Architecture guide

Deploying Ethernet Virtual Interconnect

Table of contents

Introduction

This architecture guide provides guidelines on deploying Ethernet Virtual Interconnect (EVI) in data centers (DCs) based on HP 12500 Switch Series (12504/12508/12518/12508E/12518E).

Solution overview

EVI is a Media Access Control (MAC)-over-Generic Routing Encapsulation (GRE) technology that provides Layer 2 (L2) connectivity between distant DCs across any WAN transport as shown in figure 1 to build a virtual DC.

Figure 1. EVI

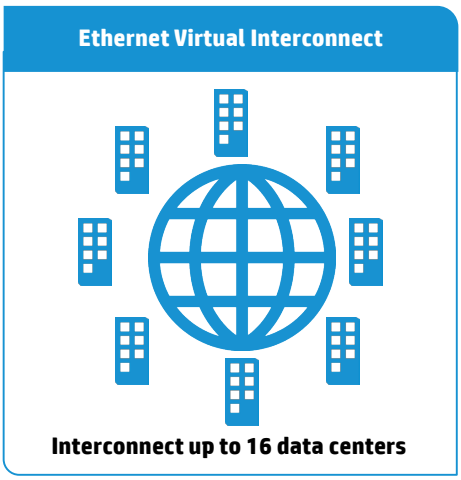

Summary of EVI benefits:

- Reduces Data Center Interconnect (DCI) configuration setup time to minutes
- Simplifies deployment but no changes to the existing networking infrastructure
- Expands to new data centers without disruptions seamlessly
- Provides automatic loop avoidance and failure domain isolation
- Enables automatic high availability (HA) and load balancing when used with Multichassis Link Aggregation Group (MLAG) and HP Intelligent Resilient Framework (IRF)
- Scales up to 16 DCs and 4K of Virtual Local Area Networks (VLANs)
- Reduces ARP request broadcasts on the EVI network with Address Resolution Protocol (ARP) flooding suppression
- Provides enhanced server-to-client traffic paths with Active/Active DCs
- Supports applications, physical servers, and virtual machines (VM) that require network connectivity within the same subnet across DCs
- Supports VMware long distance vMotion and Microsoft® Hyper-V live migration

EVI fundamentals

- The EVI Edge Device (ED) implements EVI, performs L2 switching towards the DC LAN, performs routing on the WAN interfaces, and encapsulates/de-encapsulates Ethernet frames in GRE over the WAN.
- The EVI network provides VLAN extension, formed by the EVI-tunnel interfaces and EVI links.
- The EVI-connect interface is a Layer 3 (L3) interface, it is the source of EVI (GRE) tunneled traffic, it can be from either a physical or logical interface, and is shared by multiple EVI networks as shown in figure 2.

Figure 2. EVI fundamentals

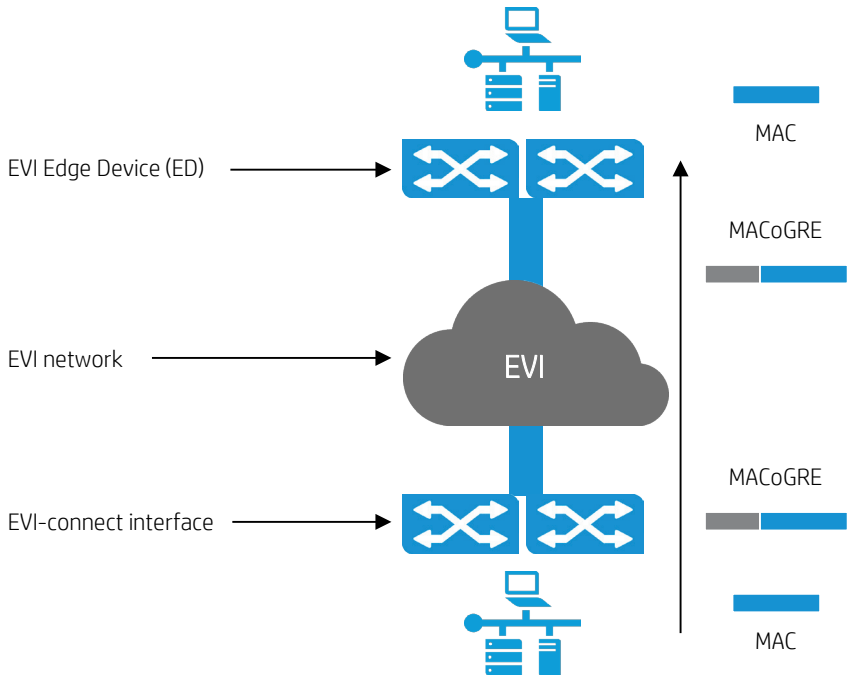

Up to 32 EVI networks can be created to provide multitenancy and to limit VLANs only to certain DCs. Each EVI network has a separate control/forwarding plane, and a unique network ID to extend a list of VLANs as shown in figure 3.

Figure 3. EVI networks with different VLANs and DCs

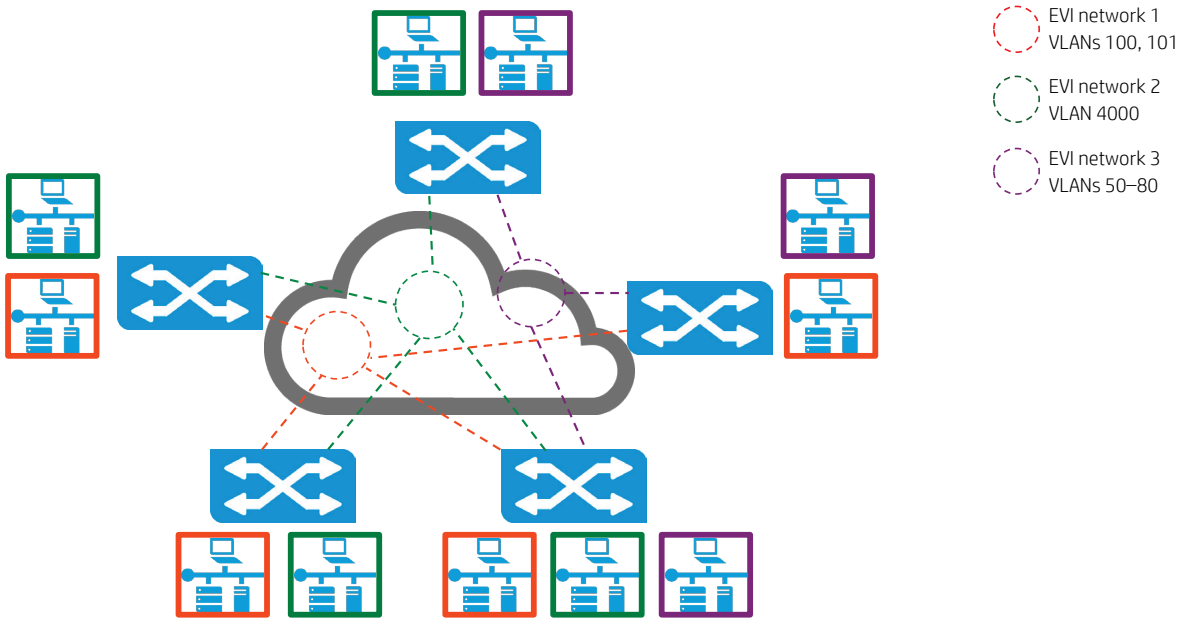

EVI control plane—Neighbor Discovery

The EVI Neighbor Discovery Protocol (ENDP) automatically discovers sites when EDs are setup in EVI Neighbor Discovery Server (ENDS)/EVI Neighbor Discovery Client (ENDC) mode. ENDP sets up/maintains EVI links between neighbors as shown in figure 4.

Two active ENDSs can be deployed for HA. This requires ENDCs to be configured with IPs of both ENDS. Security adjacency authentication can be implemented if required as shown in figure 5.

