
  description OTV Local
  otv join-interface port-channel100
  otv control-group 239.1.1.1
  otv data-group 239.2.1.0/28
  otv extend-vlan 10-19
  no shutdown
otv site-identifier 3434.3434.3434

WG-N7K1-OTV#
feature otv

otv site-vlan 200

interface Overlay1
  description OTV Global
  otv join-interface port-channel200
  otv control-group 239.1.1.2
  otv data-group 239.2.2.0/28
  otv extend-vlan 10-19
  no shutdown
otv site-identifier 3434.3434.3434
```

```
WL-N7K2-OTV#
feature otv

otv site-vlan 100

interface Overlay1
  description OTV Local
  otv join-interface port-channel100
  otv control-group 239.1.1.1
  otv data-group 239.2.1.0/28
  otv extend-vlan 10-19
  no shutdown
otv site-identifier 3434.3434.3434

WG-N7K2-OTV#
feature otv

otv site-vlan 200

interface Overlay1
  description OTV Global
  otv join-interface port-channel200
  otv control-group 239.1.1.2
  otv data-group 239.2.2.0/28
  otv extend-vlan 10-19
  no shutdown
otv site-identifier 3434.3434.3434
```

5.1.2 Core Layer Validation

In this section, the test cases listed in Table 4 are run to validate the status of the core layer setup.

Table 4. Test Cases to Validate Status of Core Layer Setup

Section	Case Index	Test Case	Command
5.1.2.1 Check OTV Local Status	5.1.2.1.1	Check overly interface status	**Show otv overlay 1**
	5.1.2.1.2	Check OTV adjacency	**Show otv adjacency**
	5.1.2.1.3	Check OTV load balance	**Show otv vlan**
	5.1.2.1.4	Connectivity test	**ping**
5.1.2.2 Check OTV Global Status	5.1.2.2.1	Check overly interface status	**Show otv overlay 1**
	5.1.2.2.2	Check OTV adjacency	**Show otv adjacency**
	5.1.2.2.3	Check OTV load balance	**Show otv vlan**
	5.1.2.2.4	connectivity test	**ping**

5.1.2.1 Check OTV Local Status

The tests in this section check the OTV Local status; determine whether the interface overlay is up, adjacency has been formed, and the AED is functional; and ping the extended VLANs.

5.1.2.1.1 Check Overly Interface Status

This test checks the overlay interface status on WL-N7K1-OTV and WL-N7K2-OTV to make sure that the overlay interface is up and that the status is displayed correctly.

```
WL-N7K1-OTV# show otv overlay 1

OTV Overlay Information
Site Identifier 3434.3434.3434

Overlay interface Overlay1

 VPN name            : Overlay1
 VPN state           : UP
 Extended vlans      : 10-19 (Total:10)
 Control group       : 239.1.1.1
 Data group range(s) : 239.2.1.0/28
 Join interface(s)   : Po100 (10.4.2.2)
 Site vlan           : 100 (up)
 AED-Capable         : Yes
 Capability          : Multicast-Reachable

WL-N7K2-OTV# sh otv overlay 1

OTV Overlay Information
Site Identifier 3434.3434.3434
```

```
Overlay interface Overlay1

 VPN name             : Overlay1
 VPN state            : UP
 Extended vlans       : 10-19 (Total:10)
 Control group        : 239.1.1.1
 Data group range(s)  : 239.2.1.0/28
 Join interface(s)    : Po100 (10.3.2.2)
 Site vlan            : 100 (up)
 AED-Capable          : Yes
 Capability           : Multicast-Reachable
```

5.1.2.1.2 Check OTV Adjacency

This test checks whether all OTV adjacency neighbors are connected in OTV Local.

```
WL-N7K1-OTV# show otv adjacency

Overlay Adjacency database

Overlay-Interface Overlay1  :
Hostname                       System-ID      Dest Addr       Up Time   State
WL-N7K2-OTV                    0024.f718.99c2 10.3.2.2        01:21:17  UP
W1-N7K1-OTV                    f866.f20e.4342 10.5.1.2        01:21:14  UP
W1-N7K2-OTV                    f866.f20e.4343 10.5.2.2        01:21:14  UP
W5-N7K1-OTV                    0024.f718.9943 10.6.1.2        01:21:17  UP
W5-N7K2-OTV                    0024.f718.9944 10.6.2.2        01:21:15  UP

WL-N7K2-OTV# show otv adjacency

Overlay Adjacency database

Overlay-Interface Overlay1  :
Hostname                       System-ID      Dest Addr       Up Time   State
WL-N7K1-OTV                    0024.986f.3e43 10.4.2.2         01:25:49  UP
W1-N7K1-OTV                    f866.f20e.4342 10.5.1.2         1d08h     UP
W1-N7K2-OTV                    f866.f20e.4343 10.5.2.2         1d08h     UP
W5-N7K1-OTV                    0024.f718.9943 10.6.1.2         1d08h     UP
W5-N7K2-OTV                    0024.f718.9944 10.6.2.2         1d08h     UP
```

5.1.2.1.3 Check OTV Load Balance

This test checks whether all the VLAN AEDs are load-balanced between WL-N7K1-OTV and WL-N7K2-OTV.

