Virtual Connect and HP A-Series switches IRF Integration Guide

Technical white paper

Feedback: vcirf-feedback@hp.com

Table of contents

Introduction

Intelligent Resilient Framework (IRF) is an innovative HP switch platform virtualization technology that allows dramatic simplification of the design and operations of data center and campus Ethernet networks. IRF overcomes the limitations of traditional STP (Spanning Tree Protocol) based and legacy competitive designs by delivering new levels of network performance and resiliency.

Virtual Connect is an industry standard-based implementation of server-edge virtualization. It cleanly separates server enclosure administration from LAN and SAN administration and allows you to add, move, or replace servers without impacting production LAN and SAN availability

This document provides detailed configuration and test information for the following items: (Please note, although A5820 was chosen as the platform of testing but IRF design concept should remain the same for other A-series switches)

IRF and Virtual Connect setup configurations

Failover tests

Images of IMC (Intelligent Management Center) and Insight Control for vCenter network monitoring

Design scenarios

Two typical design scenarios are available to connect Virtual Connect with network switches.

A common misunderstanding people tend to have when connecting Virtual Connect with IRF or Cisco vPC/VSS switches is described in the following page. The design does not work.

The above concepts apply to all Virtual Connect models providing ethernet connectivity, which include VC 1/10-F, VC Flex-10 and VC Flexfabric modules.

<table>
<tr>
<td> 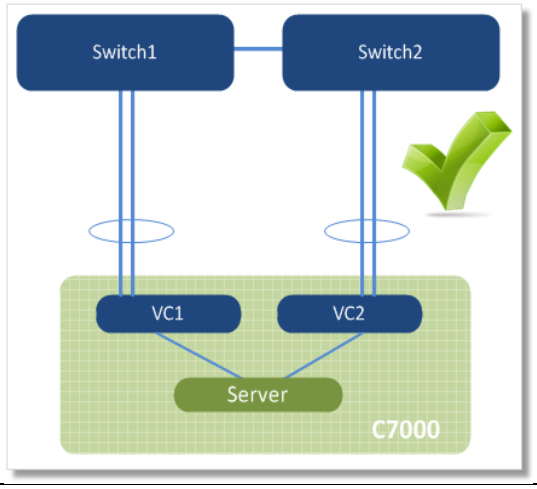
</td>
<td> 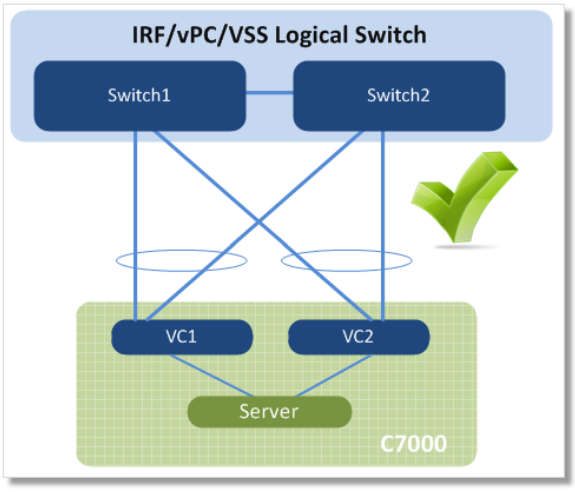
</td>
</tr>
<tr>
<td>Scenario 1—This is a typical connection scenario, in which Virtual Connect modules connect with non-IRF/vPC/VSS capable switches.

Virtual Connect needs to configure one SUS (Shared Uplink Set) per Virtual Connect module (two total). Switch 1 and switch 2 each have one port channel configured to peer with Virtual Connect SUS.</td>
<td>Scenario 2—This is the recommended connection scenario, in which Virtual Connect modules connect with IRF/vPC/VSS logical switch.

Virtual Connect needs to configure one SUS per Virtual Connect module (two total). The logical switch also has two port channels configured to peer with Virtual Connect SUS, which is known as Active/Active Virtual Connect design. Active/Standby Virtual Connect design is also available, but because it does not use all available uplink bandwidth, it is not discussed here in more detail. For more information on Active/Standby design, see scenario1:4 in the HP Virtual Connect Ethernet Cookbook (http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01990371/c01990371.pdf).

This design provides two main benefits over the previous design:
• If either switch fails, traffic remains on the same port channel and rehashes to the remaining physical link in less than one second. The server does not require failover tests.
• For the incoming traffic from upstream core switch to server direction, all traffic can be sent to Virtual Connect. Previously, if the destination MAC (media access control) was on the other switch, the traffic would have to traverse the inter-switch trunk, so the flow was not optimized.</td>
</tr>
</table>

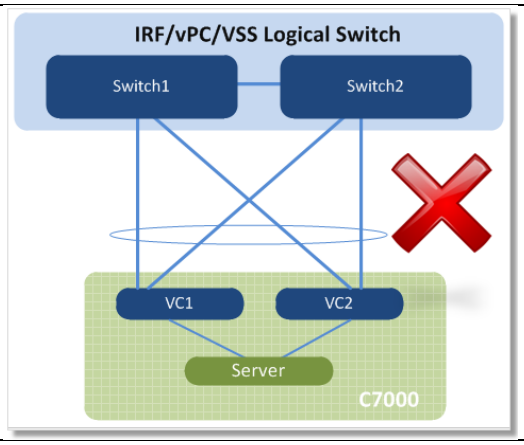

Scenario 3—This configuration does not work. Configuring one port channel on a logical switch side and one SUS on a Virtual Connect side does not move traffic forward on all four links. Virtual Connect does not support port channels across different modules. Some links will go into standby and not form port channels.
See **Appendix 2** (on page 44) for the results of this scenario.

Network topology

Physical diagram

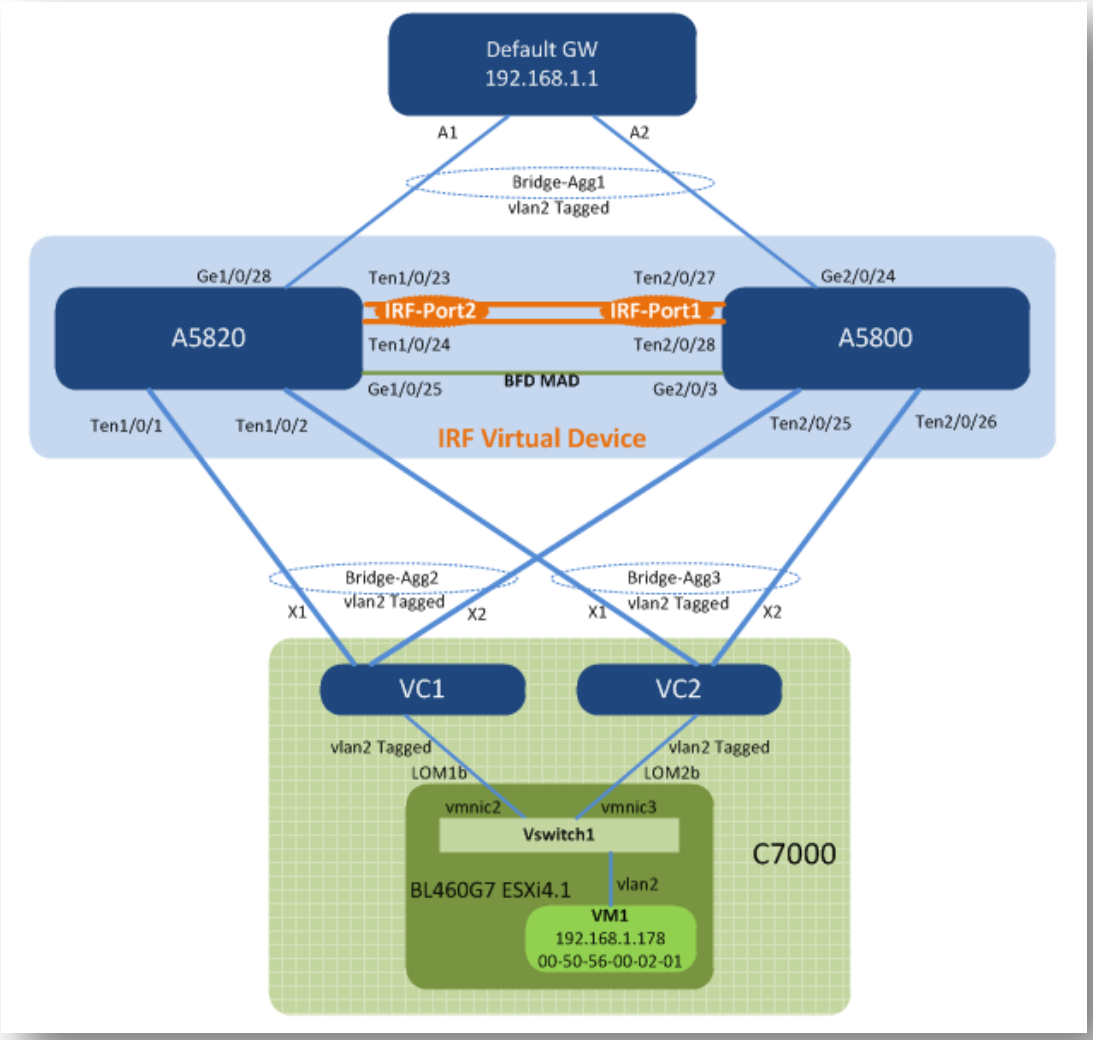

The IRF cluster consists of one A5820 switch and one A5800-32C switch. Comware software supports IRF clustering on different switch models if they are compatible with each other for IRF.

The A5820 and A5800 switches form an IRF bundle link between them with two 10G links. The A5820 switch is switch 1, the master of the domain, and has logical port IRF-Port2. The A5800 switch is switch 2, the slave of the domain, and has logical port IRF-Port1, defined originally before merging with the A5820 switch.

The A5820 and A5800 switches use one Gigabit Ethernet link as a BFD MAD link for MAD.

VC1 and VC2 are Flex-10 modules in interconnect bays 1 and 2 of the HP BladeSystem c7000 Enclosure. Each Flex-10 module has a SUS connecting to an IRF virtual device. A SUS consists of two 10G links terminated on A5820 and A5800 switches. With IRF, these two 10G links form one bridge-aggregation bundle (the same as port channel on Cisco NX-OS and etherchannel on Cisco IOS). VC1 connects the IRF cluster with the Bridge-Aggregation 2 interface, and VC2 connects the IRF cluster with the Bridge-Aggregation 3 interface. Bridge-Aggregation 1 forms a virtual port channel

between the IRF cluster and the virtual machine's default gateway (simulated by an HP E-Series switch).

Traffic flow testing uses ping packets from VM1 (192.168.1.178) to its default gateway (192.168.1.1). The VM traffic has two paths to reach its default gateway, depending on how the vSwitch hashes VM traffic to a specific vmnic.

Logical diagram

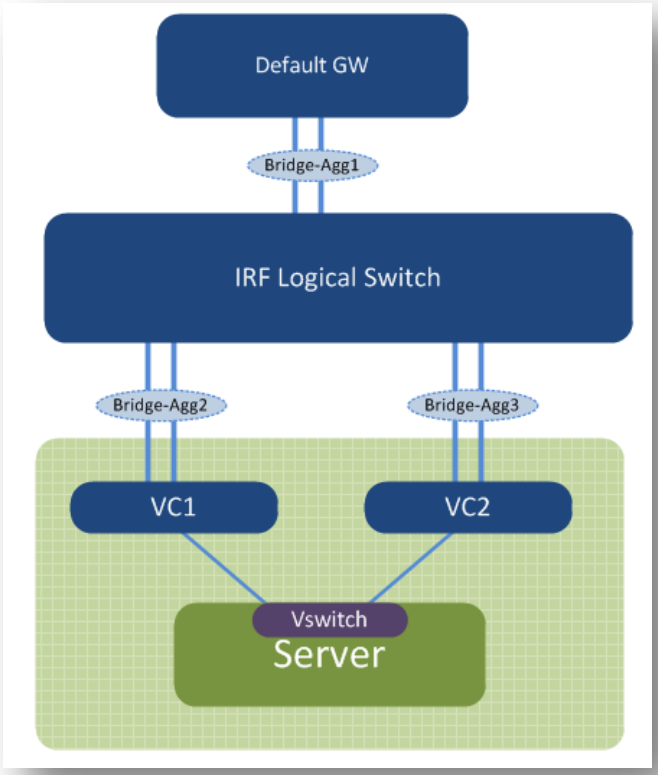

Two bundle interfaces (Bridge-Aggregation 2 and Bridge-Aggregation 3) exist between the Virtual Connect and the IRF logical switch because Virtual Connect currently does not support link bundling across two different physical modules.

