



Overlay Transport Virtualization (OTV) Unicast-Mode Transport Infrastructure Deployment

July 24, 2012

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

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.

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. (1005R)

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.

Overlay Transport Virtualization (OTV) Unicast-Mode Transport Infrastructure Deployment

© 2012 Cisco Systems, Inc. All rights reserved.



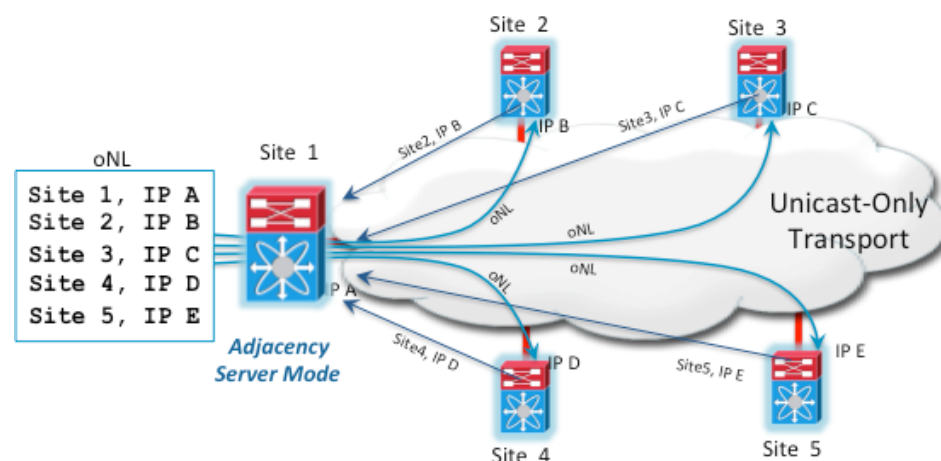
OTV Unicast-Mode Transport Infrastructure Deployment

Starting with NX-OS 5.2(1) release, OTV has the option to be deployed with unicast-only transport instead of multicast transport. The OTV control plane over a unicast-only transport works generally the same way as OTV with multicast mode. The only difference is that each OTV device would need to create multiple copies of each control plane packet and unicast them to each remote OTV device part of the same logical overlay. The operational simplification brought by the unicast-only model (removing the need for multicast deployment) can make this deployment option very appealing in scenarios where LAN extension connectivity is required only between few (2-3) DC sites.

To be able to communicate with all the remote OTV devices, each OTV node needs to know a list of neighbors to replicate the control packets to. Rather than statically configuring in each OTV node the list of all neighbors, a simple dynamic means is used to provide this information. This is achieved by designating one (or more) OTV Edge device to perform as an Adjacency Server. Every OTV device wishing to join a specific OTV logical overlay, needs to first “register” with the Adjacency Server (by sending OTV Hello messages to it). All other OTV neighbor addresses are discovered dynamically through the Adjacency Server. Thereby, when the OTV service needs to be extended to a new DC site, only the OTV edge devices for the new site need to be configured with the Adjacency Server addresses. No other sites need additional configuration.

Figure 1 shows each OTV Edge device registering with the Adjacency Server.

Figure 1 Adjacency Server Functionality



OTV Adjacency Server Configuration Procedure

The adjacency server is configured directly on the OTV device. Configure the local edge device to act as an adjacency server, or configure a remote adjacency server.

The following procedure, using the appropriate corresponding NX-OS commands, allows you to enable the adjacency server. Configuration examples are included.

Step 1 Enable the OTV feature.

Step 2 Enter global configuration mode.

configure terminal

```
switch# configure terminal
switch(config)#
```

Step 3 Create an OTV overlay interface and enters interface configuration mode.

interface overlay interface

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

Step 4 (Optional) Configure the local edge device to act as an adjacency server.

otv adjacency-server unicast-only

```
switch(config-if-overlay)# otv adjacency-server unicast-only
```

Step 5 (Optional) Configure the local edge device to use a remote adjacency server. The IP address format is in dotted decimal notation. The *secondary-ip-address* argument is the IP address of the backup adjacency server, if you have configured a backup adjacency server.

otv use-adjacency-server primary-ip-address [secondary-ip-address] unicast-only

```
switch(config-if-overlay)# otv use-adjacency-server 192.0.2.1 unicast-only
```

Step 6 (Optional) Display the OTV adjacency information. The *if-number* range is from 0 to 65503. The *vpn-name* is any case-sensitive, alphanumeric string up to 80 characters.

show otv adjacency [overlay if-number | vpn vpn-name] [detail]

```
switch(config-if-overlay)# show otv adjacency overlay 1
```

Step 7 (Optional) Copy the running configuration to the startup configuration.

copy running-config startup-config

```
switch(config-if-overlay)# copy running-config startup-config
```

Step 8 Exit configuration mode.

exit

OTV Adjacency Server Configuration Example

A pair of Adjacency Servers can be deployed for redundancy purposes. These Adjacency Server devices are completely stateless between them, which implies that every OTV edge device (OTV clients) should register its existence with both of them. For this purpose, the primary and secondary Adjacency Servers are configured in each OTV edge device. However, an OTV client will not process an alternate server's replication list until it detects that the primary Adjacency Server has timed out. Once that happens, each OTV edge device will start using the replication list from the secondary Adjacency Server and push the difference to OTV. OTV will stale the replication list entries with a timer of 10 minutes. If the Primary Adjacency Server comes back up within 10 mins, OTV will always revert back to the primary replication list. In case the Primary Adjacency Server comes back up after replication list is deleted, a new replication list will be pushed by the Primary after learning all OTV neighbors by means of OTV Hellos that are sent periodically.