Figure 4. ENDC, ENDS, and EVI network formation

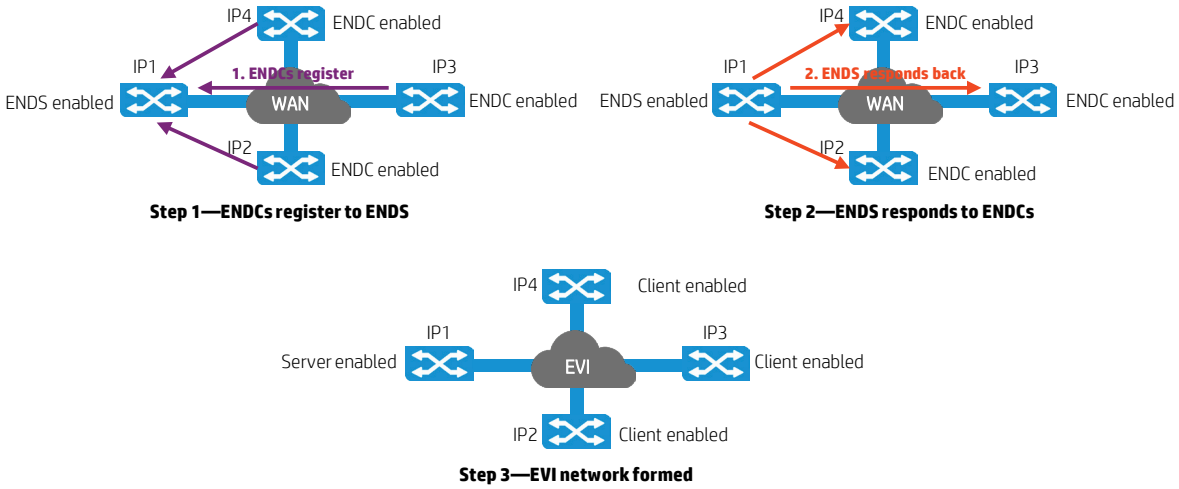

Figure 5. Multiple ENDS and EVI adjacency authentication

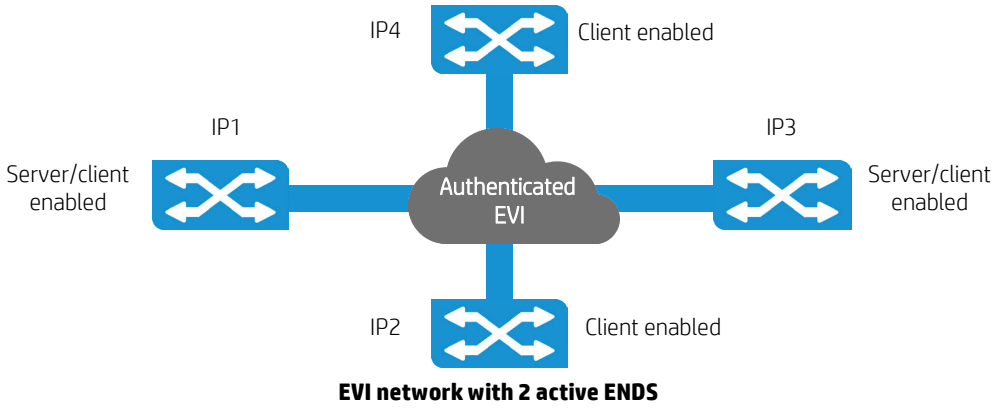

EVI control plane—MAC propagation flow

Once the EVI network is formed, MAC entries are propagated as shown in the following example and figure 6:

1. Site 1 learns new local MAC entries in VLAN 100
2. EVI Intermediate System to Intermediate System (IS-IS) process creates an LSP update
3. ED sends the LSP to its neighbors
4. Each neighbor delivers the LSP to its EVI IS-IS process
5. Each neighbor's EVI IS-IS process updates the MAC table

EDs are notified of MAC removals or age-outs in the same manner.

This MAC address change is initiated from any ED with local MAC changes.

Figure 6. MAC propagation

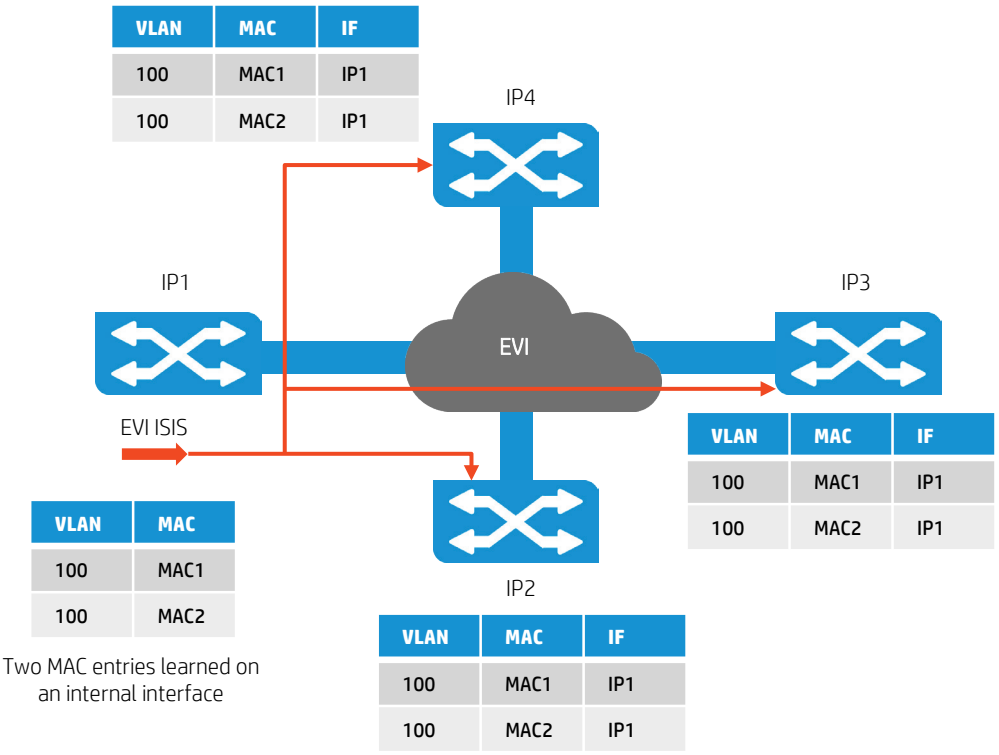

Two MAC entries learned on an internal interface

EVI data plane—Encapsulation

The ED encapsulates L2 frames in GRE, leaves the original frame intact, sets the DF bit in the outer IP header, adds the outer IP header/link-layer header/checksum and increases the packet size by 42 bytes as shown in figure 7.

Figure 7. EVI encapsulation

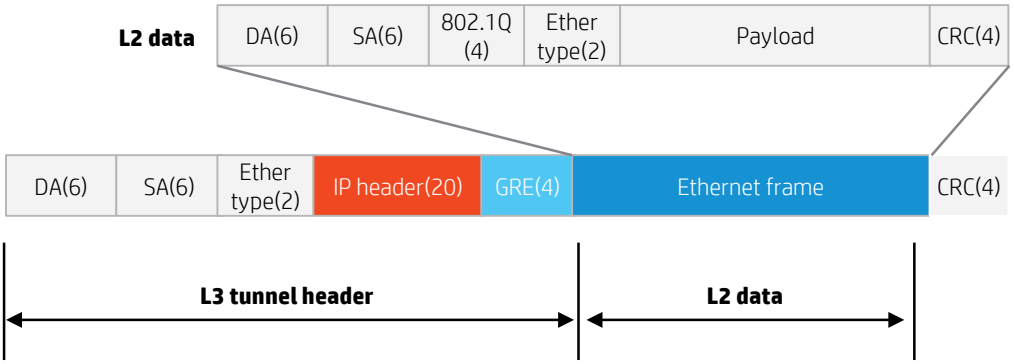

EVI data plane—Unicast forwarding

Unicast forwarding is explained below and shown in figure 8:

1. The source ED learns the source MAC address of the incoming Ethernet frame and looks up the destination MAC address in its MAC table for the outgoing interface.
2. If the outgoing interface is a local to the ED, normal switching is performed. If the outgoing interface is an EVI-link interface instead of a local port, the source ED encapsulates the frame in a GRE header and then adds an IP header and a link-layer protocol header. In the outer IP header, the source IP address is the source ED's tunnel source IP address and the destination IP address is the destination ED's tunnel source IP address.
3. The source ED forwards the encapsulated packet out of the EVI link to the destination ED across the IP transport network.
4. The destination ED removes the headers of the original Ethernet frame, looks up the destination MAC address in the MAC address table, and sends the frame out of the matching outgoing interface.