```
WL-N7K1-OTV# show otv vlan

OTV Extended VLANs and Edge Device State Information (* - AED)

VLAN   Auth. Edge Device                    Vlan State       Overlay
----   -----------------------------------  ----------       -------
  10*  WL-N7K1-OTV                          active           Overlay1
  11   WL-N7K2-OTV                          inactive(Non AED)Overlay1
  12*  WL-N7K1-OTV                          active           Overlay1
  13   WL-N7K2-OTV                          inactive(Non AED)Overlay1
  14*  WL-N7K1-OTV                          active           Overlay1
  15   WL-N7K2-OTV                          inactive(Non AED)Overlay1
  16*  WL-N7K1-OTV                          active           Overlay1
  17   WL-N7K2-OTV                          inactive(Non AED)Overlay1
  18*  WL-N7K1-OTV                          active           Overlay1
  19   WL-N7K2-OTV                          inactive(Non AED)Overlay1

WL-N7K2-OTV# show otv vlan

OTV Extended VLANs and Edge Device State Information (* - AED)

VLAN   Auth. Edge Device                    Vlan State       Overlay
----   -----------------------------------  ----------       -------
  10   WL-N7K1-OTV                          inactive(Non AED)Overlay1
  11*  WL-N7K2-OTV                          active           Overlay1
  12   WL-N7K1-OTV                          inactive(Non AED)Overlay1
  13*  WL-N7K2-OTV                          active           Overlay1
  14   WL-N7K1-OTV                          inactive(Non AED)Overlay1
  15*  WL-N7K2-OTV                          active           Overlay1
  16   WL-N7K1-OTV                          inactive(Non AED)Overlay1
  17*  WL-N7K2-OTV                          active           Overlay1
  18   WL-N7K1-OTV                          inactive(Non AED)Overlay1
  19*  WL-N7K2-OTV                          active           Overlay1
```

5.1.2.1.4 Connectivity Test

This test need to configure the SVI for extended VLANs on WC-N7K1 and WC-N7K2 for verification. Here, VLAN 10 is used as an example, and VLAN 10 SVI on W1-N7K1-AGG and W1-N7K2-AGG is pinged from WC-N7K1 and WC-N7K2.

```
W1-N7K1-AGG#
interface Vlan10
```

```
  no shutdown
  ip address 150.0.10.51/24
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 1
    preempt
    priority 110
ip 150.0.10.100

W1-N7K2-AGG#
interface Vlan10
  no shutdown
  ip address 150.0.10.52/24
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 1
    preempt
    priority 90
    ip 150.0.10.100

WC-N7K1#
interface Vlan10
  no shutdown
  ip address 150.0.10.11/24

WC-N7K2#
interface Vlan10
  no shutdown
  ip address 150.0.10.12/24

WC-N7K1#ping 150.0.10.51
PING 150.0.10.51 (150.0.10.51): 56 data bytes
64 bytes from 150.0.10.51: icmp_seq=0 ttl=254 time=1.343 ms
64 bytes from 150.0.10.51: icmp_seq=1 ttl=254 time=0.98 ms
64 bytes from 150.0.10.51: icmp_seq=2 ttl=254 time=0.796 ms
64 bytes from 150.0.10.51: icmp_seq=3 ttl=254 time=0.848 ms
64 bytes from 150.0.10.51: icmp_seq=4 ttl=254 time=0.857 ms

--- 150.0.10.51 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.796/0.964/1.343 ms

WC-N7K1#ping 150.0.10.52
PING 150.0.10.52 (150.0.10.52): 56 data bytes
```

```
64 bytes from 150.0.10.52: icmp_seq=0 ttl=254 time=1.314 ms
64 bytes from 150.0.10.52: icmp_seq=1 ttl=254 time=0.964 ms
64 bytes from 150.0.10.52: icmp_seq=2 ttl=254 time=0.863 ms
64 bytes from 150.0.10.52: icmp_seq=3 ttl=254 time=0.89 ms
64 bytes from 150.0.10.52: icmp_seq=4 ttl=254 time=0.894 ms

--- 150.0.10.52 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.863/0.984/1.314 ms

WC-N7K2# ping 150.0.10.51
PING 150.0.10.51 (150.0.10.51): 56 data bytes
64 bytes from 150.0.10.51: icmp_seq=0 ttl=254 time=1.4 ms
64 bytes from 150.0.10.51: icmp_seq=1 ttl=254 time=0.949 ms
64 bytes from 150.0.10.51: icmp_seq=2 ttl=254 time=0.833 ms
64 bytes from 150.0.10.51: icmp_seq=3 ttl=254 time=0.851 ms
64 bytes from 150.0.10.51: icmp_seq=4 ttl=254 time=0.847 ms

--- 150.0.10.51 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.833/0.976/1.4 ms

WC-N7K2# ping 150.0.10.52
PING 150.0.10.52 (150.0.10.52): 56 data bytes
64 bytes from 150.0.10.52: icmp_seq=0 ttl=254 time=1.355 ms
64 bytes from 150.0.10.52: icmp_seq=1 ttl=254 time=0.939 ms
64 bytes from 150.0.10.52: icmp_seq=2 ttl=254 time=0.963 ms
64 bytes from 150.0.10.52: icmp_seq=3 ttl=254 time=0.849 ms
64 bytes from 150.0.10.52: icmp_seq=4 ttl=254 time=0.847 ms

--- 150.0.10.52 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.847/0.99/1.355 ms
```

5.1.2.2 Check OTV Global Status

5.1.2.2.1 Check Overly Interface Status

This test checks the overlay interface status on both WG-N7K1-OTV and WG-N7K2-OTV, to make sure that the overlay interface is up and that its status is displayed correctly.