IRF and MAD technology overview

IRF (Intelligent Resilient Framework)

IRF	IRF creates one logical switch from two or more physical switches. The A5820 switch can support up to nine switches in one IRF domain. The logical switch uses standard LACP to connect to any vendor, core, distribution, or edge switches with a failure convergence time of less than 40 milliseconds. The switch acts as the following: • Single IP address for management • Single layer 2 switch • Single layer 3 router (all protocols) Implementation is available across multiple products from core to access platforms A12500, A10500 A9500, A7500, A5820, A5800, and A5500 series switches. With IRF technology, the network is transformed as shown in the following diagram. 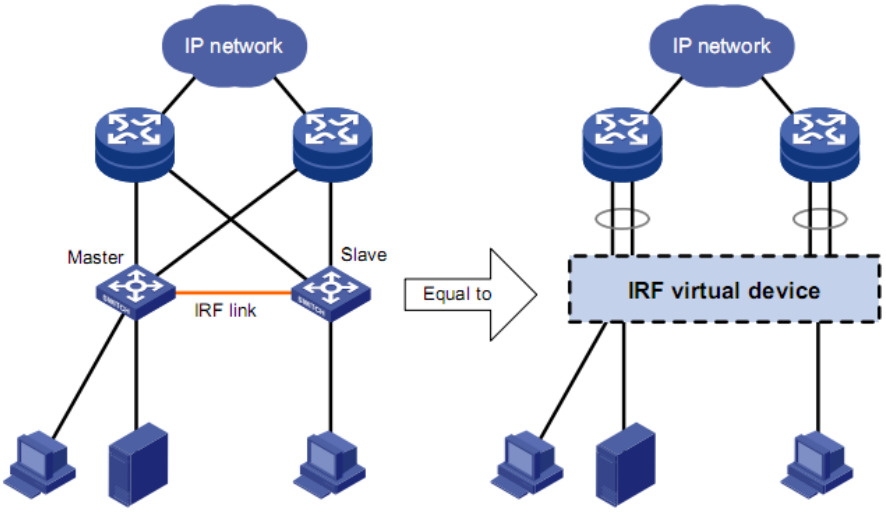
Role	Member switches form an IRF virtual device. Each of them performs one of the following two roles: • Master—manages the IRF virtual device • Subordinate—members that are backups of the master If the master fails, the IRF virtual device automatically elects a new master from one of the subordinates. Masters and subordinates are elected through the role election mechanism. An IRF virtual device has only one master at a time.
IRF port	An IRF port is a logical port dedicated to the internal connection of an IRF virtual device. An IRF port can be numbered as IRF-port1 or IRF-port2. An IRF port is effective only after it is bound to a physical port. **Important:** An IRF-Port1 on one device can only be connected to the physical port bound to the IRF-Port2 of a neighboring device; otherwise, an IRF virtual device cannot be formed.

	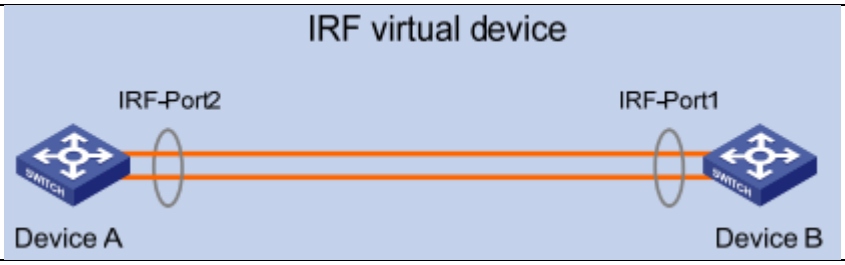
Physical IRF port	Physical IRF ports are physical (copper or fiber) ports bound to an IRF port. They perform the following functions: • Connect IRF member switches • Forward IRF protocol packets and data packets between IRF member switches
Priority	Member priority determines the role of a member during a role election process. A member with a higher priority is more likely to be a master. The priority of a switch defaults to 1.
Member ID	An IRF virtual device uses member IDs to uniquely identify its members. Configuration Information such as port (physical or logical) numbers, port configurations, and member priorities relate to member IDs.
Domain ID	Each switch belongs to one IRF domain. By default, the domain ID is 0. Although switches with different domain IDs can form an IRF virtual device, HP recommends assigning the same domain ID to the members of the same IRF virtual device. Otherwise, the LACP MAD detection cannot function properly.

MAD (Multi-Active Detection)

MAD	MAD protects IRF link failure when both switches with the same configuration meet the criteria for master switch. In this case, MAD shuts down one of the switches according to role election. The switch with a higher priority becomes the master, and then the local interfaces for switch 2 are shut down. When an IRF link is down as a result of MAD, switch 1 continues to run. Switch 2 inactivates all local interfaces. 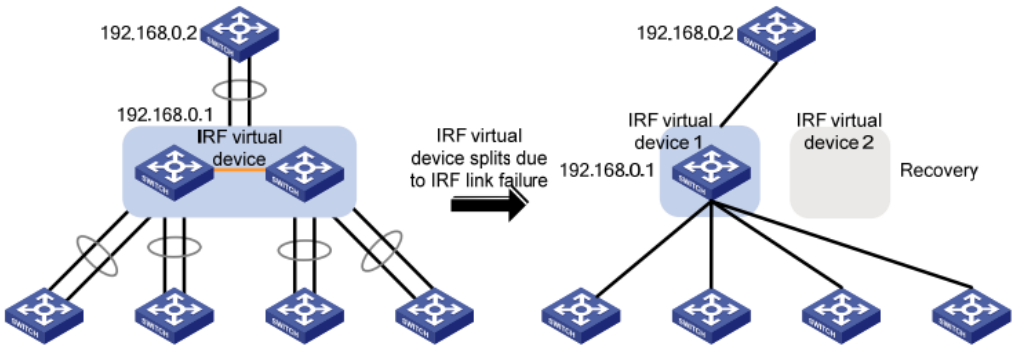 MAD detects multiple active IRFs using one of three methods: • LACP • BFD • ARP
LACP MAD	• Most widely deployed • Fastest convergence time • Needs only one CLI "MAD enable" under bridge aggregation interface • Needs a third switch (Typically HPN A-series) to understand extended LACPDU (Link Aggregation Control Protocol Data Unit) packets

<table>
<tr><td></td><td> 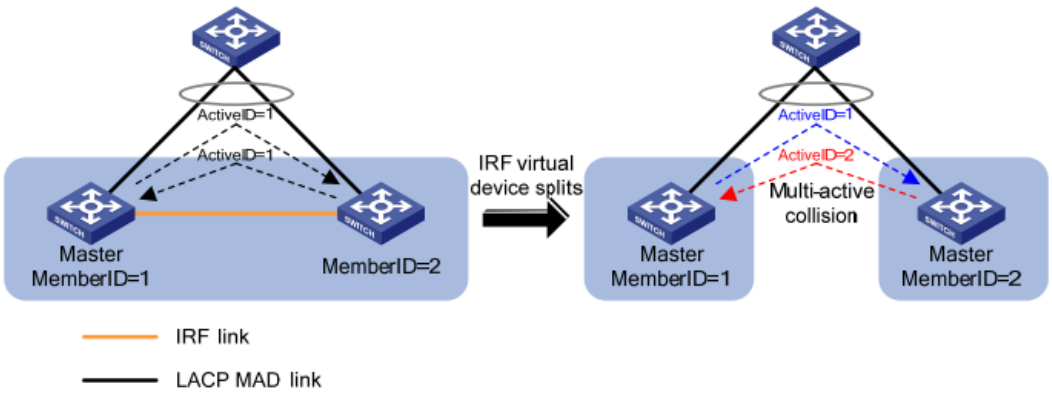
</td></tr>
<tr><td>BFD MAD</td><td>

Fast convergence time
Needs a separate link between two switches to act as a BFD MAD link
Does not require switches outside the IRF domain

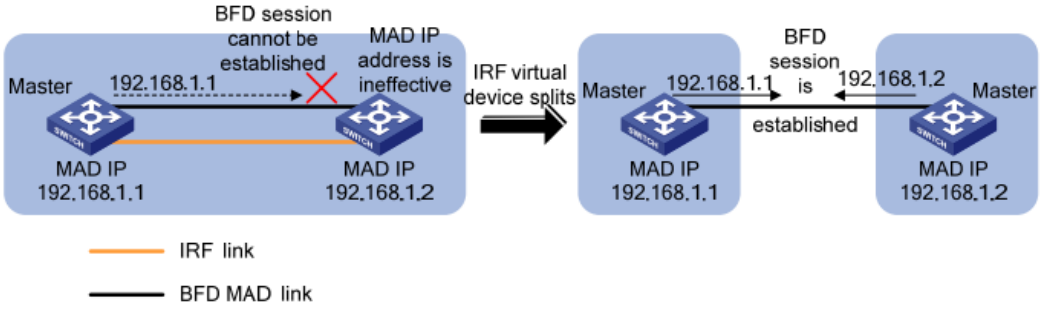

</td></tr>
<tr><td>ARP MAD</td><td>This MAD is not widely deployed. For more information, see the IRF configuration guide.
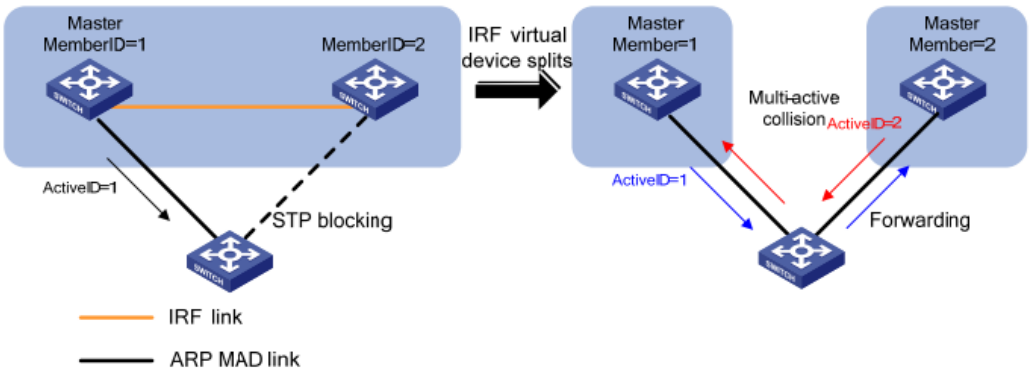
</td></tr>
</table>

For more information on IRF and MAD, see the *H3C S5820X & S5800 Series Ethernet Switches IRF Configuration Guide*.
http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c02648772/c02648772.pdf

IRF and Virtual Connect setup configurations

Quick CLI reference table

HPN A-Series Comware CLI is similar to the Cisco IOS/NX-OS format. The following table gives a quick comparison of A-Series Comware CLI and Cisco CLI, related to this setup.

Comware	Cisco
system	config terminal
undo	no
quit	exit
save force	wr mem
reset saved-config	wr erase
reboot	reload
display current	show run
display saved-configuration	show startup
display int brief	show ip int brief
display logbuffer	show log
display link-aggregation	show etherchannel/port-channel
display this (show current interface config)	
sysname	hostname
port link-mode bridge	switchport
port link-mode route	no switchport
port link-type access	switchport mode access
port link-type trunk	switchport mode trunk
port access vlan x	switchport access vlan x
port trunk permit vlan x	switchport trunk allowed vlan x
port link-aggregation group x	channel-group x
interface Bridge-Aggregation x	int port-channel x

A5820 switch: Convert standalone switches to IRF logical switch

This conversion procedure assumes that two standalone switches start from a clean factory-default startup configuration. If not, enter **reset saved-config** (**write erase** on Cisco) to reset startup-config to factory default.

A5820 (switch 1)

1. Change switch 1 IRF priority to 10**.** The default value is 1, and the higher priority is selected to be the IRF master and active switch when MAD is detected.
[H3C]**irf member 1 priority 10**
2. Shut down the IRF physical ports to prepare them to be included under the IRF logical port "irf-port 1/2" configuration. Otherwise, when trying to include these interfaces later under IRF-Port, Comware will indicate that the physical interfaces are not shut down.
[H3C]**int ten1/0/23** [H3C-Ten-GigabitEthernet1/0/23]**shut** [H3C-Ten-GigabitEthernet1/0/23]**int ten1/0/24** [H3C-Ten-GigabitEthernet1/0/24]**shut**
3. Create Logical port "irf-port 1/2" and include ten1/0/23 and ten1/0/24. Note: If you create "irf-port 1/2" on switch 1, you must use "irf-port 2/1" on switch 2. Alternatively, create local "irf-port 1/1" and use "irf-port 2/2" on switch 2. The following two scenarios do not work: • "irf-port 1/1"--- "irf-port 2/1" • "irf-port 1/2"---"irf-port 2/2"
[H3C]**irf-port 1/2** [H3C-irf-port1/2]**port group interface ten1/0/23** [H3C-irf-port1/2]**port group interface ten1/0/24**
4. While ten1/0/23 and ten1/0/24 are shut down, go to Switch 2 (page 12) to configure it to peer with Switch 1. Then, complete the remaining steps in this procedure.
5. Unshut ten1/0/23 and ten1/0/24 to bring up the irf-link. After the links and interfaces appear, proceed to the next step. Nothing happens until step 6 is executed.
[H3C]**int ten1/0/23** [H3C-Ten-GigabitEthernet1/0/23]**undo shut** [H3C-Ten-GigabitEthernet1/0/23]**int ten1/0/24** [H3C-Ten-GigabitEthernet1/0/24]**undo shut**
6. Activate the irf-port configuration to start IRF peering between the two switches.
[H3C]**irf-port-configuration active**
After several seconds, switch 2 reloads. When switch 2 comes back on, two switches are merged into one virtual IRF switch. You can use the three IRF commands to verify the running status for this virtual IRF switch. See the output following A5800 (switch 2).