Figure 8. EVI unicast forwarding

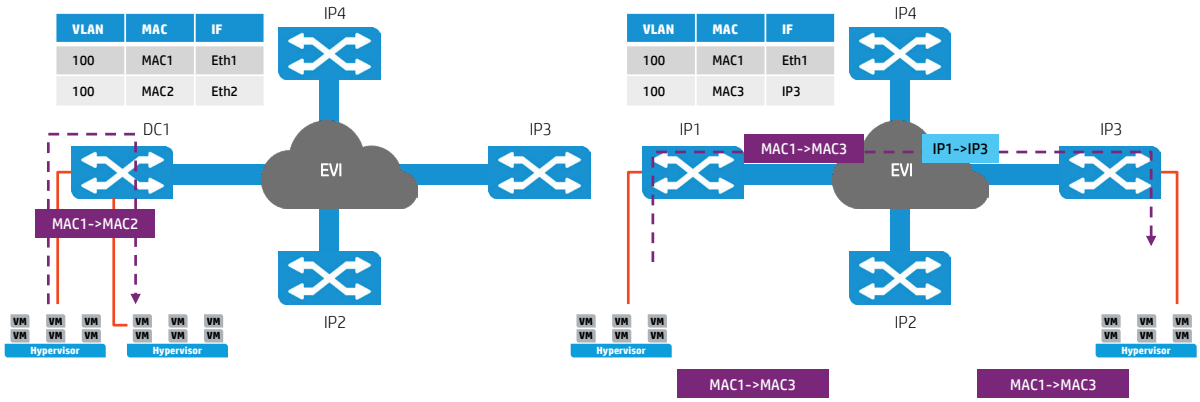

EVI data plane—Multicast forwarding

EVI does not require multicast support in the WAN transport to forward multicast traffic between sites. In order to forward multicast traffic, multicast MAC addresses need to be enabled with selective flooding within EVI. The ED uses head-end replication to tunnel Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD) and Protocol Independent Multicast (PIM) protocol packets to all remote EDs over EVI links as shown in figure 9.

Multicast forwarding is explained below:

1. The multicast Designated Router (DR) in a site sends out a multicast frame.
2. The source ED copies the frame and encapsulates one copy on each multicast member EVI-link interface.
3. The source ED unicasts the encapsulated frames to the destination EDs over the EVI links.
4. Each destination ED removes the headers of the multicast frame and copies the multicast frame on each multicast member interface.
5. Each destination ED sends the multicast frame out of all member interfaces to the destination hosts.

Figure 9. EVI multicast forwarding

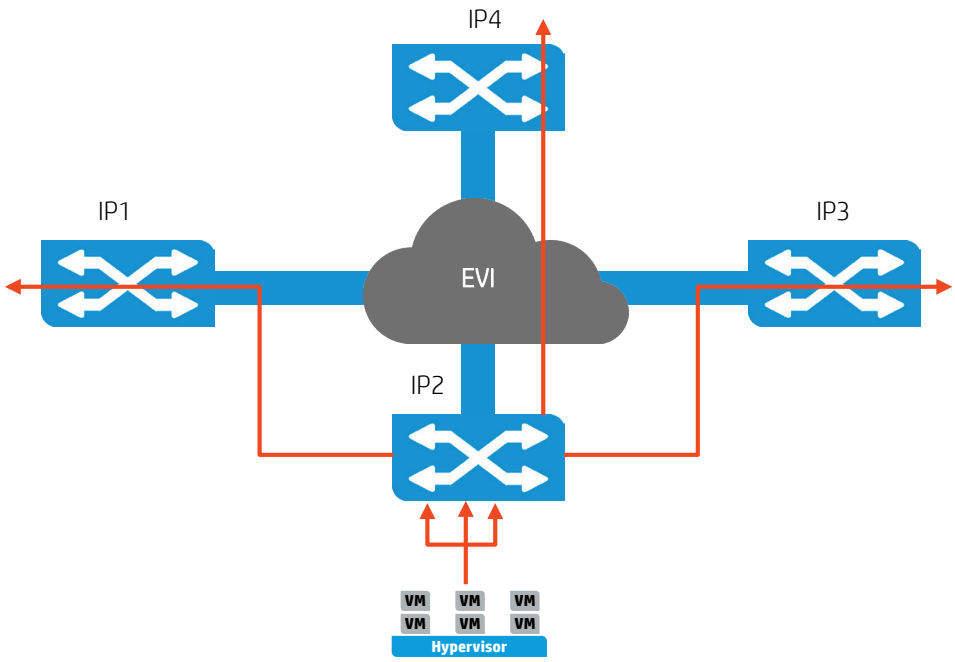

EVI features

Loop prevention

EVI by default implements split horizon by preventing frames received from EVI tunnels from being forwarded back into the WAN to other EDs as shown in figure 10. This functionality is built-in into EVI and no additional configuration is required.

Figure 10. EVI-loop prevention feature

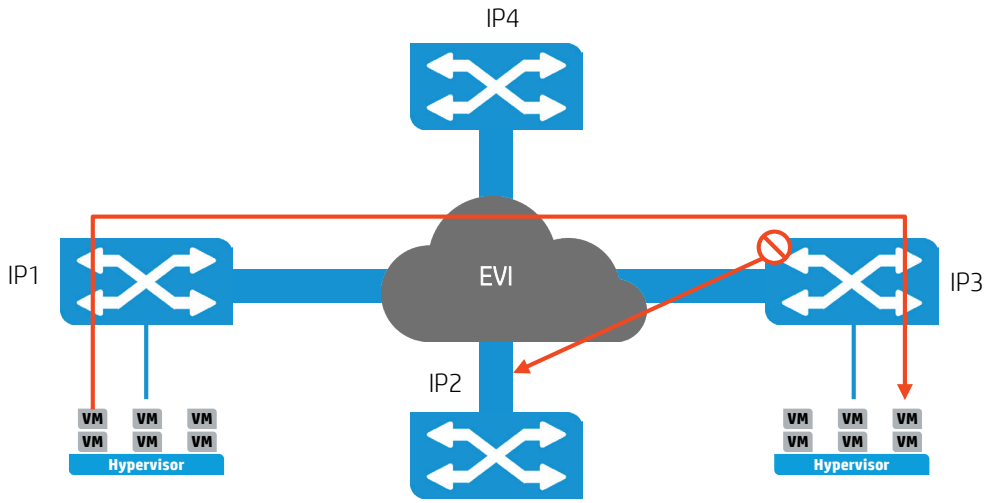

STP isolation

STP is disabled on EVI links, preventing BPDUs from being extended across sites as shown in figure 11. This ensures STP topology changes are contained within sites and allows different STP protocols to be turned on within each site. This functionality is built-in into EVI and no additional configuration is required.

Figure 11. STP-isolation feature

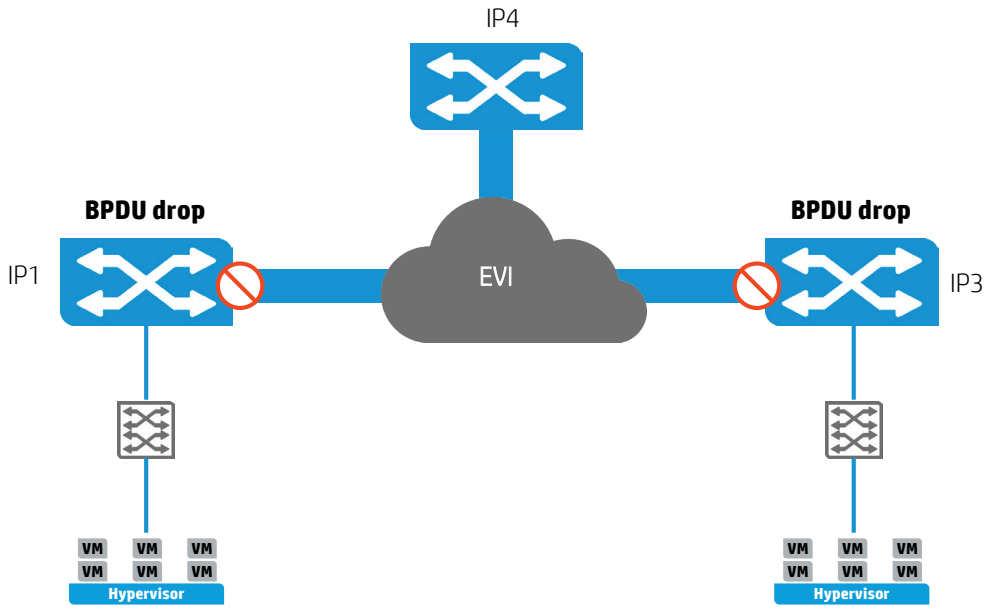

Unknown frame treatment

EVI helps minimize unnecessary WAN utilization by handling destination unknown frames (unicast, multicast) in the following manner. Unknown frames are flooded to internal interfaces and dropped on EVI-tunnel interfaces as shown in figure 12. This functionality is built-in into EVI and no additional configuration is required.

