Configure a Layer 2 vPC Data Center Interconnect on a Nexus 7000 Series Switch

TAC **Document ID: 118934**

Contributed by Richard Rutkowski and Yogesh Ramdoss, Cisco TAC Engineers, and Clark Dyson, Cisco Advanced Services.
Aug 14, 2015

Contents

Introduction

This document describes how to configure a Layer 2 (L2) Data Center Interconnect (DCI) with the use of a Virtual Port-Channel (vPC).

Prerequisites

It is assumed that vPC and Hot Standby Routing Protocol (HSRP) are already configured on the devices that are used in the examples provided in this document.

Note: Link Aggregation Control Protocol (LACP) should be used on the vPC link, which acts as the DCI.

Tip: MACSec encryption requires a LAN advanced services license in versions prior to Version 6.1(1) and has linecard-specific limitations. Refer to the Guidelines and Limitations for Cisco TrustSec section of the **Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 6.x** for additional information.

Requirements

Cisco recommends that you have knowledge of these topics:

- vPC

- HSRP
- Spanning-Tree Protocol (STP)
- MACSec Encryption (optional)

Components Used

The information in this document is based on a Cisco Nexus 7000 Series switch that runs software Version 6.2(8b).

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Background Information

The purpose of a DCI is to extend specific VLANs between different data centers, which offers L2 adjacency for servers and Network-Attached Storage (NAS) devices that are separated by large distances.

The vPC presents the benefit of STP isolation between the two sites (no Bridge Protocol Data Unit (BPDU) across the DCI vPC), so any outage in a data center is not propagated to the remote data center because redundant links are still provided between the data centers.

Note: The vPC can be used in order to interconnect a maximum of two data centers. If more than two data centers must be interconnected, Cisco recommends that you use Overlay Transport Virtualization (OTV).

A DCI vPC etherchannel is typically configured with this information in mind:

- First Hop Redundancy Protocol (FHRP) isolation: Prevent sub-optimal routing with the use of a dedicated gateway for each data center. Configurations vary dependent upon the location of the FHRP gateway.
- STP isolation: As previously mentioned, this prevents the propagation of outages from one data center to another.
- Broadcast storm control: This is used in order to minimize the amount of broadcast traffic between the data centers.
- MACSec Encryption (optional): This encrypts the traffic in order to prevent intrusion between the two facilities.

Configure

Use the information that is described in this section in order to configure an L2 DCI with the use of a vPC.

Note: Use the Command Lookup Tool (registered customers only) in order to obtain more information on the commands used in this section.

FHRP Isolation

This section describes two scenarios for which FHRP isolation can be implemented.

Dual L2/L3 POD Interconnect

This is the topology that is used in this scenario:

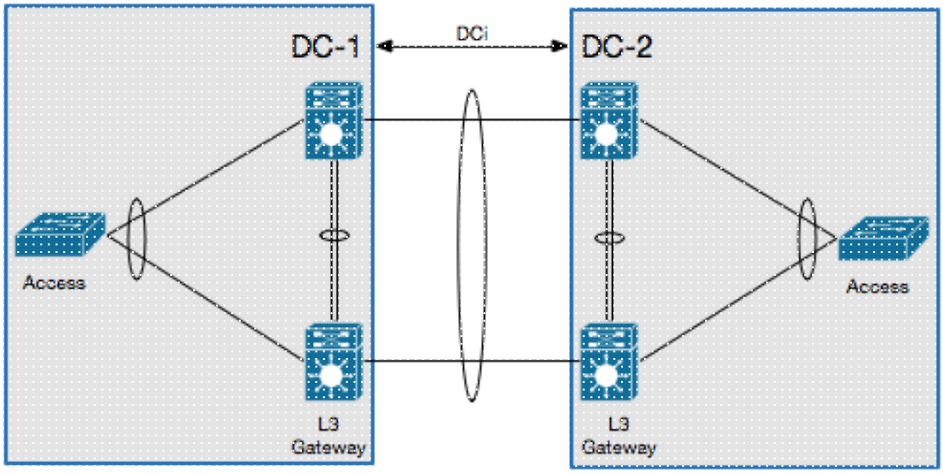

In this scenario, the Layer 3 (L3) gateway is configured on the same vPC pair and acts as the DCI. In order to isolate the HSRP, you must configure a Port Access Control List (PACL) on the DCI port-channel and disable HSRP Gratuitous Address Resolution Protocols (ARPs) (GARPs) on the Switched Virtual Interfaces (SVIs) for the VLANs that move across the DCI.