A5800 (switch 2)

1. Change switch 2 member ID from default 1 to 2.

```
[H3C]irf member 1 renumber 2
```

2. Before continuing with the following steps, reboot the switch to make all interface numbering changes from 1/x/y to 2/x/y. This command is executed when the switch is not in system mode.

```
<H3C>reboot
```

After rebooting

3. Shut down the IRF physical ports to prepare them to be included under the IRF logical port "irf-port 2/1" configuration. Otherwise, when trying to include these interfaces later under IRF-Port, Comware will indicate that the physical interfaces are not shut down.

```
[H3C]int ten2/0/27
[H3C-Ten-GigabitEthernet2/0/27]shut
[H3C-Ten-GigabitEthernet2/0/27]int ten2/0/28
[H3C-Ten-GigabitEthernet2/0/28]shut
```

4. Create Logical port "irf-port 2/1" and include ten2/0/27 and ten2/0/28.

```
[H3C]irf-port 2/1
[H3C-irf-port2/1]port group interface ten2/0/27
[H3C-irf-port2/1]port group interface ten2/0/28
```

5. Unshut ten2/0/27 and ten2/0/28 to bring up the irf-link. After the links and interfaces appear, proceed to the next step. Nothing happens until step 6 is executed.

```
[H3C]int ten2/0/27
[H3C-Ten-GigabitEthernet2/0/27]undo shut
[H3C-Ten-GigabitEthernet2/0/27]int ten2/0/28
[H3C-Ten-GigabitEthernet2/0/28]undo shut
```

6. Activate irf port configuration to start IRF peering between two switches. At this moment, nothing happens because both switch 1 IRF physical links are still shut down.

```
[H3C]irf-port-configuration active
```

7. Go to Switch 1 (page 11) to start IRF physical links and activate the IRF-link configuration. Several seconds later, switch 2 reloads itself with the message below (only part of the booting message is shown here for reference).

```
IRF port 1 is up.

Starting......


    ****************************************************************************

    *                                                                          *

    *                 H3C S5800-32C BOOTROM, Version 205                       *

    *                                                                          *

    ****************************************************************************
        Copyright (c) 2004-2010 Hangzhou H3C Technologies Co., Ltd.
```

After merging, IRF status checks the output. For the complete logical switch configuration, see Appendix 2 (on page 44).

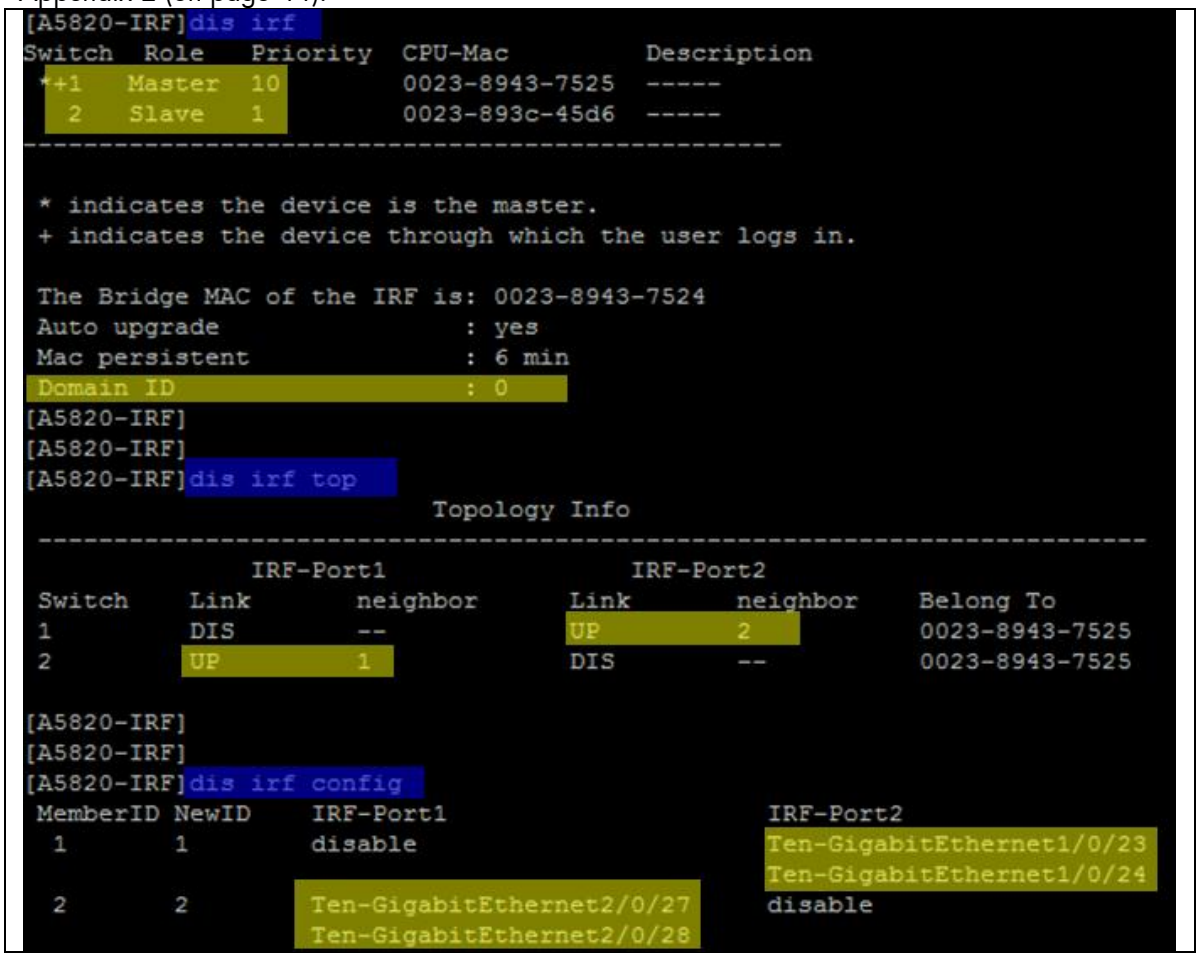

A5820: BFD MAD configuration

```
#
vlan 100
#
interface vlan-interface100
 mad bfd enable
 mad ip address 100.100.100.1 255.255.255.0 member 1
 mad ip address 100.100.100.2 255.255.255.0 member 2
#
interface GigabitEthernet1/0/25
 port link-mode bridge
 port access vlan 100
 stp disable
#
interface GigabitEthernet2/0/3
 port link-mode bridge
 port access vlan 100
 stp disable
```

```
[A5820-IRF-Vlan-interface100]dis mad verbose
Current MAD status: Detect
Excluded ports(configurable):
Excluded ports(can not be configured):
  Ten-GigabitEthernet1/0/23
  Ten-GigabitEthernet1/0/24
  Ten-GigabitEthernet2/0/27
  Ten-GigabitEthernet2/0/28
MAD ARP disabled.
MAD LACP disabled.
MAD BFD enabled interface:
  Vlan-interface100
    mad ip address 100.100.100.1 255.255.255.0 member 1
    mad ip address 100.100.100.2 255.255.255.0 member 2
```

To disable STP for the BFD MAD interface, issue the command **`stp disable`**. The BFD MAD interface is a dedicated interface and should not run any other services/features.

A5820: LLDP

LLDP (Link-layer Discovery Protocol) is the IEEE standard protocolused by network devices for advertising their identity, capabilities, and neighbors. LLDP performs functions similar to some other proprietary protocols, such as Cisco Discovery Protocol(CDP).

LLDP transmits and receives are enabled by default on A5820 interfaces. No configuration is required.

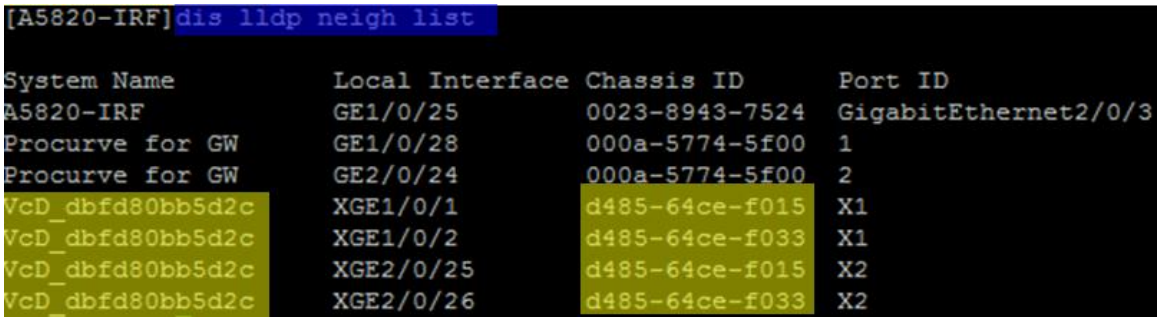

```
[A5820-IRF]dis lldp neigh list

System Name          Local Interface Chassis ID      Port ID
A5820-IRF            GE1/0/25        0023-8943-7524  GigabitEthernet2/0/3
Procurve for GW      GE1/0/28        000a-5774-5f00  1
Procurve for GW      GE2/0/24        000a-5774-5f00  2
VcD_dbfd80bb5d2c     XGE1/0/1        d485-64ce-f015  X1
VcD_dbfd80bb5d2c     XGE1/0/2        d485-64ce-f033  X1
VcD_dbfd80bb5d2c     XGE2/0/25       d485-64ce-f015  X2
VcD_dbfd80bb5d2c     XGE2/0/26       d485-64ce-f033  X2
```

The "VcD_xyz" string is the unique Virtual Connect domain ID generated internally when creating Virtual Connect. VC1 and VC2 share the same LLDP "System Name" because they are in the same Virtual Connect domain. To determine which physical Virtual Connect module is the LLDP neighbor, use the "Chassid ID" field. This is the Virtual Connect module system MAC address. To determine the system MAC address for a particular Virtual Connect module, log into Virtual Connect by SSH (Secure Shell) and use the **`show interconnect`** command.

```
->show interconnect enc0:1
ID                : enc0:1
Enclosure         : oa8
Bay               : 1
Type              : VC-ENET
Product Name      : HP VC Flex-10 Enet Module
Role              : Primary
Status            : OK
Comm Status       : OK
OA Status         : OK
Power State       : On
MAC Address       : d4:85:64:ce:f0:15
Node WWN          : -- --
Firmware Version  : 3.15 2010-10-09T07:18:16Z
Manufacturer      : HP
Part Number       : 455880-B21
Spare Part Number : 456095-001
Rack Name         : R8-9-10
Serial Number     : 3C4031000B
UID               : Off
```

Flex-10: LLDP

LLDP transmits and receives are enabled by default on all Virtual Connect modules interfaces, including Flex-10 and Flexfabric. No configuration is required.
Trunk-A and Trunk-B are defined in the following LACP sections. All links will show as active only after finishing the LACP configuration on the switch and Virtual Connect.

VC1 connects with IRF logical switch ten1/0/1 and ten2/0/25.

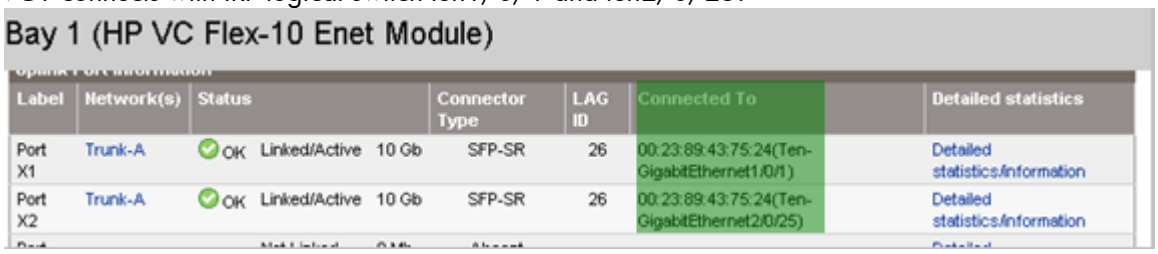

Bay 1 (HP VC Flex-10 Enet Module)

Label	Network(s)	Status	Connector Type	LAG ID	Connected To	Detailed statistics
Port X1	Trunk-A	OK Linked/Active 10 Gb	SFP-SR	26	00:23:89:43:75:24(Ten-GigabitEthernet1/0/1)	Detailed statistics/information
Port X2	Trunk-A	OK Linked/Active 10 Gb	SFP-SR	26	00:23:89:43:75:24(Ten-GigabitEthernet2/0/25)	Detailed statistics/information

VC2 connects with IRF logical switch ten1/0/2 and ten2/0/26.