```
WG-N7K1-OTV# show otv overlay 1

OTV Overlay Information
Site Identifier 3434.3434.3434
```

```
Overlay interface Overlay1

 VPN name            : Overlay1
 VPN state           : UP
 Extended vlans      : 10-19 (Total:10)
 Control group       : 239.1.1.2
 Data group range(s) : 239.2.2.0/28
 Join interface(s)   : Po200 (10.4.3.2)
 Site vlan           : 200 (up)
 AED-Capable         : Yes
 Capability          : Multicast-Reachable

WG-N7K2-OTV# show otv overlay 1

OTV Overlay Information
Site Identifier 3434.3434.3434

Overlay interface Overlay1

 VPN name            : Overlay1
 VPN state           : UP
 Extended vlans      : 10-19 (Total:10)
 Control group       : 239.1.1.2
 Data group range(s) : 239.2.2.0/28
 Join interface(s)   : Po200 (10.3.3.2)
 Site vlan           : 200 (up)
 AED-Capable         : Yes
 Capability          : Multicast-Reachable
```

5.1.2.2.2 Check OTV Adjacency

This test checks the OTV adjacency for OTV Global to make sure that all the neighbors are up.

```
WG-N7K1-OTV# show otv adjacency

Overlay Adjacency database

Overlay-Interface Overlay1  :
Hostname                       System-ID      Dest Addr       Up Time   State
EG-N7K1-OTV                    001b.54c2.1e43 10.0.200.2      17:48:10  UP
EG-N7K2-OTV                    001b.54c2.1e44 10.0.202.2      17:48:10  UP
WG-N7K2-OTV                    0024.f718.99c3 10.3.3.2        17:48:05  UP

WG-N7K2-OTV# show otv adjacency

Overlay Adjacency database
```

```
Overlay-Interface Overlay1  :
Hostname                        System-ID      Dest Addr       Up Time    State
EG-N7K1-OTV                     001b.54c2.1e43 10.0.200.2      1d22h      UP
EG-N7K2-OTV                     001b.54c2.1e44 10.0.202.2      1d22h      UP
WG-N7K1-OTV                     0024.986f.3e42 10.4.3.2        17:52:53   UP
```

5.1.2.2.3 Check OTV Load Balance

```
WG-N7K1-OTV# show otv vlan

OTV Extended VLANs and Edge Device State Information (* - AED)

VLAN    Auth. Edge Device                      Vlan State        Overlay
----    -----------------------------------    ----------        -------
  10*   WG-N7K1-OTV                            active            Overlay1
  11    WG-N7K2-OTV                            inactive(Non AED) Overlay1
  12*   WG-N7K1-OTV                            active            Overlay1
  13    WG-N7K2-OTV                            inactive(Non AED) Overlay1
  14*   WG-N7K1-OTV                            active            Overlay1
  15    WG-N7K2-OTV                            inactive(Non AED) Overlay1
  16*   WG-N7K1-OTV                            active            Overlay1
  17    WG-N7K2-OTV                            inactive(Non AED) Overlay1
  18*   WG-N7K1-OTV                            active            Overlay1
  19    WG-N7K2-OTV                            inactive(Non AED) Overlay1

WG-N7K2-OTV# show otv vlan

OTV Extended VLANs and Edge Device State Information (* - AED)

VLAN    Auth. Edge Device                      Vlan State        Overlay
----    -----------------------------------    ----------        -------
  10    WG-N7K1-OTV                            inactive(Non AED) Overlay1
  11*   WG-N7K2-OTV                            active            Overlay1
  12    WG-N7K1-OTV                            inactive(Non AED) Overlay1
  13*   WG-N7K2-OTV                            active            Overlay1
  14    WG-N7K1-OTV                            inactive(Non AED) Overlay1
  15*   WG-N7K2-OTV                            active            Overlay1
  16    WG-N7K1-OTV                            inactive(Non AED) Overlay1
  17*   WG-N7K2-OTV                            active            Overlay1
  18    WG-N7K1-OTV                            inactive(Non AED) Overlay1
  19*   WG-N7K2-OTV                            active            Overlay1
```

5.1.2.2.4 Connectivity Test

This test need to configure the SVI for extended VLANs on core routers in Data Center West and Data Center East to test connectivity. Here, VLAN 10 is used as an example, and the VLAN 10 SVI on EC-N7K1 and EC-N7K2 is pinged from WC-N7K1 and WC-N7K2.

```
EC-N7K1#
interface Vlan10
 no shutdown
 ip address 150.0.10.21/24

EC-N7K2#
interface Vlan10
 no shutdown
 ip address 150.0.10.22/24

WC-N7K1#
interface Vlan10
  no shutdown
  ip address 150.0.10.11/24

WC-N7K2#
interface Vlan10
  no shutdown
  ip address 150.0.10.12/24

WC-N7K1#ping 150.0.10.21
PING 150.0.10.21 (150.0.10.21): 56 data bytes
64 bytes from 150.0.10.21: icmp_seq=0 ttl=254 time=1.579 ms
64 bytes from 150.0.10.21: icmp_seq=1 ttl=254 time=9.114 ms
64 bytes from 150.0.10.21: icmp_seq=2 ttl=254 time=0.798 ms
64 bytes from 150.0.10.21: icmp_seq=3 ttl=254 time=0.763 ms
64 bytes from 150.0.10.21: icmp_seq=4 ttl=254 time=0.787 ms

--- 150.0.10.21 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.763/2.608/9.114 ms

WC-N7K1#ping 150.0.10.22
PING 150.0.10.22 (150.0.10.22): 56 data bytes
64 bytes from 150.0.10.22: icmp_seq=0 ttl=254 time=1.379 ms
64 bytes from 150.0.10.22: icmp_seq=1 ttl=254 time=0.899 ms
64 bytes from 150.0.10.22: icmp_seq=2 ttl=254 time=0.857 ms
64 bytes from 150.0.10.22: icmp_seq=3 ttl=254 time=0.851 ms
64 bytes from 150.0.10.22: icmp_seq=4 ttl=254 time=0.851 ms
```