OTV also uses graceful exit of Adjacency Server. When a Primary Adjacency Server is de-configured or is rebooted, it can let its client know about it and can exit gracefully. Following this, all OTV clients can start using alternate Server's replication list without waiting for primary Adjacency Server to time out.

There is no designation in the configuration to denote the distinction between the primary and secondary adjacency server when configuring the adjacency server itself. When pointing the edge devices to the adjacency server, the primary and secondary designation is implied in the command (see table 1 above).

The following exhibits show sample configuration for Adjacency server (Primary and Secondary).

1. Primary Adjacency Server Configuration (20.1.1.1)

```
feature otv
otv site-vlan 1999
otv site-identifier 0x111
interface Overlay200
    otv adjacency-server unicast-only
otv join-interface Ethernet1/2
otv extend-vlan 178, 181, 2500-2563
no otv suppress-arp-nd
no shutdown
```

2. Secondary Adjacency Server Configuration (20.2.1.1)

```
feature otv
otv site-vlan 1999
otv site-identifier 0x111
interface Overlay200
otv join-interface Ethernet1/2
otv extend-vlan 178, 181, 2500-2563
otv use-adjacency-server 20.1.1.1 unicast-only
otv adjacency-server unicast-only
no otv suppress-arp-nd
no shutdown
```

The adjacency servers were placed in separate data centers on one of the two Nexus 7000 switches. It is important to note that even though the Nexus is configured as an adjacency server, it is also an OTV edge device. As shown above, the configuration on the Primary Adjacency Server is very simple and limited to enable AS functionality (otv adjacency-server command). The same command is also required on the Secondary Adjacency Server device, but also needs to point to the Primary AS (leveraging the otv use-adjacency-server command).

All OTV edge devices must be configured to point to the adjacency servers and can be configured using the following as an example:

```
feature otv
otv site-vlan 1999
```

```

otv site-identifier 0x111
interface Overlay200
  otv join-interface Ethernet1/2
  otv extend-vlan 178, 181, 2500-2563
  otv use-adjacency-server 20.1.1.1 20.2.1.1 unicast-only
  no otv suppress-arp-nd
  no shutdown

```

The OTV Edge Device must be configured to use both the Primary and Secondary Adjacency Servers. The sequence of adjacency server address in the configuration determine primary or secondary adjacency server role. This order is relevant since an OTV edge device will always use the neighbor-list provided by the Primary Adjacency Server, unless it detects that specific device is not available anymore.

OTV Adjacency Server Configuration Verification

To verify OTV is utilizing the Unicast-only mode, you can use the **show otv** command. The output shows Unicast-Only capability and addresses of the adjacency servers:

```

dc1a-agg-7k1-otv# show otv
OTV Overlay Information
Site Identifier 0000.0000.0111
Overlay interface Overlay200
VPN name          : Overlay200
VPN state         : UP
Extended vlans   : 178 2500-2563 (Total:65)
Join interface(s) : Eth1/2 (20.1.1.1)
Site vlan        : 1999 (up)
AED-Capable     : Yes
Capability       : Unicast-Only
Is Adjacency Server : Yes
Adjacency Server(s) : 20.1.1.1 / 20.2.1.1

```

To view the replication list that the OTV devices are using, use the **show otv internal adjacency-server replication-list** command.

```

dc1a-agg-7k1-otv# show otv internal adjacency-server replication-list
Overlay Repl List database
Overlay-Interface Overlay200,
Dest Addr          Capability
20.1.1.1           U
20.1.2.1           U
20.2.1.1           U
20.2.2.1           U
20.3.1.1           U

```

The same commands that are useful for multicast mode OTV display similar information in Unicast-only mode. For example, **show otv adjacency** shows the other OTV edge devices in the network and their status.

```

dc1a-agg-7k1-otv# show otv adjacency
Overlay Adjacency database
Overlay-Interface Overlay200 :
Hostname          System-ID      Dest Addr      Up Time  State
dc1a-agg-7k2-otv  001b.54c2.e142 20.1.2.1      6d03h   UP
dc2a-agg-7k1-otv  0023.ac64.7042 20.2.1.1      6d03h   UP
dc2a-agg-7k2-otv  0022.5579.0f42 20.2.2.1      5d01h   UP
dc3a-agg-7k-otv   0026.982f.a0c2 20.3.1.1      4d23h   UP

```

This document assumes a working knowledge of OTV. Other required portions of the OTV configuration were not in the scope of this document. For more information on OTV configurations and design considerations, please refer to:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DCI/whitepaper/DCI3_OTV_Intro_WP.pdf



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.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn is a service mark; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0805R)