Here is an example configuration:

```
ip access-list DENY_HSRP_IP
  10 deny udp any 224.0.0.2/32 eq 1985
  20 deny udp any 224.0.0.102/32 eq 1985
  30 permit ip any any

interface <DCI-Port-Channel>
  ip port access-group DENY_HSRP_IP in

interface Vlan <x>
  no ip arp gratuitous hsrp duplicate
```

Note: The previous configuration can also be used with Nexus 9000 switches.

Multilayer vPC for Aggregation and DCI

This is the topology that is used in this scenario:

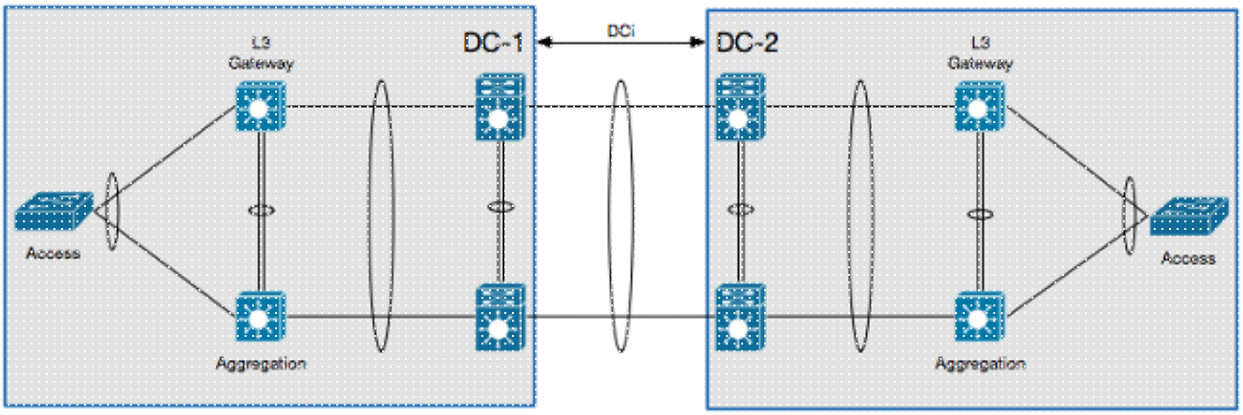

In this scenario, the DCI is isolated on its own L2 Virtual Device Context (VDC), and the L3 gateway is on an aggregation layer device. In order to isolate the HSRP, you must configure a VLAN Access Control List (VACL) that blocks the HSRP control traffic and an ARP inspection filter that blocks the HSRP GARPs on the L2 DCI VDC.

Here is an example configuration:

```
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
vlan access-map HSRP_Localization 10
        match ip address HSRP_IP
        match mac address HSRP_VMAC
        action drop
        statistics per-entry
vlan access-map HSRP_Localization 20
        match ip address ALL_IPs
        match mac address ALL_MACs
        action forward
        statistics per-entry
vlan filter HSRP_Localization vlan-list <DCI_Extended_VLANs>

feature dhcp

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

ip arp inspection filter HSRP_VMAC_ARP vlan <DCI_Extended_VLANs>
```

Additional Isolation Configuration

This section provides an example configuration that:

- Allows only the VLANs that are needed at the remote data center to be extended.
- Isolates the STP at each data center.
- Drops the broadcast traffic that exceeds 1% of the total link speed.

Here is the example configuration:

```
 interface <DCI-Port-Channel>
  switchport trunk allowed vlan <DCI_Extended_VLANs>
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  storm-control broadcast level 1.0
```

Note: Storm control for multicast traffic can also be configured, but it must have the same percentage as the broadcast traffic.

MACSec Encryption

Note: The configuration that is described in this section is optional.

Use this information in order to configure MACSec encryption:

```
feature dot1x
feature cts

! MACSec requires 24 additional bytes for encapsulation.
interface <DCI-Port-Channel>
  mtu 1524

interface <DCI-Physical-Port>
  cts manual
    no propagate-sgt
    sap pmk <Preshared-Key>
```

Note: The interface must be flapped in order for MACSec authorization to occur.