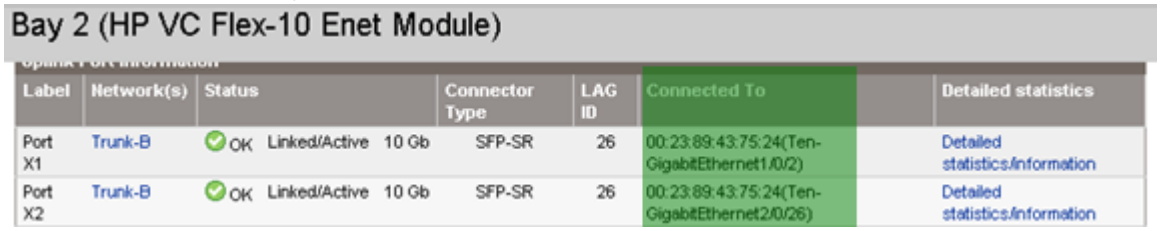

Bay 2 (HP VC Flex-10 Enet Module)

Label	Network(s)	Status	Connector Type	LAG ID	Connected To	Detailed statistics
Port X1	Trunk-B	OK Linked/Active 10 Gb	SFP-SR	26	00:23:89:43:75:24(Ten-GigabitEthernet1/0/2)	Detailed statistics/information
Port X2	Trunk-B	OK Linked/Active 10 Gb	SFP-SR	26	00:23:89:43:75:24(Ten-GigabitEthernet2/0/26)	Detailed statistics/information

A5820: LACP

The Bridge-Aggregation interface is equal to the port channel interface on Cisco to bundle multiple physical links.

```
interface Bridge-Aggregation2
 port link-type trunk
 port trunk permit vlan 1 to 2
 link-aggregation mode dynamic
 stp edged-port enable
#
interface Bridge-Aggregation3
 port link-type trunk
 port trunk permit vlan 1 to 2
 link-aggregation mode dynamic
 stp edged-port enable
#
interface Ten-GigabitEthernet1/0/1
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 to 2
 port link-aggregation group 2
#
interface Ten-GigabitEthernet1/0/2
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 to 2
 port link-aggregation group 3
#
interface Ten-GigabitEthernet2/0/25
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 to 2
 port link-aggregation group 2
#
interface Ten-GigabitEthernet2/0/26
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 to 2
 port link-aggregation group 3
#
```

When connecting with Virtual Connect, the Spanning Tree edge ports (Cisco PortFast) feature should be enabled because Virtual Connect does not communicate STP with any network device. The command is **`stp edged-port enable`** under the interface. This can speed up network convergence time, especially when links come up.

The BPDU (Bridge Protocol Data Unit) guard feature can be enabled for more security to protect edge ports. The global command is **`stp bpdu-protection`**.

These practices are in line with networking best design when connecting with host NICs. Networking switches should treat any ports connecting with Virtual Connect as the ports connecting with regular servers.

Bridge-Aggregation interfaces commands

```
[A5820-IRF]dis link-aggregation verbose b2

Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Port Status: S -- Selected, U -- Unselected
Flags:  A -- LACP_Activity, B -- LACP_Timeout, C -- Aggregation,
        D -- Synchronization, E -- Collecting, F -- Distributing,
        G -- Defaulted, H -- Expired

Aggregation Interface: Bridge-Aggregation2
Aggregation Mode: Dynamic
Loadsharing Type: Shar
System ID: 0x8000, 0023-8943-7524
Local:
  Port             Status  Priority Oper-Key  Flag
--------------------------------------------------------------------------------
  XGE1/0/1         S       32768    1         {ACDEF}
  XGE2/0/25        S       32768    1         {ACDEF}
Remote:
  Actor            Partner Priority Oper-Key  SystemID               Flag
--------------------------------------------------------------------------------
  XGE1/0/1         17      1        3         0x1   , d485-64ce-f015 {ABCDEF}
  XGE2/0/25        18      1        3         0x1   , d485-64ce-f015 {ABCDEF}
[A5820-IRF]dis link-aggregation verbose b3

Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Port Status: S -- Selected, U -- Unselected
Flags:  A -- LACP_Activity, B -- LACP_Timeout, C -- Aggregation,
        D -- Synchronization, E -- Collecting, F -- Distributing,
        G -- Defaulted, H -- Expired

Aggregation Interface: Bridge-Aggregation3
Aggregation Mode: Dynamic
Loadsharing Type: Shar
System ID: 0x8000, 0023-8943-7524
Local:
  Port             Status  Priority Oper-Key  Flag
--------------------------------------------------------------------------------
  XGE1/0/2         S       32768    2         {ACDEF}
  XGE2/0/26        S       32768    2         {ACDEF}
Remote:
  Actor            Partner Priority Oper-Key  SystemID               Flag
--------------------------------------------------------------------------------
  XGE1/0/2         17      1        7         0x1   , d485-64ce-f033 {ABCDEF}
  XGE2/0/26        18      1        7         0x1   , d485-64ce-f033 {ABCDEF}
```

Flex-10: LACP

Trunk uplink config on VC1

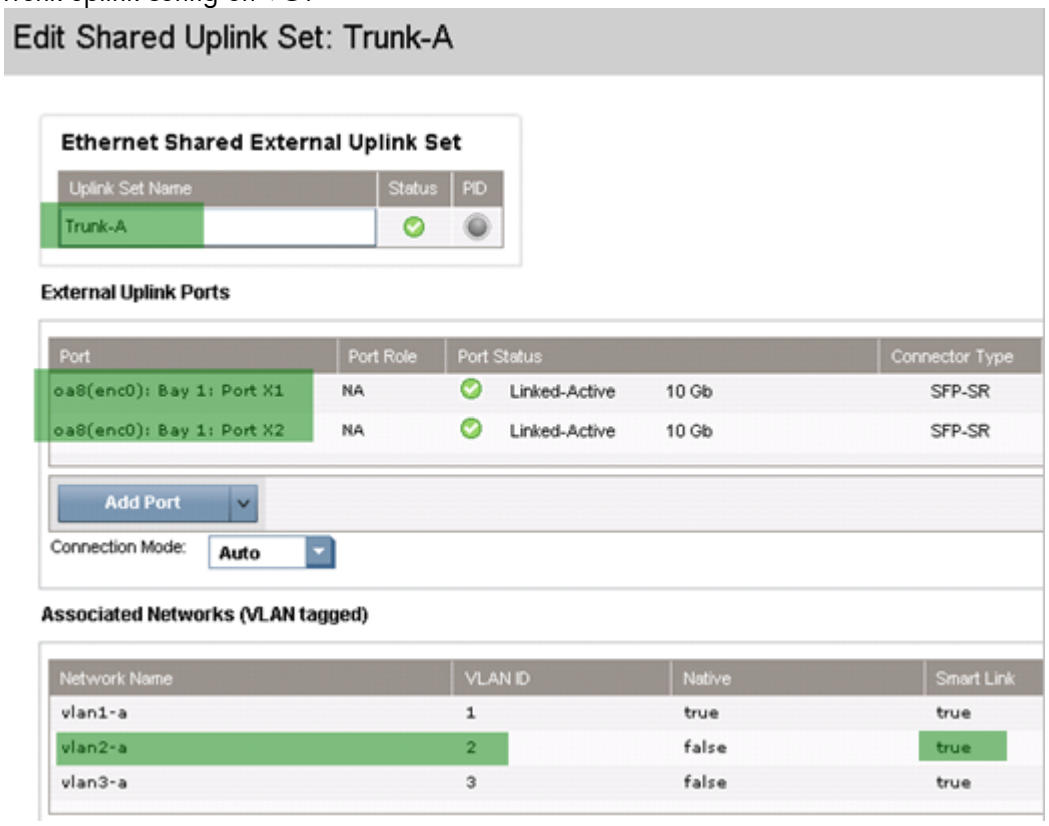

Trunk uplink config on VC2

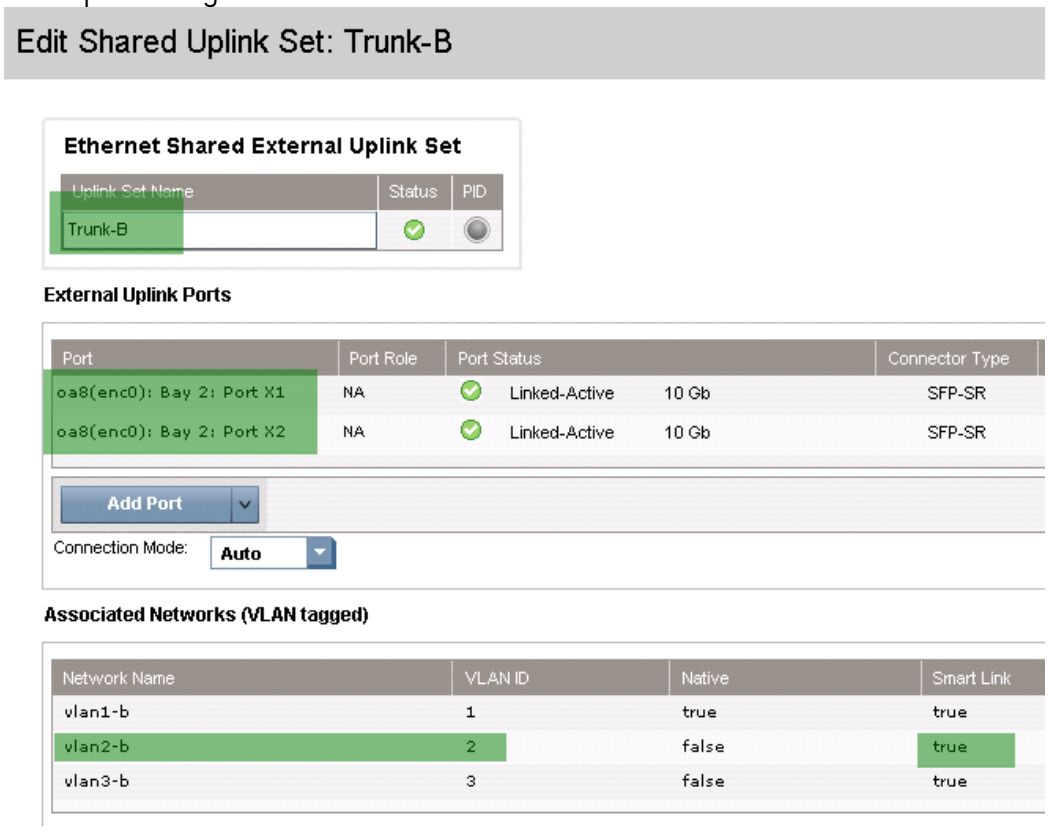

Trunk uplinks monitoring on Virtual Connect

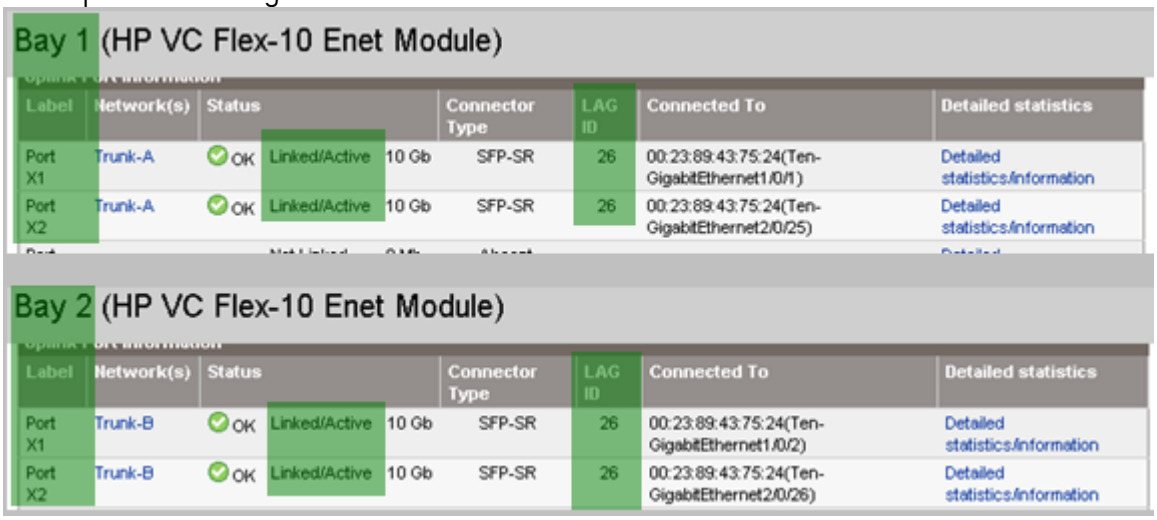

Bay 1 (HP VC Flex-10 Enet Module)

Label	Network(s)	Status		Connector Type	LAG ID	Connected To	Detailed statistics
Port X1	Trunk-A	OK	Linked/Active	10 Gb SFP-SR	26	00:23:89:43:75:24(Ten-GigabitEthernet1/0/1)	Detailed statistics/information
Port X2	Trunk-A	OK	Linked/Active	10 Gb SFP-SR	26	00:23:89:43:75:24(Ten-GigabitEthernet2/0/25)	Detailed statistics/information

Bay 2 (HP VC Flex-10 Enet Module)

Label	Network(s)	Status		Connector Type	LAG ID	Connected To	Detailed statistics
Port X1	Trunk-B	OK	Linked/Active	10 Gb SFP-SR	26	00:23:89:43:75:24(Ten-GigabitEthernet1/0/2)	Detailed statistics/information
Port X2	Trunk-B	OK	Linked/Active	10 Gb SFP-SR	26	00:23:89:43:75:24(Ten-GigabitEthernet2/0/26)	Detailed statistics/information

Both trunks show active/active. Also LAG (Link Aggregation Group) ID shows that a LACP bundle has been established with IRF Virtual Switch. Both channels use LAG 26. Since they are on different modules, Virtual Connect can uniquely identify them.