Figure 12. Unknown frame treatment feature

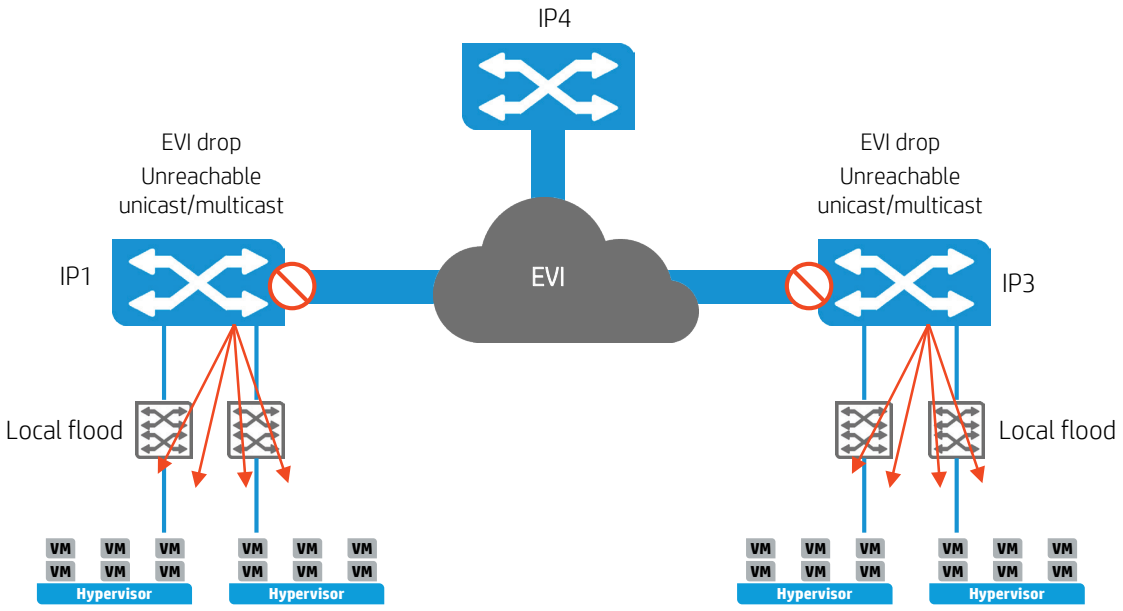

ARP flooding suppression

EVI helps minimize unnecessary WAN utilization by enabling EDs to reply to ARP requests on behalf of remote-site hosts as shown in figure 13. Additional configuration is required to enable this feature, please refer to configuration guides for details.

Figure 13. EVI ARP flooding suppression

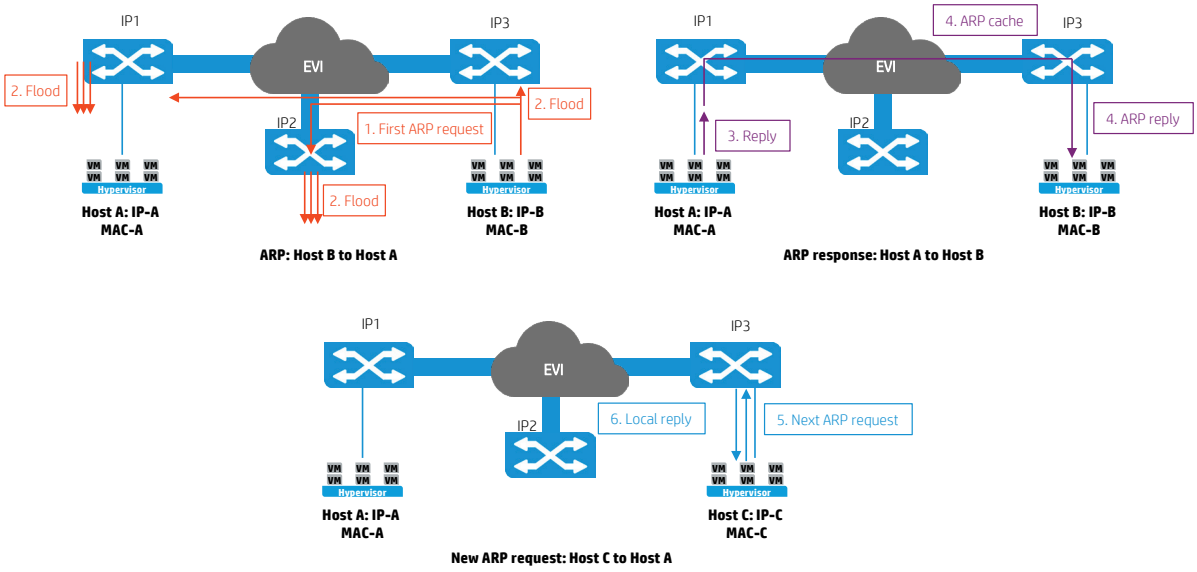

EVI HA with IRF

To provide for chassis redundancy and increased port density, IRF technology, as shown in figure 14 can be deployed together with EVI. With IRF, there exists only one appointed edge forwarder, this allows both chassis in an IRF to forward and receive traffic for the same VLANs at the same time. Additional configuration and a redundant chassis is required to enable this feature, please refer to configuration guides for details.

Figure 14. IRF, EVI, and EDs at each DC

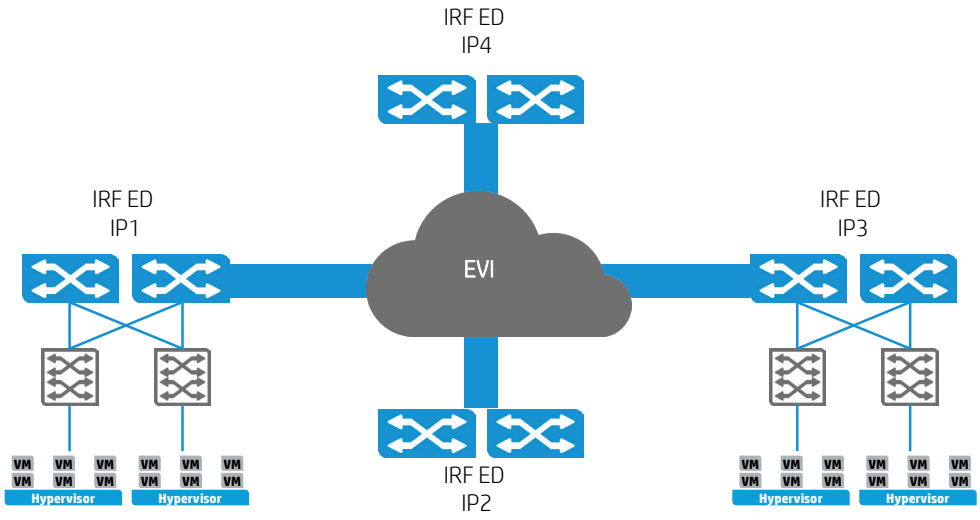

Efficient server-to-client traffic path

Inbound client-to-server traffic typically uses the shortest path from the routing table as shown in figure 15. EVI helps to optimize the outbound server to client traffic path by facilitating active default gateways at each DC. This is achieved by blocking Virtual Router Redundancy Protocol (VRRP) hellos at each DC from traversing into the DCI WAN. By blocking VRRP hellos, each active DC will maintain VRRP master state with active default gateways. This functionality is built in into EVI and no additional configuration is required. In certain environments, this leads to asymmetric routing, firewalls at each DC typically need extra configuration to allow asymmetric routed traffic.