```
--- 150.0.10.22 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.851/0.967/1.379 ms

WC-N7K2# ping 150.0.10.21
PING 150.0.10.21 (150.0.10.21): 56 data bytes
64 bytes from 150.0.10.21: icmp_seq=0 ttl=254 time=1.311 ms
64 bytes from 150.0.10.21: icmp_seq=1 ttl=254 time=0.784 ms
64 bytes from 150.0.10.21: icmp_seq=2 ttl=254 time=0.877 ms
64 bytes from 150.0.10.21: icmp_seq=3 ttl=254 time=0.883 ms
64 bytes from 150.0.10.21: icmp_seq=4 ttl=254 time=0.885 ms

--- 150.0.10.21 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.784/0.947/1.311 ms

WC-N7K2# ping 150.0.10.22
PING 150.0.10.22 (150.0.10.22): 56 data bytes
64 bytes from 150.0.10.22: icmp_seq=0 ttl=253 time=11.74 ms
64 bytes from 150.0.10.22: icmp_seq=1 ttl=253 time=0.918 ms
64 bytes from 150.0.10.22: icmp_seq=2 ttl=253 time=0.928 ms
64 bytes from 150.0.10.22: icmp_seq=3 ttl=253 time=0.947 ms
64 bytes from 150.0.10.22: icmp_seq=4 ttl=253 time=0.845 ms

--- 150.0.10.22 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.845/3.075/11.74 ms
```

5.2 Layer 2 Domain Setup, Configuration, and Validation

This section uses Domain 1 as an example and shows the Layer 3, Layer 2, and OTV configurations within each Layer 2 domain (Figure 12). The configurations in other domains are similar.

At the Domain 1 aggregation-layer device, an additional VDC dedicated to OTV is created to form a virtual appliance on a stick, and OTV is configured under this separated VDC.

This step is needed because the current implementation of OTV requires separation between the SVIs and the Layer 2 extension, which is now performed by the OTV VDC. In the current release, the SVIs cannot be defined after the device configures OTV to extend the same VLAN.

Figure 12. Domain 1 in Data Center West

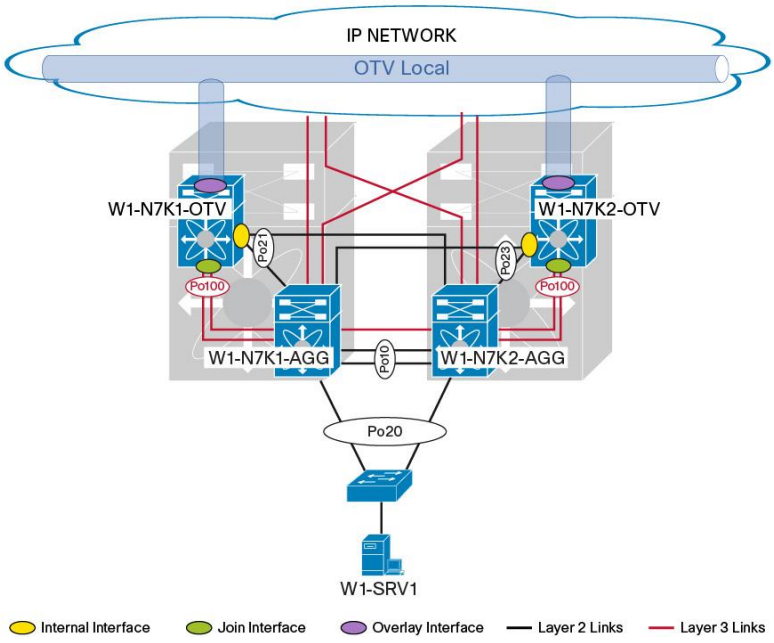

5.2.1 Layer 2 Domain Setup and Configuration

Within the data center, OSPF is used as the routing protocol of choice to provide Layer 3 connectivity and fast convergence among Layer 2 domains. In each Layer 2 domain, the Cisco Nexus 7000 Series Switches are configured in such a way that OSPF Process 1 is running on all Layer 3 links, loopback interfaces, and SVIs for the extended VLANs.

Layer 3 Configuration

```
W1-N7K1-AGG#
feature ospf

router ospf 1
  log-adjacency-changes

interface port-channel100
  description L3 port-channel to
W1-N7K1-OTV
  mtu 9216
  ip address 10.5.1.1/24
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode

interface Vlan10
  no shutdown
  ip address 150.0.10.51/24
```

```
W1-N7K2-AGG#
feature ospf

router ospf 1
  log-adjacency-changes

interface port-channel100
  description L3 port-channel to
W1-N7K2-OTV
  mtu 9216
  ip address 10.5.2.1/24
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip pim sparse-mode

interface Vlan10
  no shutdown
  ip address 150.0.10.52/24
```

```
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 1
    preempt
    priority 110
    ip 150.0.10.100

W1-N7K1-OTV#
feature ospf

router ospf 1
  log-adjacency-changes

interface port-channel100
  description L3 port-channel connect
to W1-N7K1-AGG
  mtu 9216
ip port access-group deny-hsrp in
  ip address 10.5.1.2/24
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip igmp version 3
```

```
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 1
    preempt
    priority 90
    ip 150.0.10.100

W1-N7K2-OTV#
feature ospf

router ospf 1
  log-adjacency-changes

interface port-channel100
  description L3 port-channel connect
to W1-N7K2-AGG
  mtu 9216
ip port access-group deny-hsrp in
  ip address 10.5.2.2/24
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip igmp version 3
```

HSRP Isolation