Flex-10: Server Profile

Server profile configuration

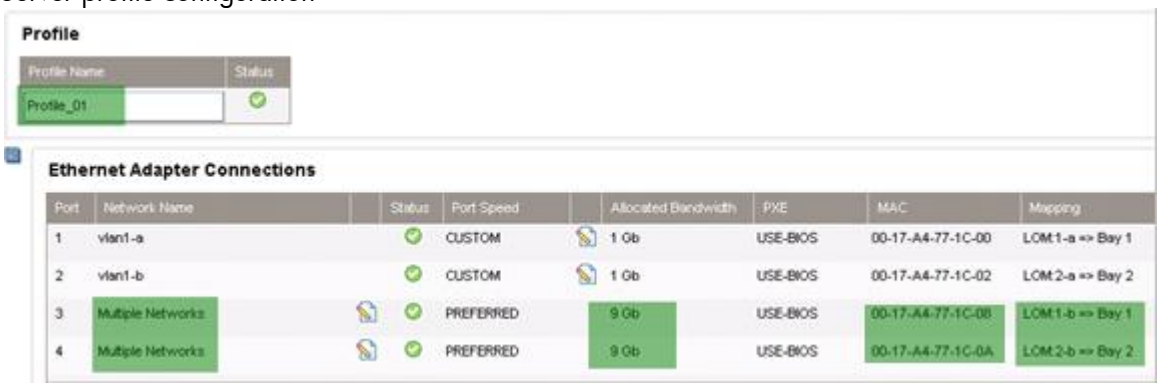

Port 3 “Multiple Networks” configuration

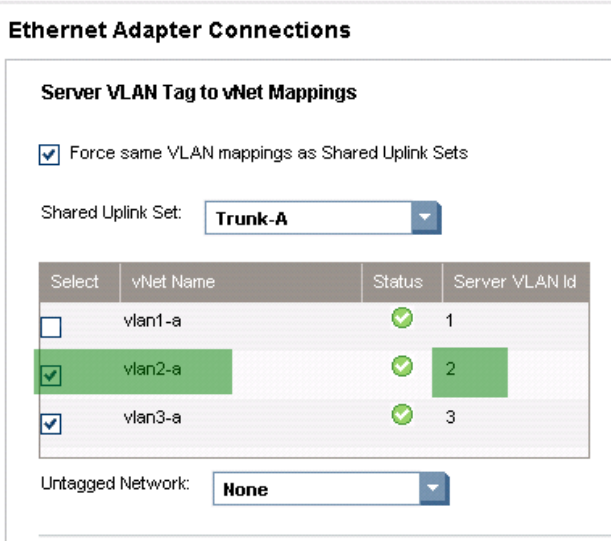

Port 4 “Multiple Networks” configuration

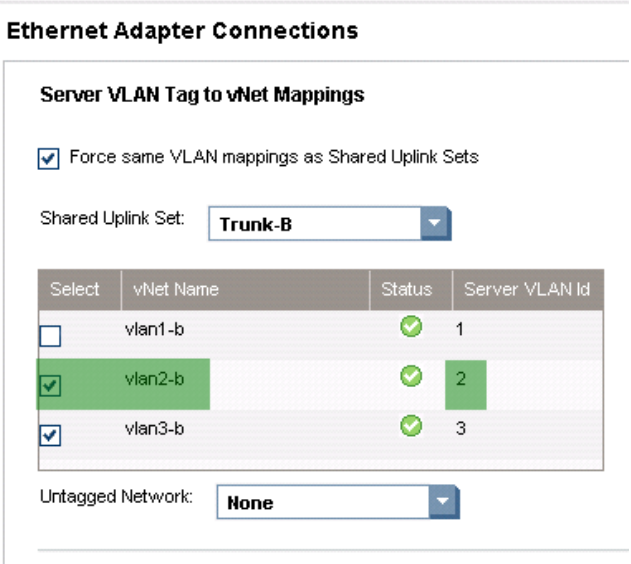

ESXi configuration

Host adapter

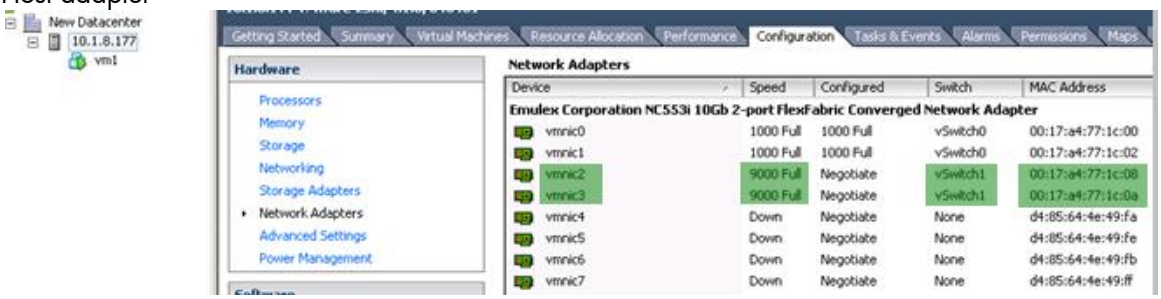

Switch1 configuration

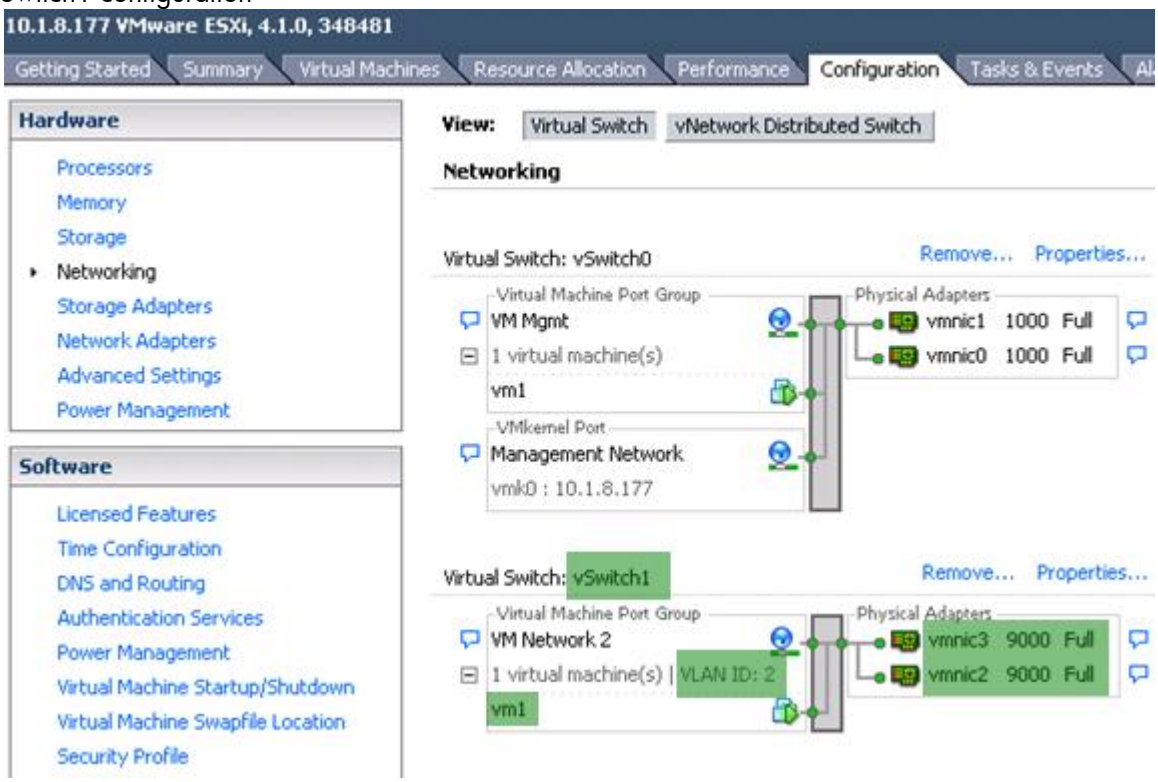

VM1 network adapter configuration for VLAN 2

```
Ethernet adapter Local Area Connection:

   Connection-specific DNS Suffix  . :
   Description . . . . . . . . . . . : Intel(R) PRO/1000 MT Network Connection
   Physical Address. . . . . . . . . : 00-50-56-00-02-01
   DHCP Enabled. . . . . . . . . . . : No
   Autoconfiguration Enabled . . . . : Yes
   Link-local IPv6 Address . . . . . : fe80::a94a:8645:dfcc:3343%10(Preferred)
   IPv4 Address. . . . . . . . . . . : 192.168.1.178(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . :
   DNS Servers . . . . . . . . . . . : fec0:0:0:ffff::1%1
                                       fec0:0:0:ffff::2%1
                                       fec0:0:0:ffff::3%1
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

Failover tests

Uplink failure

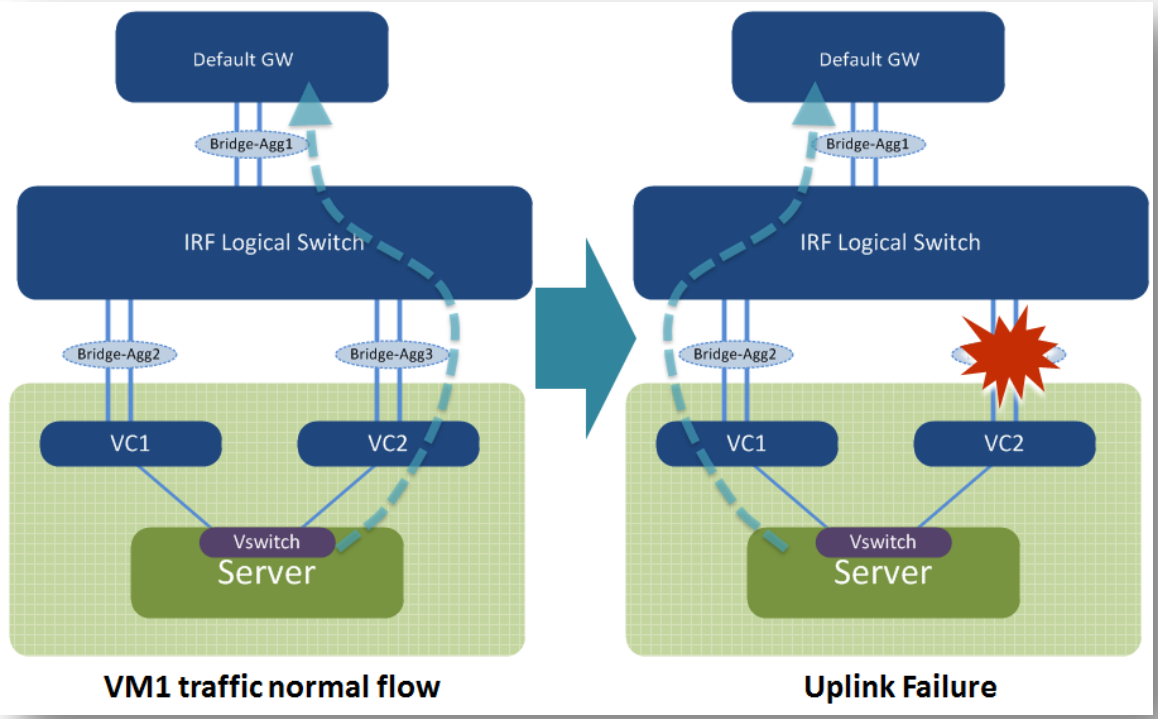

VM1 has a continuous ping to its default GW 192.168.1.1. Under normal conditions, vSwitch hashes the traffic from this VM to the vmnic3, which is mapped to the VC2 and then enters the Bridge-Aggregate3 interface in the IRF logical switch.

The test issued a `shut down` command under interface b3. From the `display MAC address` command, we can see the traffic failed over to the other path.

Test Result:

- Shut down int b3: about 3-4 seconds packets loss.
- Undo shut int b3: about 1-2 seconds packets loss with "stp edged-port enable." Without it, about 30 seconds of packet loss occurs due to the regular STP learning stage.

Note:

IRF convergence time is much faster than three seconds, typically less than 50 milliseconds. The overall three second convergence time is related to Virtual Connect convergence around the smartlink to notify the server link in the event of uplink downtime, which then triggers vSwitch to converge the packet flow. Even with a regular switch without IRF (verified in the lab), three seconds is the expected Virtual Connect/vSwitch convergence time in similar topology.

Shut int b3