Figure 15. Inbound client-to-server traffic and improved outbound server-to-client traffic

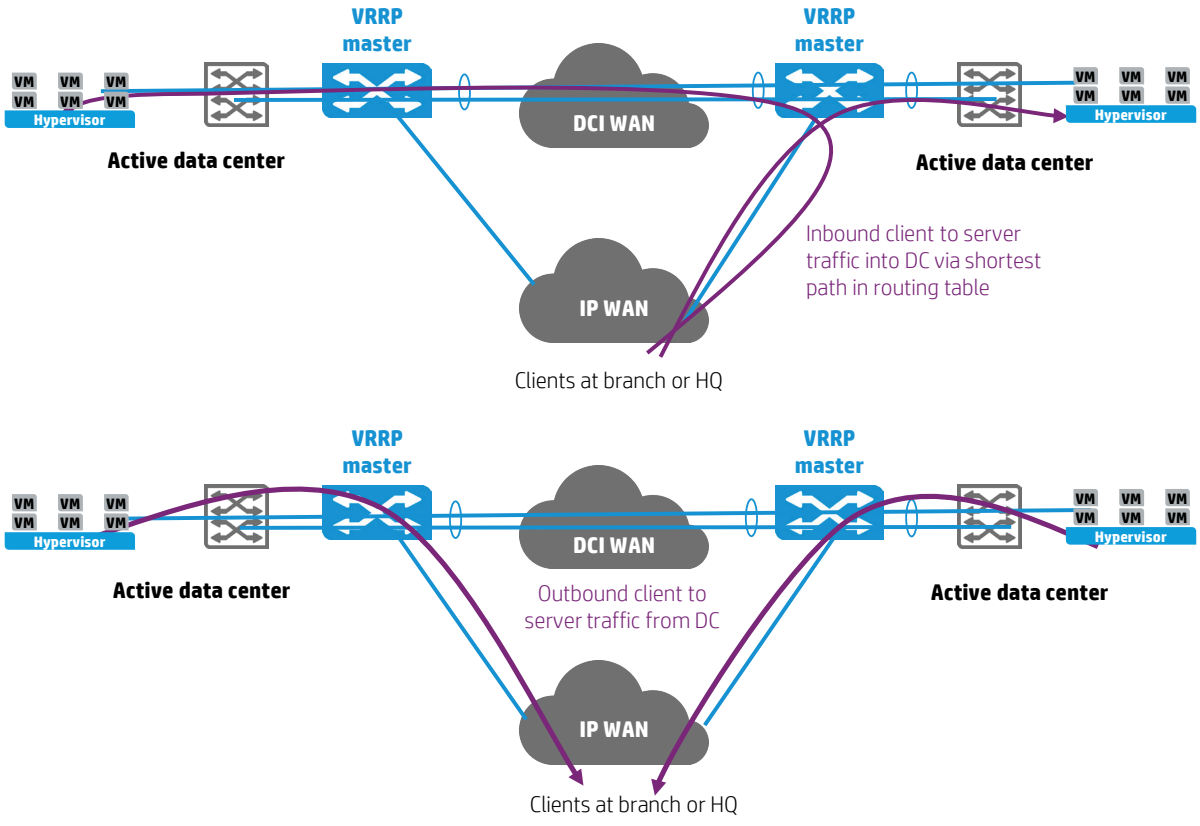

EVI deployment best practices

This section highlights some best practices when deploying EVI.

Utilize different ports for LAN/WAN interfaces

It is recommended that different ports be used for the LAN and WAN segments, e.g., LAN port = G1/4/0/1, WAN port = G1/4/0/15

Do not use the same port with 802.1Q trunking for both LAN/WAN.

```
interface GigabitEthernet1/4/0/1
 port link-mode bridge
 port link-type trunk
 undo port trunk permit vlan 1
 port trunk permit vlan 10 to 30
#
interface GigabitEthernet1/4/0/15
 port link-mode route
 ip address 200.200.200.1 255.255.255.252
 evi enable
```

EVI-connect interface

It is recommended to use the loopback interface as the EVI-connect interface. This can help ensure the EVI tunnel stays up even though physical interface or chassis goes down.

L2 multicast traffic extension across DCs

EVI does not forward multicast traffic by default, to enable flooding of certain multicast MAC addresses, this has to be explicitly enabled. For example:

```
[SwitchA-Tunnel0] evi selective-flooding mac-address 0100-5e00-0001 vlan 100
```

Take note that all multicast applications/MACs on that VLAN will need to be added as shown above. If the objective is to allow multicast applications/MACs to be forwarded without explicitly specifying each multicast MAC, the following can be enabled for the multicast applications.

```
[SwitchA-Tunnel0] evi flooding enable
```

Take note that “evi flooding enable” does not forward all multicast MACs, specifically control plane protocols such as: PIM hellos (destination MAC 0100-5e00-0001 and destination IP 224.0.0.1) and IGMP general query packets (destination MAC 0100-5e00-000D and destination IP 224.0.0.13). These destination MAC addresses cannot be learned in the data plane. Selective flood has to be enabled for these packets to be sent between DCs.

```
[SwitchA-Tunnel0] evi selective-flooding mac-address 0100-5e00-0001 vlan 100
[SwitchA-Tunnel0] evi selective-flooding mac-address 0100-5e00-000D vlan 100
```

Another example would be to allow OSPF adjacencies across EVI.

```
[SwitchA-Tunnel0] evi selective-flooding mac-address 0100-5e00-0005 vlan 100
[SwitchA-Tunnel0] evi selective-flooding mac-address 0100-5e00-0006 vlan 100
```

L3 multicast routing on a separate device

If multicast routing is required (e.g., with PIM SM/IGMP querier functionality as shown in sample configuration below), it is recommended that a different physical device/MDC be used instead.

```
[Router] multicast routing
[Router] vlan 618
[Router] interface vlan-interface 618
[Router-Vlan-interface618] ip address 155.111.142.10 255.255.255.192
[Router-Vlan-interface618] pim sm
[Router-Vlan-interface618] igmp enable
```

WAN MTU

For 1500 byte traffic, EVI can increase the maximum frame size to 1542 with DF bit set (42 byte for additional GRE header). Users should ensure that the DCI WAN can support the above MTU size in order to forward applications that requires 1500 with DF bit set.

For example: Users could request for a 1600 MTU WAN between DCs for long distance vMotion to work.

In order to support Jumbo Frames across EVI, the DCI WAN provider would need to support the appropriate MTU size between DCs.

In summary: Jumbo frames are supported if the WAN is able to accommodate the extra MTU overhead, fragmentation is not supported.

High availability

Depending on EVI deployment, the following could be useful to help minimize impact of traffic loss across EVI due to link and chassis failures:

- IRF 2 x 12500s into 1 EVI ED for chassis redundancy
- Ensure at least 2 ENDS are enabled for an EVI network at different locations
- Link delay 0 for interfaces to speed up link failover

```
interface GigabitEthernet X/X/X/X
  description WAN
  link-delay 0 mode updown
```

- BFD for BGP to speed up failure detection

```
bgp X
  peer X.X.X.X bfd
```

- OSPF graceful restart to speed up IRF master failover

```
ospf X
  graceful-restart ietf
```

- EVI ISIS graceful restart to speed up IRF master failover

```
evi-isis 1
  graceful-restart
```

WAN link-load sharing

Due to the load balancing hashing algorithm or routing policies implemented by the DCI WAN service provider, EVI traffic between DCs using a pair of loopback IP address (connect interfaces) might stick to the same WAN path. It might be beneficial to spread out VLANs to other EVI networks and connect interfaces as shown in figure 16 for more efficient WAN load distribution.

Figure 16. EVI network load sharing across DCI WAN using different loopback addresses

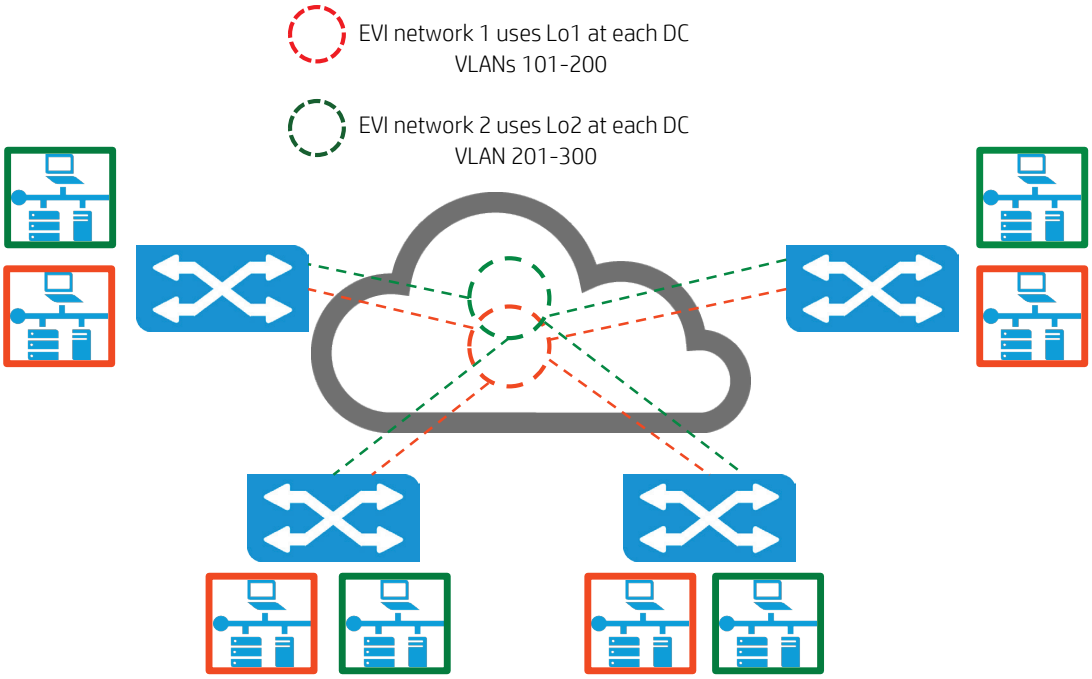

If BGP is used to interconnect DCs, the following sample deployment and configuration as shown in figure 17 can be used to help set EVI-connect interfaces route preferences into the DCI WAN.