```
W1-N7K1-OTV#

ip access-list ALL_IPs
10 permit ip any any
!
mac access-list ALL_MACs
10 permit 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
   !
mac access-list HSRP_VMAC
10 permit 0000.0c07.ac00 0000.0000.00ff
any
20 permit 0000.0c9f.f000 0000.0000.0fff
any
!
arp access-list HSRP_VMAC_ARP
10 deny ip any mac 0000.0c07.ac00
```

```
W1-N7K2-OTV#

ip access-list ALL_IPs
10 permit ip any any
!
mac access-list ALL_MACs
10 permit 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
   !
mac access-list HSRP_VMAC
10 permit 0000.0c07.ac00 0000.0000.00ff
any
20 permit 0000.0c9f.f000 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 permit ip any mac any
!
vlan access-map HSRP_Localization 10
   match mac address HSRP_VMAC
match ip address HSRP_IP
action drop
vlan access-map HSRP_Localization 20
match mac address ALL_MACs
match ip address ALL_IPs
action forward
!
feature dhcp
ip arp inspection filter HSRP_VMAC_ARP 10-19
vlan filter HSRP_Localization vlan-list 10-19
!
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 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
```

```
ffff.ffff.ff00
20 deny ip any mac 0000.0c9f.f000 ffff.ffff.f000
   30 permit ip any mac any
!
vlan access-map HSRP_Localization 10
   match mac address HSRP_VMAC
match ip address HSRP_IP
action drop
vlan access-map HSRP_Localization 20
match mac address ALL_MACs
match ip address ALL_IPs
action forward
!
feature dhcp
ip arp inspection filter HSRP_VMAC_ARP 10-19
vlan filter HSRP_Localization vlan-list 10-19
!
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 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
```

Layer 2 Configuration

```
W1-N7K1#
feature vpc

vpc domain 1
  role priority 100
  peer-keepalive destination 30.0.0.2 source 30.0.0.1 vrf keepalive

interface port-channel10
  switchport
```

```
W1-N7K2#
feature vpc

vpc domain 1
  role priority 200
  peer-keepalive destination 30.0.0.1 source 30.0.0.2 vrf keepalive

interface port-channel10
  switchport
```

```
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
  spanning-tree port type network
  vpc peer-link

interface port-channel20
  description vpc to host
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10-19
  vpc 20

interface port-channel21
  description vpc to W1-N7K1-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
  vpc 21

interface port-channel23
  description vpc to W1-N7K2-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
  vpc 23

interface Ethernet3/15
  vrf member keepalive
  ip address 30.0.0.1/30
  no shutdown

W1-N7K1-OTV#
interface port-channel21
  description vpc 21
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
```

```
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
  spanning-tree port type network
  vpc peer-link

interface port-channel20
  description vpc to host
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10-19
  vpc 20

interface port-channel21
  description vpc to W1-N7K1-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
  vpc 21

interface port-channel23
  description vpc to W1-N7K2-OTV
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
  vpc 23

interface Ethernet3/25
  vrf member keepalive
  ip address 30.0.0.2/30
  no shutdown

W1-N7K2-OTV#
interface port-channel23
  description vpc 23
  switchport
  switchport mode trunk
  switchport trunk allowed vlan
10-19,100
```

OTV Configuration

```
W1-N7K1-OTV#
feature otv

otv site-vlan 100

interface Overlay1
  otv join-interface port-channel100
  otv control-group 239.1.1.1
  otv data-group 239.2.1.0/28
  otv extend-vlan 10-19
  no shutdown
otv site-identifier 1111.1111.1111
```

```
W1-N7K2-OTV#
feature otv

otv site-vlan 100

interface Overlay1
  otv join-interface port-channel100
  otv control-group 239.1.1.1
  otv data-group 239.2.1.0/28
  otv extend-vlan 10-19
  no shutdown
otv site-identifier 1111.1111.1111
```

5.2.2 Layer 2 Domain Validation

In this section, the test cases listed in Table 5 are run to validate the status of the Layer 2 domain.

Table 5. Layer 2 Domain Validation

Case Index	Test Case	Command
5.2.2.1	Check overlay interface status	**Show otv overlay 1**
5.2.2.2	Check OTV adjacency	**Show otv adjacency**
5.2.2.3	Check OTV load balance	**Show otv vlan**
5.2.2.4	Connectivity test	**ping**
5.2.2.5	Check HSRP isolation	**show hsrp interface vlan 10**

5.2.2.1 Check Overly Interface Status

This test checks the overlay interface status on W1-N7K1-OTV and W1-N7K2-OTV to make sure that the overlay interface is up and that its status is displayed correctly.

```
W1-N7K1-OTV# show otv overlay 1

OTV Overlay Information
Site Identifier 1111.1111.1111

Overlay interface Overlay1

 VPN name            : Overlay1
 VPN state           : UP
 Extended vlans      : 10-19 (Total:10)
 Control group       : 239.1.1.1
 Data group range(s) : 239.2.1.0/28
 Join interface(s)   : Po100 (10.5.1.2)
 Site vlan           : 100 (up)
 AED-Capable         : Yes
```

```
 Capability          : Multicast-Reachable

W1-N7K2-OTV# show otv overlay 1

OTV Overlay Information
Site Identifier 1111.1111.1111

Overlay interface Overlay1

 VPN name            : Overlay1
 VPN state           : UP
 Extended vlans      : 10-19 (Total:10)
 Control group       : 239.1.1.1
 Data group range(s) : 239.2.1.0/28
 Join interface(s)   : Po100 (10.5.2.2)
 Site vlan           : 100 (up)
 AED-Capable         : Yes
 Capability          : Multicast-Reachable
```

5.2.2.2 Check OTV Adjacency

This test checks the OTV adjacency for OTV Local and makes sure that all the neighbors are up.