```
<A5820-IRF>
<A5820-IRF>dis mac-address dynamic vlan 2
MAC ADDR       VLAN ID  STATE          PORT INDEX               AGING TIME(s)
000a-5774-5f01 2        Learned        Bridge-Aggregation1      AGING
0050-5600-0201 2        Learned        Bridge-Aggregation3      AGING

  ---  2 mac address(es) found  ---

<A5820-IRF>int b3
           ^
 % Unrecognized command found at '^' position.
<A5820-IRF>sys
System View: return to User View with Ctrl+Z.
[A5820-IRF]int b3
[A5820-IRF-Bridge-Aggregation3]shut
[A5820-IRF-Bridge-Aggregation3]dis mac-add dynamic vlan 2
MAC ADDR       VLAN ID  STATE          PORT INDEX               AGING TIME(s)
000a-5774-5f01 2        Learned        Bridge-Aggregation1      AGING
0050-5600-0201 2        Learned        Bridge-Aggregation2      AGING

  ---  2 mac address(es) found  ---
```

```
2011/03/11 12:18:55 : Reply[1785] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
2011/03/11 12:18:55 : Reply[1786] from 192.168.1.1: bytes=32 time=1.6 ms TTL=64
2011/03/11 12:18:55 : Reply[1787] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
2011/03/11 12:18:55 : Reply[1788] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
2011/03/11 12:18:56 : Reply[1789] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
2011/03/11 12:18:57 : 192.168.1.1: request timed out
2011/03/11 12:18:58 : 192.168.1.1: request timed out
2011/03/11 12:18:59 : 192.168.1.1: request timed out
2011/03/11 12:19:00 : 192.168.1.1: request timed out
2011/03/11 12:19:00 : Reply[1794] from 192.168.1.1: bytes=32 time=1.6 ms TTL=64
2011/03/11 12:19:00 : Reply[1795] from 192.168.1.1: bytes=32 time=1.9 ms TTL=64
2011/03/11 12:19:00 : Reply[1796] from 192.168.1.1: bytes=32 time=1.4 ms TTL=64
2011/03/11 12:19:00 : Reply[1797] from 192.168.1.1: bytes=32 time=1.4 ms TTL=64
2011/03/11 12:19:00 : Reply[1798] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
```

Undo shut int b3

```
[A5820-IRF-Bridge-Aggregation3]dis mac-add dynamic vlan 2
MAC ADDR       VLAN ID  STATE          PORT INDEX               AGING TIME(s)
000a-5774-5f01 2        Learned        Bridge-Aggregation1      AGING
0050-5600-0201 2        Learned        Bridge-Aggregation3      AGING

  ---  2 mac address(es) found  ---

[A5820-IRF-Bridge-Aggregation3]int b3
[A5820-IRF-Bridge-Aggregation3] shut
[A5820-IRF-Bridge-Aggregation3]dis mac-add dynamic vlan 2
MAC ADDR       VLAN ID  STATE          PORT INDEX               AGING TIME(s)
000a-5774-5f01 2        Learned        Bridge-Aggregation1      AGING
0050-5600-0201 2        Learned        Bridge-Aggregation2      AGING

  ---  2 mac address(es) found  ---

[A5820-IRF-Bridge-Aggregation3]undo shut
[A5820-IRF-Bridge-Aggregation3]dis mac-add dynamic vlan 2
MAC ADDR       VLAN ID  STATE          PORT INDEX               AGING TIME(s)
000a-5774-5f01 2        Learned        Bridge-Aggregation1      AGING
0050-5600-0201 2        Learned        Bridge-Aggregation3      AGING
```

```
2011/03/11 12:26:45 : Reply[482] from 192.168.1.1: bytes=32 time=1.4 ms TTL=64
2011/03/11 12:26:46 : 192.168.1.1: request timed out
2011/03/11 12:26:47 : 192.168.1.1: request timed out
2011/03/11 12:26:47 : Reply[485] from 192.168.1.1: bytes=32 time=1.8 ms TTL=64
2011/03/11 12:26:47 : Reply[486] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
```

Switch failure

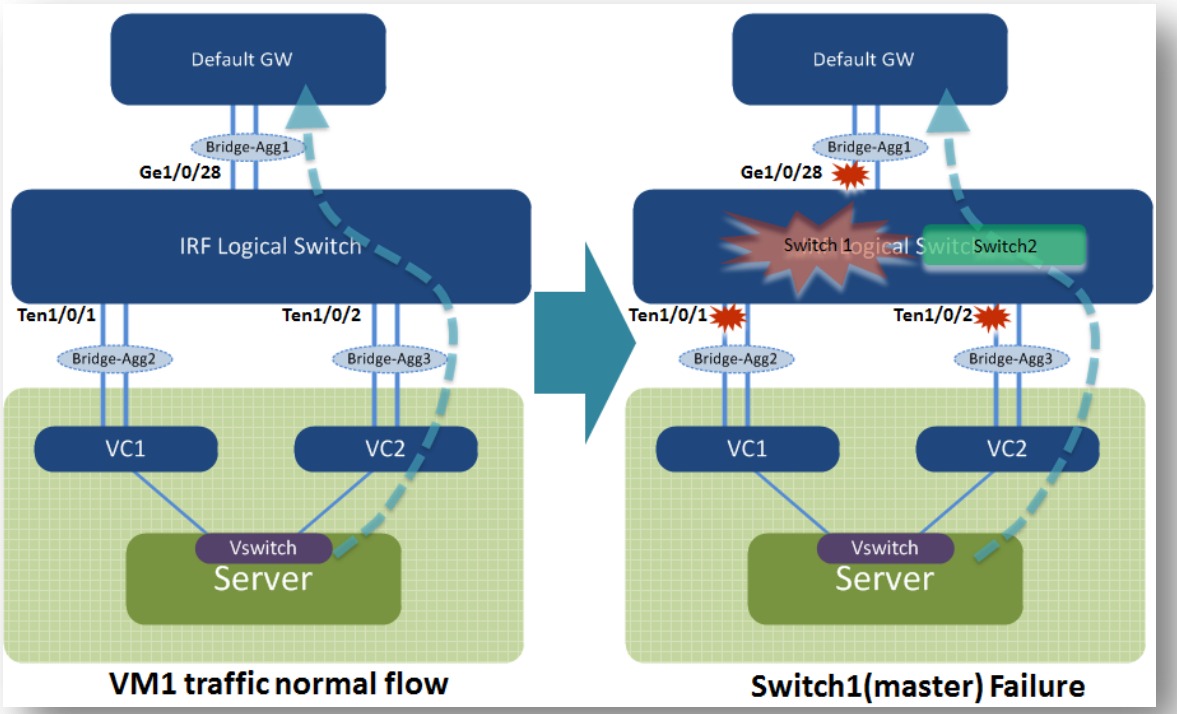

VM1 has a continuous ping to its default GW 192.168.1.1. Under normal conditions, vSwitch hashes the traffic from this VM to the vmnic3, which is mapped to VC2 and then enters the Bridge-Aggregation3 interface in the IRF logical switch.

The test issues a `reboot` command on Switch1 A5820. Switch 2 takes over as the new master and any interface related to 1/y/z is shut down.

Test Result:
Switch1 down: Ping packet loss did not occur, so the convergence time was less than one second.
Switch1 up: Ping packet loss did not occur, so the convergence time was less than one second.

Note:
The convergence time remained less than one second because the traffic flow did not switch over to the other path. It still used int b3 because even with switch 1 and all 1/y/z interfaces down, int b3 still had the other interface ten2/0/26 up. So, the convergence time is the result of LACP rehashing the traffic to the other remaining link, which is typically less than one second.

For this scenario, IRF does not change the traffic flow path, even when losing one switch. The two uplinks operate at 10G each.

After switch 1 comes back up, it remains the slave to prevent traffic switch-over again, even though it has higher priority.

```
<A5820-IRF>dis irf
Switch  Role    Priority  CPU-Mac         Description
   1    Slave   10        0023-8943-7525  -----
 *+2    Master  1         0023-893c-45d6  -----
-------------------------------------------------

 * indicates the device is the master.
 + indicates the device through which the user logs in.

 The Bridge MAC of the IRF is: 0023-8943-7524
 Auto upgrade                : yes
 Mac persistent              : 6 min
 Domain ID                   : 0
```

IRF link failure

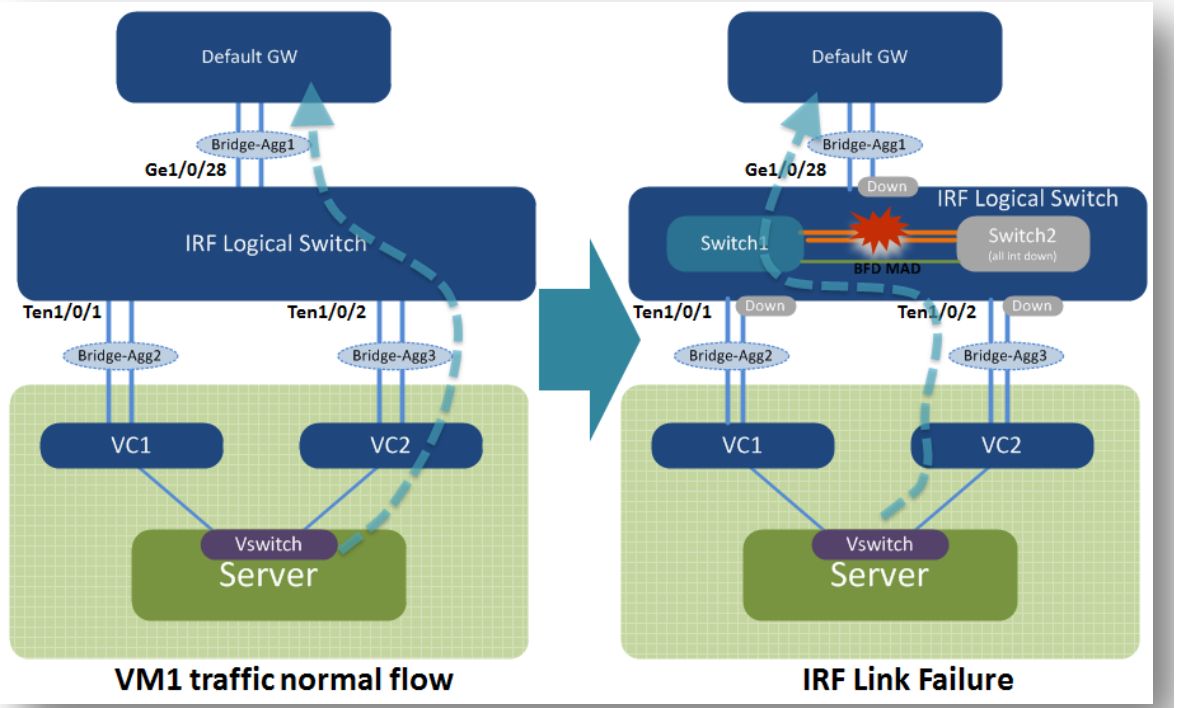

VM1 has a continuous ping to its default GW 192.168.1.1. Under normal conditions, the vSwitch hashes the traffic from this VM to the vmnic3, which is mapped to the VC2 and then enters the Bridge-Aggregation3 interface in the IRF logical switch.

The test issued a shut down command under switch1 A5820 IRF1/2 to simulate IRF link failure.

Test Result:

- Shut irf-port 1/2: Ping packet loss did not occur, so the convergence time was less than one second.
- No shut irf-port 1/2: About one second packet loss after switch 2 rebooted and came back up to join the IRF domain

Note:

Upon losing the IRF link, MAD initiates and elects one master for the domain, and the other switch (switch 2 with lower IRF priority) shuts down all its local interfaces to prevent a dual active (split brain) scenario. When the IRF link is restored, switch 2 reboots itself and rejoins the IRF domain.

Packet loss when Switch2 (A5800) came back and joined IRF domain:

```
2011/03/11 14:14:19 : Reply[1694] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
2011/03/11 14:14:19 : Reply[1695] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
2011/03/11 14:14:20 : 192.168.1.1: request timed out
2011/03/11 14:14:20 : Reply[1697] from 192.168.1.1: bytes=32 time=2.6 ms TTL=64
2011/03/11 14:14:20 : Reply[1698] from 192.168.1.1: bytes=32 time=1.6 ms TTL=64
```

Switch2 (A5800) view after IRF link failure with BFD MAD protection

```
<A5820-IRF>dis irf
Switch   Role    Priority  CPU-Mac         Description
 *+2     Master  1         0023-893c-45d6  -----
-------------------------------------------------

 * indicates the device is the master.
 + indicates the device through which the user logs in.