Verify

Use the information that is described in this section in order to confirm that your configuration works properly.

FHRP Isolation

Enter the **show hsrp br** command into the CLI in order to verify that the HSRP gateway is active at both data centers:

```
!DC-1
N7K-A# show hsrp br
*:IPv6 group   #:group belongs to a bundle
                     P indicates configured to preempt.
                     |
 Interface   Grp  Prio P State    Active addr      Standby addr     Group addr
  Vlan10      10   120    Active   local            10.1.1.3         10.1.1.5
     (conf)

!DC-2
N7K-C# show hsrp br
*:IPv6 group   #:group belongs to a bundle
                     P indicates configured to preempt.
                     |
 Interface   Grp  Prio P State    Active addr      Standby addr     Group addr
  Vlan10      10   120    Active   local            10.1.1.3         10.1.1.5
     (conf)
```

Enter this command into the CLI in order to verify the ARP filter:

```
N7K-D# show log log | i DUP_VADDR
2015 Apr 10 21:16:45 N7K-A %ARP-3-DUP_VADDR_SRC_IP:  arp [7915]  Source address of
 packet received from 0000.0c9f.f00a on Vlan10(port-channel102) is duplicate of local
 virtual ip, 10.1.1.5
```

If an output similar to this appears, then the GARPs between the two active gateways is not properly isolated.

Additional Isolation

Enter the **show spanning-tree root** command into the CLI in order to verify that the STP root does not point towards the DCI port-channel:

```
N7K-A# show spanning-tree root

                                        Root  Hello Max Fwd
Vlan                 Root ID            Cost  Time  Age Dly  Root Port
---------------- -------------------- ------- ----- --- ---  ----------------
VLAN0010          4106 0023.04ee.be01       0    2   20  15  This bridge is root
```

Enter this command into the CLI in order to verify that storm control is properly configured:

```
N7K-A# show interface <DCI-Port-Channel> counters storm-control

--------------------------------------------------------------------------------
Port        UcastSupp %     McastSupp %     BcastSupp %     TotalSuppDiscards
--------------------------------------------------------------------------------
Po103            100.00          100.00            1.00                     0
```

MACSec Encryption

Enter this command into the CLI in order to verify that MACSec encryption is properly configured:

```
N7K-A# show cts interface <DCI-Physical-Port>
CTS Information for Interface Ethernet3/41:
...
    SAP Status:             CTS_SAP_SUCCESS
      Version: 1
      Configured pairwise ciphers: GCM_ENCRYPT
      Replay protection: Enabled
      Replay protection mode: Strict
      Selected cipher: GCM_ENCRYPT
      Current receive SPI: sci:e4c7220b98dc0000 an:0
      Current transmit SPI: sci:e4c7220b98d80000 an:0
...
```

Troubleshoot

There is currently no specific troubleshooting information available for the FHRP or additional isolation configurations.

For MACSec configuration, if the pre-shared key is not agreed upon on both sides of the link, you see an output similar to this when you enter the **show interface <DCI-Physical-Port>** command into the CLI:

```
N7K-A# show interface <DCI-Physical-Port>
Ethernet3/41 is down (Authorization pending)
admin state is up, Dedicated Interface
```

Note: The key must be the same on both sides of the connection.

Caveats

Note: Caveats for the related products are not included.

These caveats are related to the use of a DCI on the Cisco Nexus 7000 Series switch:

- Cisco bug ID CSCur69114 - *HSRP PACL Filter Broken - Packets are flooded to layer2 domain*. This bug is found only in software Version 6.2(10).

- Cisco bug ID CSCut75457 - *HSRP VACL Filter Broken*. This bug is found only in software Versions 6.2(10) and 6.2(12).

- Cisco bug ID CSCut43413 - *DCi: HSRP Virtual MAC Flapping through FHRP Isolation PACL*. This bug is due to a hardware limitation.

Related Information

- **Data Center Designs: Data Center Interconnect**
- **OTV Technology Introduction and Deployment Considerations**
- **Cisco Virtualized Workload Mobility Design Considerations**
- **Technical Support & Documentation - Cisco Systems**

Updated: Aug 14, 2015 Document ID: 118934