```
W1-N7K1-OTV# show otv adjacency
Overlay Adjacency database

Overlay-Interface Overlay1  :
Hostname                       System-ID      Dest Addr      Up Time    State
WL-N7K2-OTV                    0024.f718.99c2 10.3.2.2        04:16:05  UP
WL-N7K1-OTV                    0024.986f.3e43 10.4.2.2        04:16:04  UP
W1-N7K2-OTV                    f866.f20e.4343 10.5.2.2        04:16:05  UP
W5-N7K1-OTV                    0024.f718.9943 10.6.1.2        04:16:08  UP
W5-N7K2-OTV                    0024.f718.9944 10.6.2.2        04:16:07  UP

W1-N7K2-OTV# show otv adjacency

Overlay Adjacency database

Overlay-Interface Overlay1  :
Hostname                       System-ID      Dest Addr      Up Time    State
WL-N7K2-OTV                    0024.f718.99c2 10.3.2.2        2d03h     UP
WL-N7K1-OTV                    0024.986f.3e43 10.4.2.2        19:49:09  UP
W1-N7K1-OTV                    f866.f20e.4342 10.5.1.2        04:18:22  UP
W5-N7K1-OTV                    0024.f718.9943 10.6.1.2        1w0d      UP
W5-N7K2-OTV                    0024.f718.9944 10.6.2.2        1w0d      UP
```

5.2.2.3 Check OTV Load Balance

```
W1-N7K1-OTV# show otv vlan

OTV Extended VLANs and Edge Device State Information (* - AED)

VLAN   Auth. Edge Device                      Vlan State       Overlay
----   -------------------------------------  ----------       -------
  10*  W1-N7K1-OTV                            active           Overlay1
  11   W1-N7K2-OTV                            inactive(Non AED)Overlay1
  12*  W1-N7K1-OTV                            active           Overlay1
  13   W1-N7K2-OTV                            inactive(Non AED)Overlay1
  14*  W1-N7K1-OTV                            active           Overlay1
  15   W1-N7K2-OTV                            inactive(Non AED)Overlay1
  16*  W1-N7K1-OTV                            active           Overlay1
  17   W1-N7K2-OTV                            inactive(Non AED)Overlay1
  18*  W1-N7K1-OTV                            active           Overlay1
  19   W1-N7K2-OTV                            inactive(Non AED)Overlay1

W1-N7K2-OTV# show otv vlan

OTV Extended VLANs and Edge Device State Information (* - AED)

VLAN   Auth. Edge Device                      Vlan State       Overlay
----   -------------------------------------  ----------       -------
  10   W1-N7K1-OTV                            inactive(Non AED)Overlay1
  11*  W1-N7K2-OTV                            active           Overlay1
  12   W1-N7K1-OTV                            inactive(Non AED)Overlay1
  13*  W1-N7K2-OTV                            active           Overlay1
  14   W1-N7K1-OTV                            inactive(Non AED)Overlay1
  15*  W1-N7K2-OTV                            active           Overlay1
  16   W1-N7K1-OTV                            inactive(Non AED)Overlay1
  17*  W1-N7K2-OTV                            active           Overlay1
  18   W1-N7K1-OTV                            inactive(Non AED)Overlay1
  19*  W1-N7K2-OTV                            active           Overlay1
```

5.2.2.4 Connectivity Test

This test uses VLAN 10 as an example and pings the VLAN 10 SVI on WC-N7K1 and WC-N7K2 from W1-N7K1-AGG and W1-N7K2-AGG.

```
W1-N7K1-AGG#
interface Vlan10
  no shutdown
  ip address 150.0.10.51/24
```

```
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 1
    preempt
    priority 110
ip 150.0.10.100

W1-N7K2-AGG#
interface Vlan10
  no shutdown
  ip address 150.0.10.52/24
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 1
    preempt
    priority 90
    ip 150.0.10.100

WC-N7K1#
interface Vlan10
  no shutdown
  ip address 150.0.10.11/24

WC-N7K2#
interface Vlan10
  no shutdown
  ip address 150.0.10.12/24

W1-N7K1-AGG#ping 150.0.10.11
PING 150.0.10.11 (150.0.10.11): 56 data bytes
64 bytes from 150.0.10.11: icmp_seq=0 ttl=254 time=1.324 ms
64 bytes from 150.0.10.11: icmp_seq=1 ttl=254 time=0.907 ms
64 bytes from 150.0.10.11: icmp_seq=2 ttl=254 time=0.872 ms
64 bytes from 150.0.10.11: icmp_seq=3 ttl=254 time=0.889 ms
64 bytes from 150.0.10.11: icmp_seq=4 ttl=254 time=0.886 ms

--- 150.0.10.11 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.872/0.975/1.324 ms

W1-N7K1-AGG#ping 150.0.10.12
PING 150.0.10.12 (150.0.10.12): 56 data bytes
64 bytes from 150.0.10.12: icmp_seq=0 ttl=254 time=1.271 ms
64 bytes from 150.0.10.12: icmp_seq=1 ttl=254 time=0.835 ms
```

```
64 bytes from 150.0.10.12: icmp_seq=2 ttl=254 time=0.865 ms
64 bytes from 150.0.10.12: icmp_seq=3 ttl=254 time=0.862 ms
64 bytes from 150.0.10.12: icmp_seq=4 ttl=254 time=0.894 ms