 The Bridge MAC of the IRF is: 0023-8943-7524
 Auto upgrade                : yes
 Mac persistent              : 6 min
 Domain ID                   : 0
<A5820-IRF>dis mad ver
Current MAD status: Recovery
Excluded ports(configurable):
Excluded ports(can not be configured):
  Ten-GigabitEthernet2/0/27
  Ten-GigabitEthernet2/0/28
MAD ARP disabled.
MAD LACP disabled.
MAD BFD enabled interface:
  Vlan-interface100
    mad ip address 100.100.100.1 255.255.255.0 member 1
    mad ip address 100.100.100.2 255.255.255.0 member 2
<A5820-IRF>
```

Switch2 (A5800) view after all local interfaces were shut down to prevent a dual active scenario

```
<A5820-IRF>dis int brief
The brief information of interface(s) under route mode:
Link: ADM - administratively down; Stby - standby
Protocol: (s) - spoofing
Interface            Link Protocol Main IP         Description
M-GE0/0/0            DOWN DOWN     10.1.8.2
NULL0                DOWN UP(s)    --
Vlan100              DOWN DOWN     --

The brief information of interface(s) under bridge mode:
Link: ADM - administratively down; Stby - standby
Speed or Duplex: (a)/A - auto; H - half; F - full
Type: A - access; T - trunk; H - hybrid
Interface            Link Speed   Duplex Type PVID Description
BAGG1                DOWN auto    A      T    1
BAGG2                DOWN auto    A      T    1
BAGG3                DOWN auto    A      T    1
GE2/0/1              DOWN auto    A      A    1
GE2/0/2              DOWN auto    A      A    1
GE2/0/3              DOWN auto    A      A    100
GE2/0/4              DOWN auto    A      A    1
GE2/0/5              DOWN auto    A      A    1
GE2/0/6              DOWN auto    A      A    1
GE2/0/7              DOWN auto    A      A    1
GE2/0/8              DOWN auto    A      A    1
GE2/0/9              DOWN auto    A      A    1
GE2/0/10             DOWN auto    A      A    1
GE2/0/11             DOWN auto    A      A    1
GE2/0/12             DOWN auto    A      A    1
GE2/0/13             DOWN auto    A      A    1
GE2/0/14             DOWN auto    A      A    1
GE2/0/15             DOWN auto    A      A    1
GE2/0/16             DOWN auto    A      A    1
GE2/0/17             DOWN auto    A      A    1
GE2/0/18             DOWN auto    A      A    1
GE2/0/19             DOWN auto    A      A    1
GE2/0/20             DOWN auto    A      A    1
GE2/0/21             DOWN auto    A      A    1
GE2/0/22             DOWN auto    A      A    1
GE2/0/23             DOWN auto    A      A    1
GE2/0/24             DOWN auto    A      T    1
XGE2/0/25            DOWN auto    A      T    1
XGE2/0/26            DOWN auto    A      T    1
XGE2/0/27            DOWN --      --     --   --
XGE2/0/28            DOWN --      --     --   --
```

Switch1 (A5820) console log after IRF-link fail

Virtual Connect module failure

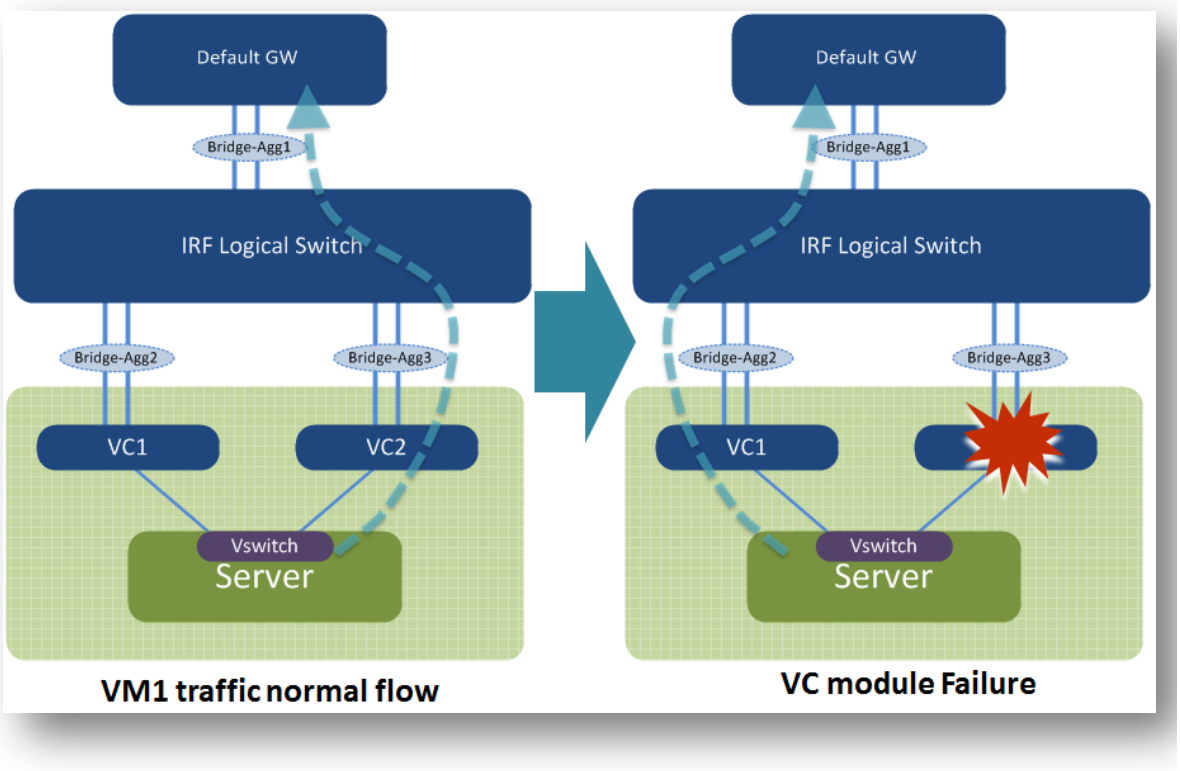

VM1 has a continuous ping to its default GW 192.168.1.1. Under normal conditions, the vSwitch hashes the traffic from this VM to the vmnic3, which is mapped to the VC2 and enters the Bridge-Aggregation 3 interface in IRF the logical switch.

The test uses the "power off" button on OA (Onboard Administrator) to shut down VC2 to simulate module failure.

Test Result:

- VC2 down: About one second packet loss
- VC2 up: About six seconds packet loss

(Please note: VC 3.30 will have the enhancement to reduce the convergence time to less than 1 sec upon VC module coming up. VC3.30 is currently scheduled to be available by the end of 08/2011.)

Note:

The VC2 up event had more convergence time because the vmnic3, which is mapped to VC2, was up. Therefore, the vSwitch started to send traffic to VC2 before VC2 was ready internally for switching traffic.

VC2 down