Figure 17. EVI network load sharing across DCI WAN using BGP AS path prepending

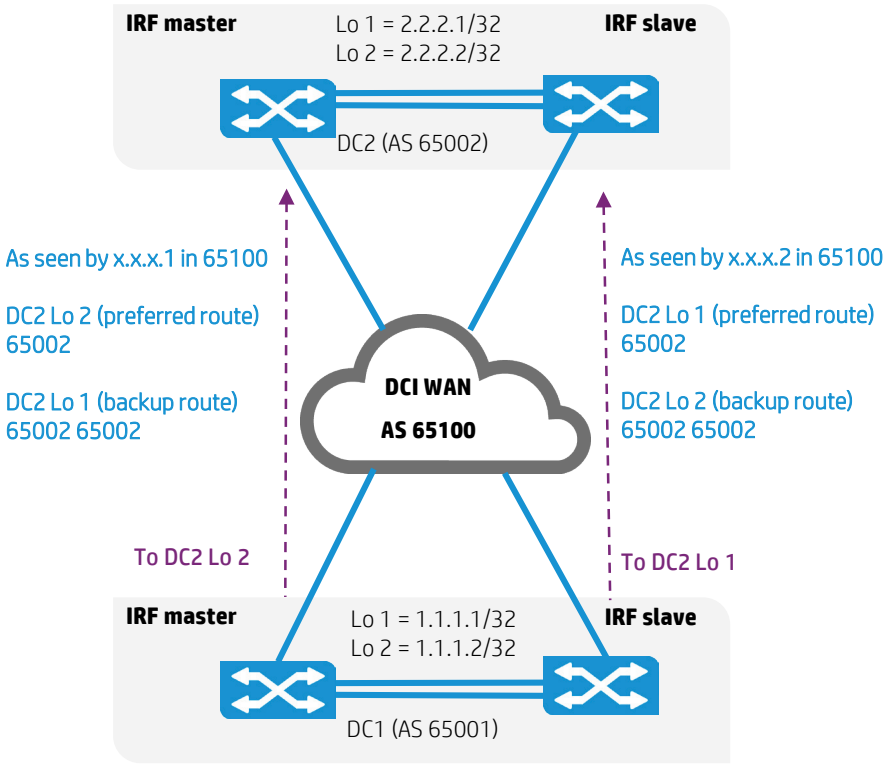

```
sysname DC2-EVI-ED
#
ip prefix-list BGP1 permit 2.2.2.1 32
ip prefix-list BGP2 permit 2.2.2.2 32
#
route-policy BGP1 permit node 1
 if-match ip address prefix-list BGP1
 apply as-path 65002
#
route-policy BGP2 permit node 1
 if-match ip address prefix-list BGP2
 apply as-path 65002
#
bgp 65002
 peer x.x.x.1 as-number 65100
 peer x.x.x.2 as-number 65100
#
 ipv4-family unicast
 network 2.2.2.1 255.255.255.255
 network 2.2.2.2 255.255.255.255
 peer x.x.x.1 enable
 peer x.x.x.1 route-policy BGP1 export
 peer x.x.x.2 enable
 peer x.x.x.2 route-policy BGP2 export
```

System working mode

Before deploying EVI, it is recommended to plan for appropriate system working modes depending on the needs of the network with regards to MAC/ARP/FIB table size. Changing the system working mode in the future would require an entire chassis reboot. Configurations might also be impacted due to a system working mode change, e.g., routed ports will be removed when the system working mode is changed from “routed” to “standard.”

EVI QoS

EVI QoS works by mapping QoS priority from LAN onto GRE encapsulated packets, this allows high-priority traffic to be forwarded as expected during congestion. An example of EVI QoS implementation is shown in figure 18. Inbound traffic from DC1 LAN into ED is either DSCP or dot1p trusted and then queued outbound on the WAN interface. The same is done for traffic from DC2 to DC1 in the opposite direction.

QoS implementation within the WAN is important: high priority GRE encapsulated EVI traffic should be given preference against lower priority traffic by trusting DSCP markings from both DCs.

Figure 18. EVI QoS deployment

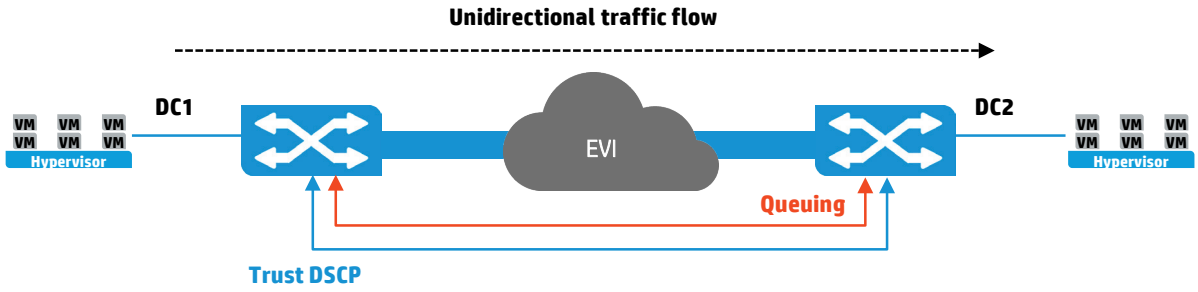

IRF/MDC planning

- EVI does not require Multitenant Device Context (MDC) in order to support both EVI and "interface VLANs" within the same physical chassis
- EVI supports 4-chassis IRF without MDCs
- EVI can work together with both IRF and MDC if desired as shown in figure 19. However when MDC is implemented, EVI is currently limited to 2-chassis IRF. It is therefore recommended to plan for future port requirements and purchase appropriate 12500 chassis, e.g., 4 slot, 8 slot, 18 slot

Figure 19. EVI network load sharing across DCI WAN

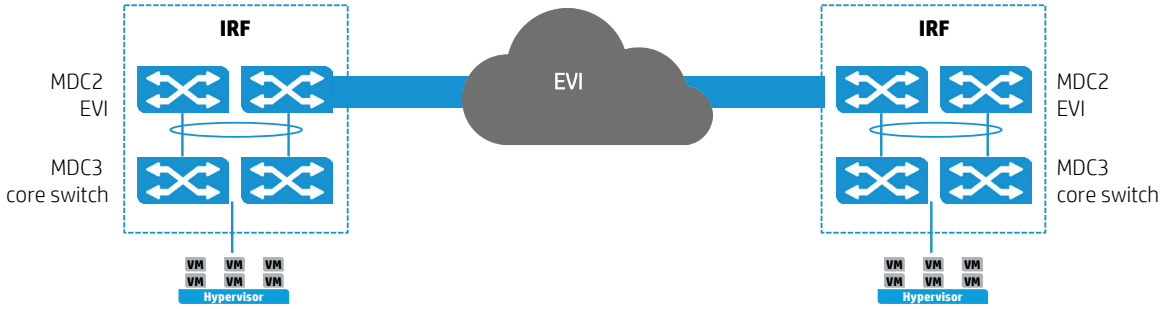

Traffic black hole prevention

To avoid traffic black holes, make sure the MAC aging timer is longer than the EVI ARP entry aging timer (fixed at 25 minutes) on all EDs. It is recommended to set the MAC aging timer to 30 minutes using:

```
mac-address timer aging 1800
```

By default, only 56K MAC addresses are synchronized to a remote ED. If there are more than 56K hosts or MAC addresses connected to the local ED that need to communicate to a remote ED, unique virtual system IDs must be configured on each ED.

```
evi-isis 1
      virtual-system 0000.0000.0010
```

VRRP IP as default gateway for VMs and servers

It is recommended to utilize VRRP IPs as default gateways for VMs and servers in each DC as the virtual MAC address for default gateways will not change even after VMs move across DCs.

Active/Active DC routing implications

Prior to actual deployment, it is recommended a test lab be setup to verify inbound/outbound traffic paths for the new Active/Active DC to perform functions as expected, especially when multiple-routing protocols with different routing protocol preference values are used simultaneously as seen in figure 20.