--- 150.0.10.12 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.835/0.945/1.271 ms

W1-N7K2-AGG# ping 150.0.10.11
PING 150.0.10.11 (150.0.10.11): 56 data bytes
64 bytes from 150.0.10.11: icmp_seq=0 ttl=254 time=1.422 ms
64 bytes from 150.0.10.11: icmp_seq=1 ttl=254 time=1.728 ms
64 bytes from 150.0.10.11: icmp_seq=2 ttl=254 time=0.953 ms
64 bytes from 150.0.10.11: icmp_seq=3 ttl=254 time=0.872 ms
64 bytes from 150.0.10.11: icmp_seq=4 ttl=254 time=1.334 ms

--- 150.0.10.11 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.872/1.261/1.728 ms

W1-N7K2-AGG#ping 150.0.10.12
PING 150.0.10.12 (150.0.10.12): 56 data bytes
64 bytes from 150.0.10.12: icmp_seq=0 ttl=254 time=1.477 ms
64 bytes from 150.0.10.12: icmp_seq=1 ttl=254 time=0.953 ms
64 bytes from 150.0.10.12: icmp_seq=2 ttl=254 time=1.058 ms
64 bytes from 150.0.10.12: icmp_seq=3 ttl=254 time=0.862 ms
64 bytes from 150.0.10.12: icmp_seq=4 ttl=254 time=0.867 ms

--- 150.0.10.12 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 0.862/1.043/1.477 ms
```

5.2.2.5 Check HSRP Isolation

```
W1-N7K1-AGG# show run interface vlan 10
!Command: show running-config interface Vlan10
!Time: Tue Jan 17 12:08:11 2012

version 5.2(3a)

interface Vlan10
  no shutdown
  ip address 150.0.10.51/24
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
```

```
  hsrp 1
    preempt
    priority 110
     ip 150.0.10.100

W1-N7K2-AGG# show run interface vlan 10
!Command: show running-config interface Vlan10
!Time: Tue Jan 17 12:09:15 2012

version 5.2(3a)

interface Vlan10
  no shutdown
  ip address 150.0.10.52/24
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 1
    preempt
    priority 90
    ip 150.0.10.100
W5-N7K1-AGG# show run interface vlan 10
!Command: show running-config interface Vlan10
!Time: Tue Jan 17 12:08:17 2012

version 5.2(3a)

interface Vlan10
  no shutdown
  ip address 150.0.10.61/24
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 1
    preempt
    priority 110
    ip 150.0.10.100

W5-N7K2-AGG# show run interface vlan 10
!Command: show running-config interface Vlan10
!Time: Tue Jan 17 12:09:55 2012

version 5.2(3a)

interface Vlan10
  no shutdown
```

```
  ip address 150.0.10.62/24
  ip ospf passive-interface
  ip router ospf 1 area 0.0.0.0
  hsrp 1
    preempt
    priority 90
    ip 150.0.10.100

W1-N7K1-AGG# show hsrp interface vlan 10
Vlan10 - Group 1 (HSRP-V1) (IPv4)
  Local state is Active, priority 110 (Cfged 110), may preempt
    Forwarding threshold(for vPC), lower: 1 upper: 110
  Hellotime 3 sec, holdtime 10 sec
  Next hello sent in 1.212000 sec(s)
  Virtual IP address is 150.0.10.100 (Cfged)
  Active router is local
  Standby router is 150.0.10.52, priority 90 expires in 0.582000 sec(s)
  Authentication text "cisco"
  Virtual mac address is 0000.0c07.ac01 (Default MAC)
  13 state changes, last state change 00:03:34
  IP redundancy name is hsrp-Vlan10-1 (default)

W1-N7K2-AGG# show hsrp interface vlan 10
Vlan10 - Group 1 (HSRP-V1) (IPv4)
  Local state is Standby, priority 90 (Cfged 90), may preempt
    Forwarding threshold(for vPC), lower: 1 upper: 90
  Hellotime 3 sec, holdtime 10 sec
  Next hello sent in 1.673000 sec(s)
  Virtual IP address is 150.0.10.100 (Cfged)
  Active router is 150.0.10.51, priority 110 expires in 2.253000 sec(s)
  Standby router is local
  Authentication text "cisco"
  Virtual mac address is 0000.0c07.ac01 (Default MAC)
  8 state changes, last state change 00:01:40
  IP redundancy name is hsrp-Vlan10-1 (default)

W5-N7K1-AGG# show hsrp interface vlan 10
Vlan10 - Group 1 (HSRP-V1) (IPv4)
  Local state is Active, priority 110 (Cfged 110), may preempt
    Forwarding threshold(for vPC), lower: 1 upper: 110
  Hellotime 3 sec, holdtime 10 sec
  Next hello sent in 1.829000 sec(s)
  Virtual IP address is 150.0.10.100 (Cfged)
  Active router is local
```

```
    Standby router is 150.0.10.62, priority 90 expires in 4.679000 sec(s)
    Authentication text "cisco"
    Virtual mac address is 0000.0c07.ac01 (Default MAC)
    14 state changes, last state change 00:02:01
    IP redundancy name is hsrp-Vlan10-1 (default)