Figure 20. Example topology with multiple-routing protocols used simultaneously

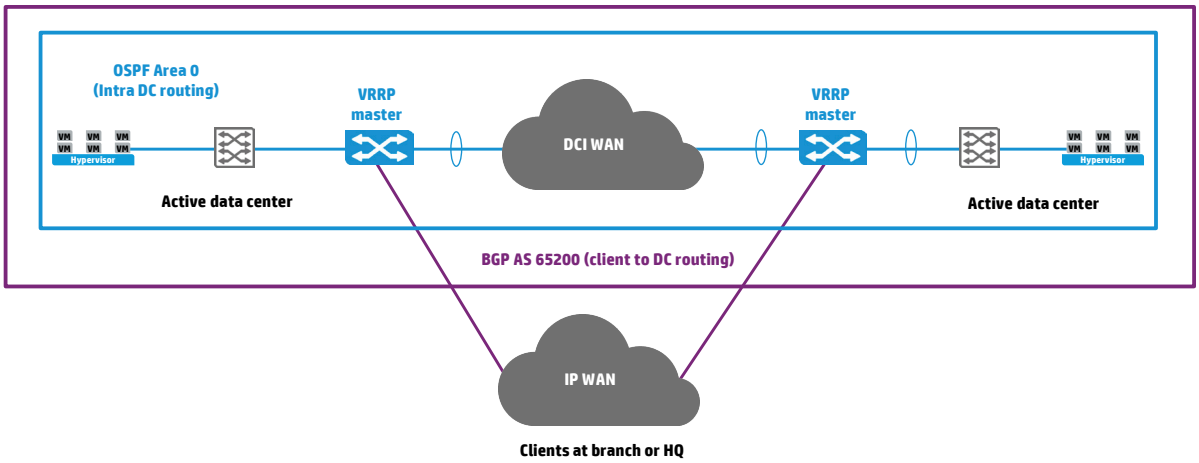

Duplicate VRRP GARP ACLs on 12500

In order to stop "Duplicate VRRP logs" on 12500 due to VRRP Gratuitous Address Resolution Protocol (GARP) between Active/Active DCs as shown in example below:

```
%Jul 22 09:02:06:504 2013 Core-DC ARP/6/DUPVRRPIP:

IP address 192.168.40.254 conflicts with VRRP virtual IP address on interface
Vlan-interface40, sourced from 0000-5e00-0128


%Jul 22 09:02:12:504 2013 Core-DC ARP/6/DUPVRRPIP:

IP address 192.168.40.254 conflicts with VRRP virtual IP address on interface
Vlan-interface40, sourced from 0000-5e00-0128
```

It is recommended to apply outbound ACLs on the VLAN as shown in sample configs below.

The VRRP MAC address to be used in ACL can be found via "dis vrrp verbose" command if the default gateway exists on a Comware device.

```
acl number 4009

 description DENY VRRP

 rule 5 deny type 0806 ffff source-mac 0000-5e00-0101 ffff-ffff-ffff

 rule 100 permit

#

packet-filter 4009 vlan X outbound
```

Appendix: Sample EVI configurations

This section provides sample EVI configurations. Please refer to configuration guides of the appropriate product for more details, URLs are provided at the end of this document.

2 DC EVI deployment

In a deployment with only 2 DCs, the bare minimum configurations required for EVI to function are shown in figure 21. 1 ED will be configured as ENDS using “EVI Neighbor Discovery Server enable” while the opposite ED will be configured as ENDC using “EVI Neighbor Client enable X.X.X.X.”

Figure 21. 2 DC EVI deployment

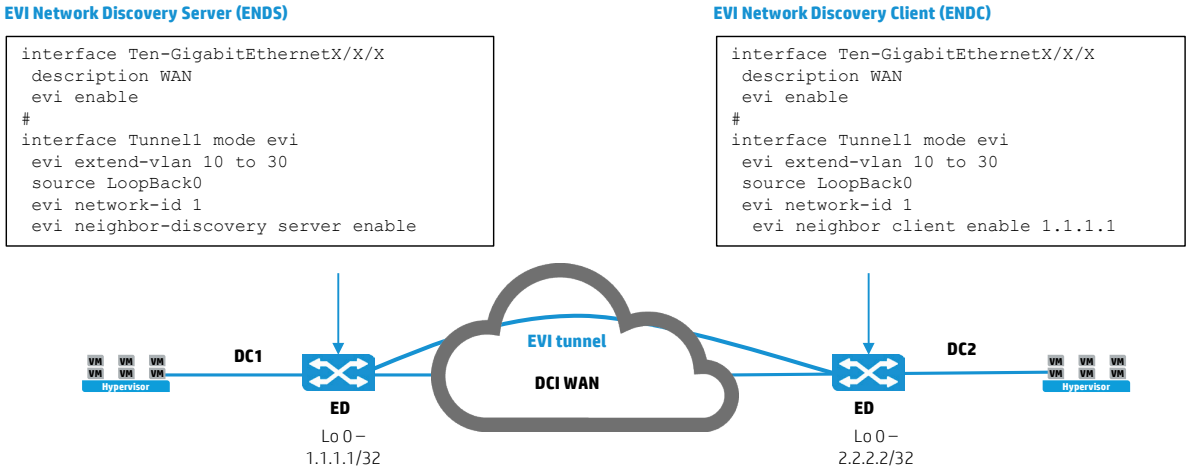

In this 2 DC EVI deployment with future DC expansion planned, the bare minimum configurations required for EVI to function with ENDS redundancy are shown in figure 22. EDs at both DCs are configured as ENDS, they also function as ENDC and need to be configured with the opposite ENDS address.

Figure 22. 2 DC EVI deployment with future expansion plans

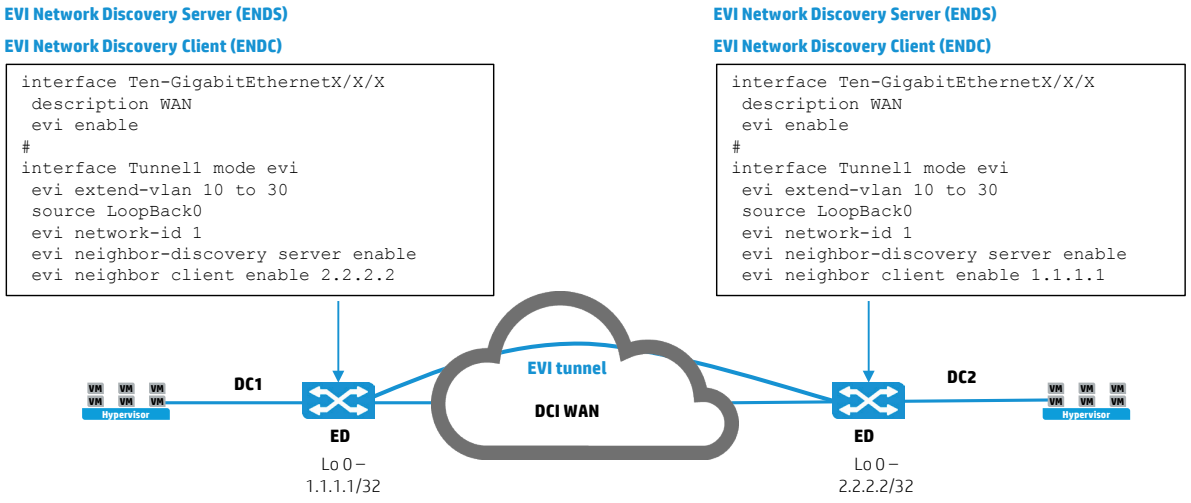

EVI expansion

This example shows how EVI can be expanded to other DCs, the bare-minimum configurations required for EVI to function are shown in figure 23. The same ENDC configuration is used at any new ED that wishes to join the EVI network. New DC additions will not disrupt traffic between existing EVI enabled DCs.

Figure 23. 3 to 16 DC EVI expansion

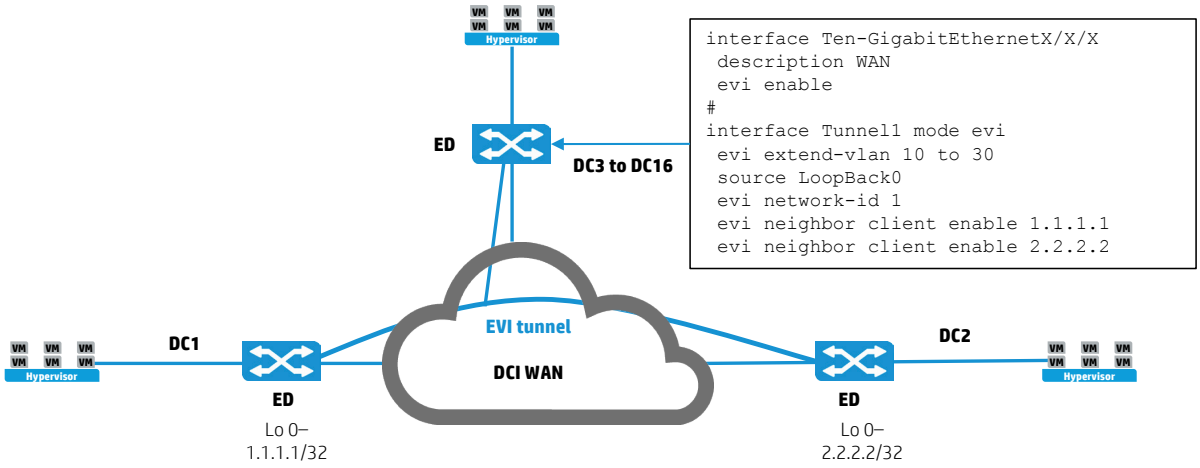

Sample 12500 config with deployment best practices

The example shown in figure 24 has a pair of 12500 switches configured as an IRF fabric, functioning as one ED with sample configurations that include some of the previously discussed deployment best practices.

Figure 24. 12500 IRF EVI with deployment best practices

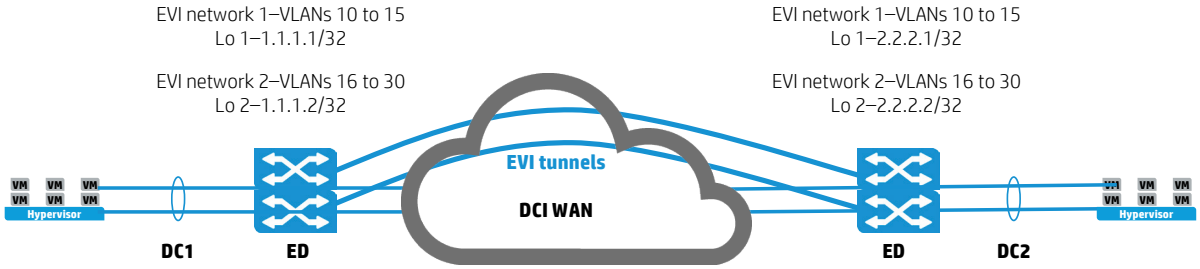


```
sysname DC1-ED
#
 vlan 2
 vlan 10 to 30
#
```

#IRF configs (please refer to 12500 IRF guide for more details)

```
irf member 1 priority 10
 irf member 2 priority 1
#
irf-port 1/1
 port group mdc 1 interface Ten-GigabitEthernet1/1/0/5
 port group mdc 1 interface Ten-GigabitEthernet1/2/0/5
#
irf-port 2/2
 port group mdc 1 interface Ten-GigabitEthernet2/1/0/5
```

```
 port group mdc 1 interface Ten-GigabitEthernet2/2/0/5
#
 stp global enable
#
interface LoopBack1
 description EVI Network 1
 ip address 1.1.1.1 255.255.255.255
#
interface LoopBack2
 description EVI Network 1
 ip address 1.1.1.2 255.255.255.255
#
```

 # IRF related configs to prevent multimaster/split brain (please refer to 12500 IRF guide for more details)

```
interface Vlan-interface2
 description BFD MAD
 mad bfd enable
 mad ip address 192.168.2.1 255.255.255.240 member 1
 mad ip address 192.168.2.2 255.255.255.240 member 2
#
interface GigabitEthernet1/4/0/25
 port link-mode bridge
 description BFD MAD
 port access vlan 2
 undo stp enable
#
interface GigabitEthernet2/4/0/25
 port link-mode bridge
 description BFD MAD
 port access vlan 2
 undo stp enable
#
interface Vlan-interface10
 description Default Gateway for VMs on VLAN 10
 ip address 192.168.10.1 255.255.255.0
 vrrp vrid 1 virtual-ip 192.168.10.254
#
interface Vlan-interface11
 description Default Gateway for VMs on VLAN 11
 ip address 192.168.11.1 255.255.255.0
 vrrp vrid 1 virtual-ip 192.168.11.254
#
```

```
# MLAG to hypervisor
interface Bridge-Aggregation1
 port link-type trunk
 undo port trunk permit vlan 1
 port trunk permit vlan 10 to 30
 link-aggregation mode dynamic
#
interface GigabitEthernet1/4/0/1
 port link-mode bridge
 port link-type trunk
 undo port trunk permit vlan 1
 port trunk permit vlan 10 to 30
# Trust dot1p markings from hypervisor
qos trust dot1p
 port link-aggregation group 1
#
interface GigabitEthernet2/4/0/1
 port link-mode bridge
 port link-type trunk
 undo port trunk permit vlan 1
 port trunk permit vlan 10 to 30
 qos trust dot1p
 port link-aggregation group 1
#
# Remark default dot1p to dscp value for 5
qos map-table inbound dot1p-dscp
 import 5 export 46
#
# Modify queuing parameters as desired
qos qmprofile QOS_QUEUING
 queue be wrr group 2 weight 5
 queue af1 wrr group 2 weight 10
 queue af2 wrr group 1 weight 1
 queue af3 wrr group 1 weight 5
 queue af4 wrr group 1 weight 10
#
interface GigabitEthernet1/4/0/15
 port link-mode route
# Increase MTU towards WAN to cater for EVI overhead
 mtu 9198
 ip address 200.200.200.1 255.255.255.252
# Set link-delay to 0 for faster link failover
link-delay 0 mode updown
```

```
# Trust inbound QoS markings from WAN
 qos trust dscp
# Set outbound queuing mechanism
 qos apply qmprofile QOS_QUEUING
# Enable EVI on physical interface towards WAN
 evi enable
#
interface GigabitEthernet2/4/0/15
 port link-mode route
 mtu 9198
 ip address 200.200.200.5 255.255.255.252
 link-delay 0 mode updown
 qos trust dscp
 qos apply qmprofile QOS_QUEUING
 evi enable
#
# Enable EVI on tunnel 1
interface Tunnel1 mode evi
# Set VLANs to be extended
 evi extend-vlan 10 to 15
# Set source connect interface
 source LoopBack1
# Set EVI network
 evi network-id 1
# Set EVI server or client mode
 evi neighbor-discovery server enable
 evi neighbor-discovery client enable 2.2.2.1
#
# Enable EVI on tunnel 2 with different source/destination connect interfaces to load balance traffic in the WAN
interface Tunnel2 mode evi
 evi extend-vlan 16 to 30
 source LoopBack2
 evi network-id 2
 evi neighbor-discovery server enable
 evi neighbor-discovery client enable 2.2.2.2
#
# Sample configs with BGP AS-path prepending if WAN utilizes BGP
ip prefix-list BGP1 permit 1.1.1.1 32
ip prefix-list BGP2 permit 1.1.1.2 32
#
route-policy BGP1 permit node 1
 if-match ip address prefix-list BGP1
 apply as-path 65001
#
```

```
route-policy BGP2 permit node 1
 if-match ip address prefix-list BGP2
 apply as-path 65001
#
bgp 65001
 peer 200.200.200.2 as-number 65100
 peer 200.200.200.2 bfd
 peer 200.200.200.6 as-number 65100
 peer 200.200.200.6 bfd
 #
 ipv4-family unicast
 network 1.1.1.1 255.255.255.255
 peer 200.200.200.2 enable
 peer 200.200.200.2 route-policy BGP1 export
 peer 200.200.200.6 enable
 peer 200.200.200.6 route-policy BGP2 export
#
# Sample configs with OSPF GR if WAN utilizes OSPF
ospf 1 router-id 1.1.1.1
 graceful-restart ietf
 area 0.0.0.0
 network 1.1.1.1 0.0.0.0
 network 200.200.200.0 0.0.0.255
#
# EVI GR to help minimize traffic loss during IRF master failover
evi-isis 1
 graceful-restart
#
# Modify MAC address timer to prevent black hole
mac-address timer aging 1800
#
# Sample configs to stop duplicate VRRP logs from appearing
acl number 4009
 description DENY VRRP
 rule 5 deny type 0806 ffff source-mac 0000-5e00-0101 ffff-ffff-ffff
 rule 100 permit
packet-filter 4009 vlan 10 to 11 outbound
```

Additional links

For more information, refer to the configuration guides of the specific product.

HP 12500 configuration guides

Learn more at
hp.com/networking

Sign up for updates
hp.com/go/getupdated

Share with colleagues

Rate this document

4AA5-6737ENW, August 2015, Rev. 1