W5-N7K2-AGG# show hsrp interface vlan 10
Vlan10 - Group 1 (HSRP-V1) (IPv4)
    Local state is Standby, priority 90 (Cfged 90), may preempt
      Forwarding threshold(for vPC), lower: 1 upper: 90
    Hellotime 3 sec, holdtime 10 sec
    Next hello sent in 1.022000 sec(s)
    Virtual IP address is 150.0.10.100 (Cfged)
    Active router is 150.0.10.61, priority 110 expires in 3.872000 sec(s)
    Standby router is local
    Authentication text "cisco"
    Virtual mac address is 0000.0c07.ac01 (Default MAC)
    15 state changes, last state change 00:01:58
    IP redundancy name is hsrp-Vlan10-1 (default)
```

6. Test Results

Convergence time was tested with the following hardware and software configuration:

- Hardware
  - Cisco Nexus 7000 Series Supervisor Module
  - Cisco Nexus 7000 Series 48-port 10/100/1000 Ethernet Module
  - Cisco Nexus 7000 Series 32-Port 10Gb Ethernet Module
- Software
  - Cisco NX-OS Software Release 5.2(3a)

The Spirent testing tool was used to generate bidirectional traffic between E1-SRV1 in Data Center East and W1-SRV1 in Data Center West. All failure cases were run on the OTV AEDs.

6.1 Convergence Test

Figure 13 and Table 6 shows the test results for double link failures, which cause AED migration. With Layer 2 or Layer 3 single link failures, convergence occurs in less than one second.

Note: There is a corner case for reloading an entire core device. When the entire core device fails, the traffic will not be converged until the new ARP request has sended. With the device up, there is no problem. This problem is inherited from the OTV technology's unidirectional traffic AED migration problem in the current software release.

Figure 13. Traffic Flow Under Test

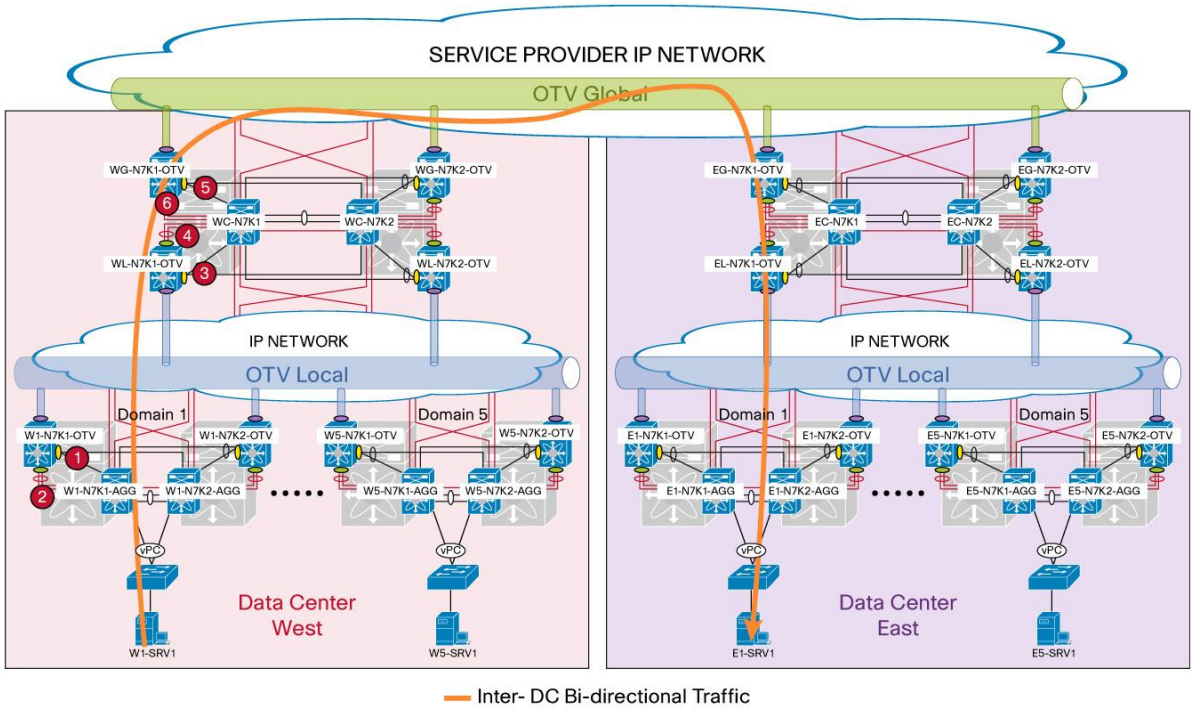

Table 6. Convergence Test Results

Index	Test Case	Failure (Seconds)		Recovery (Seconds)	
		Data Center West to Data Center East	Data Center East to Data Center West	Data Center West to Data Center East	Data Center East to Data Center West
1	W1-N7K1-OTV Layer 2 Isolation	4.85	4.513	4.534	4.536
2	W1-N7K1-OTV Layer 3 Isolation	4.91	6.02	4.091	7.628
3	WL-N7K1-OTV Layer 2 Isolation	4.627	4.894	5.10	4.819
4	WL-N7K1-OTV Layer 3 Isolation	4.795	4.795	6.715	4.908
5	WG-N7K1-OTV Layer 2 Isolation	3.895	3.922	4.155	3.898
6	WG-N7K1-OTV Layer 3 Isolation	4.603	6.09	3.893	5.361

6.2 Latency Test

Latency is tested using the Spirent test tool to generate bidirectional traffic between W-SRV1 in Data Center West and E1-SRV1in Data Center East (Table 7). Across the entire network, there are 18 hops along the path.

Table 7. Latency Test Results

Load	10%	20%	50%	70%	90%	95%
Average Latency (Microseconds)	435.798	441.323	446.349	448.896	458.688	473.438
Minimum Latency (Microseconds)	432.27	434.19	435.19	437.02	435.68	442.4
Maximum Latency (Microseconds)	449.13	459.58	492.64	500.26	502.93	518.21

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.

Printed in USA C11-702185-00 04/12