```
2011/03/16 11:31:52 : Reply[986] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
2011/03/16 11:31:53 : Reply[987] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
2011/03/16 11:31:54 : Reply[988] from 192.168.1.1: bytes=32 time=1.4 ms TTL=64
2011/03/16 11:31:55 : Reply[989] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
2011/03/16 11:31:57 : 192.168.1.1: request timed out
2011/03/16 11:31:57 : Reply[991] from 192.168.1.1: bytes=32 time=1.4 ms TTL=64
2011/03/16 11:31:58 : Reply[992] from 192.168.1.1: bytes=32 time=1.5 ms TTL=64
```

VC2 up

With VC 3.30 enhancement, VC2 up event will be reduced to around 500 msec. Please note below fping timeout and interval were set to 500 msec. In the test below, please regard 192.168.1.2 as 192.168.1.1 in previous tests due to different IP scheme in different lab setup.

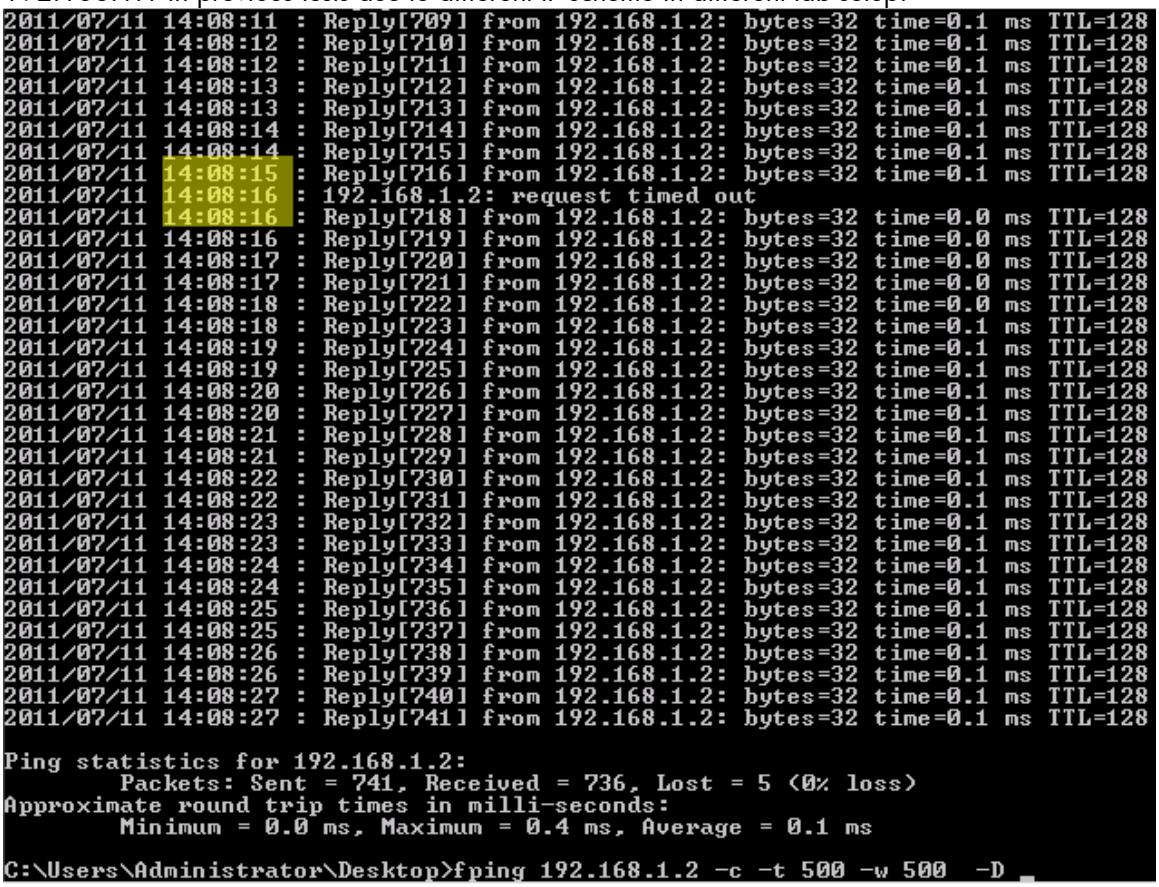

Insight Control for Vmware vCenter monitoring
Insight Control for vCenter utilizes a visual networking view from vSwitch to Virtual Connect to physical access switch. The following images provide examples of its appearance and functionality.

VM1 uses vSwitch1, which has two uplinks (vmnic2 and vmnic3). The uplinks carry tagged packets for VLAN (Virtual Local Area Network) 2 and VLAN 3. VLAN 3 is not in used in the testing but is provided to show the concept of tagged trunking between Virtual Connect and vSwitch. The graphic also displays the physical uplink ports used to connect to the access switch. The host name and MAC address of that switch are also provided, and are obtained through the use of LLDP between Virtual Connect and the network switch.

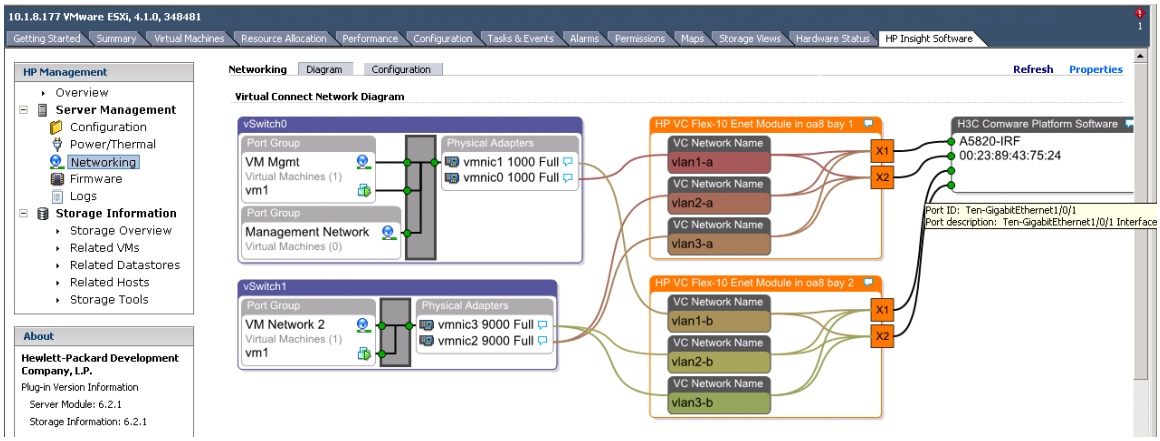

Host H/W inventory details

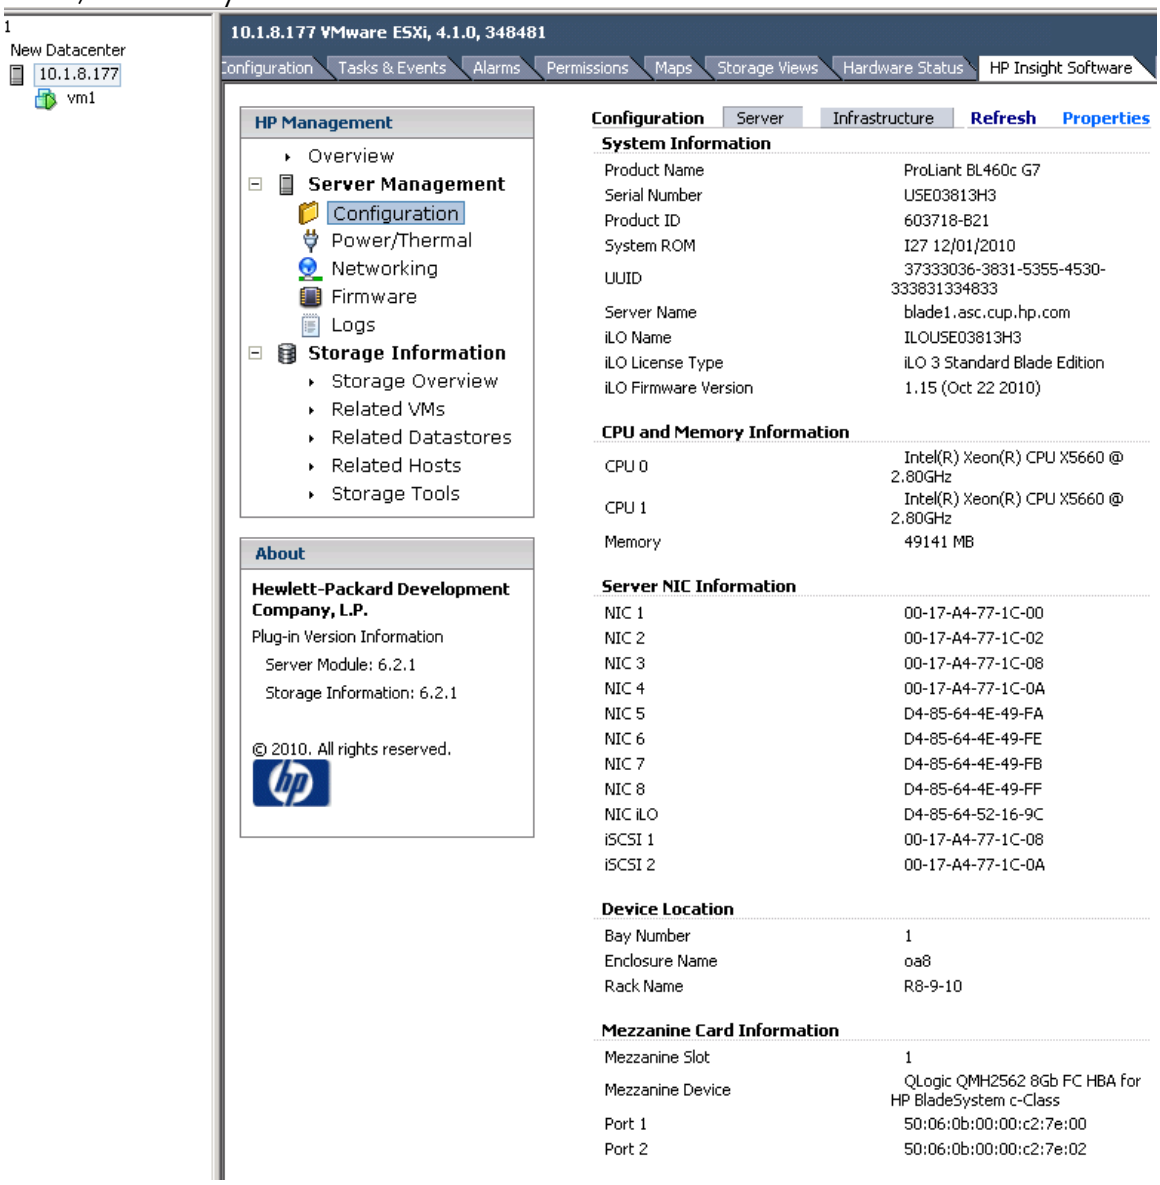

Host and Enclosure firmware version report

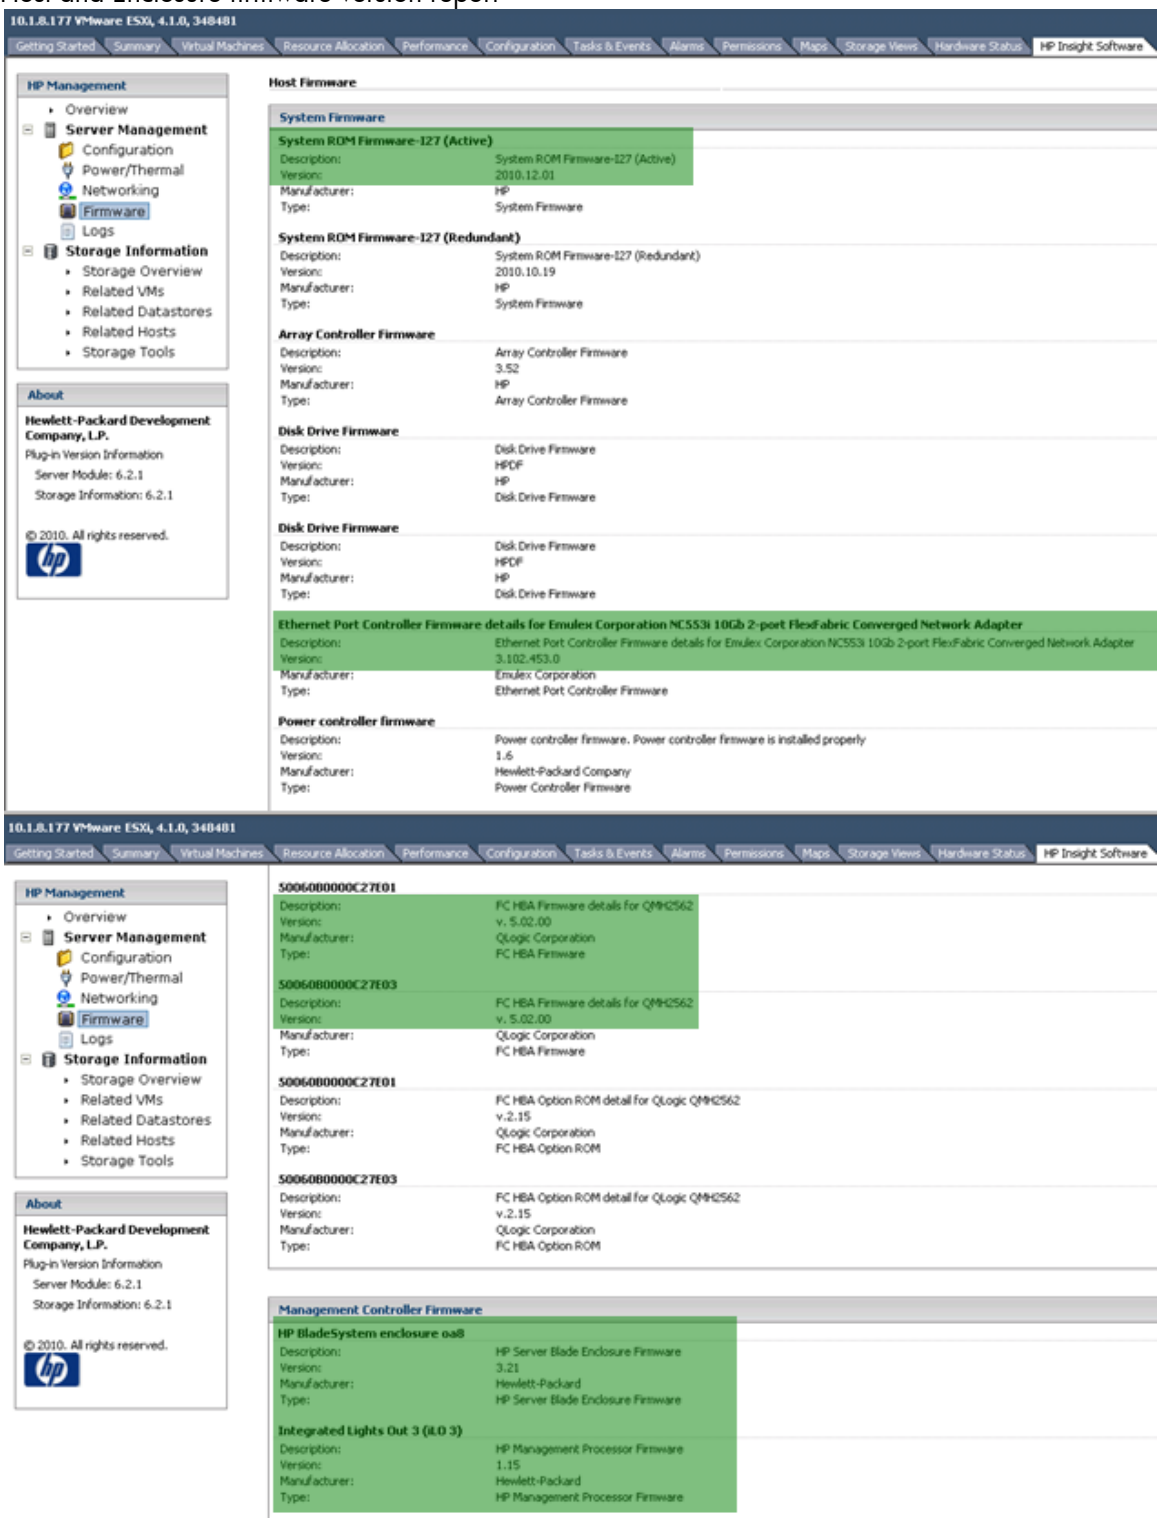

IMC network management

HP IMC is HP networking management software that supports network device configuration, accounting, performance, security management, and monitoring. It can manage HP Network devices, as well as routers and switches from other vendors.

The following images corresponding to this setup provide an overview of the appearance and functionality of IMC. It does not represent the full functionality of IMC.

For more information on IMC, see the HP website:
http://h17007.www1.hp.com/us/en/products/network-management/index.aspx

To download full-featured evaluation software, see the HP website:
https://h10145.www1.hp.com/downloads/SoftwareReleases.aspx?ProductNumber=JF377A&lang=en&cc=us&prodSeriesId=4176535

Overview page (can customize layout)

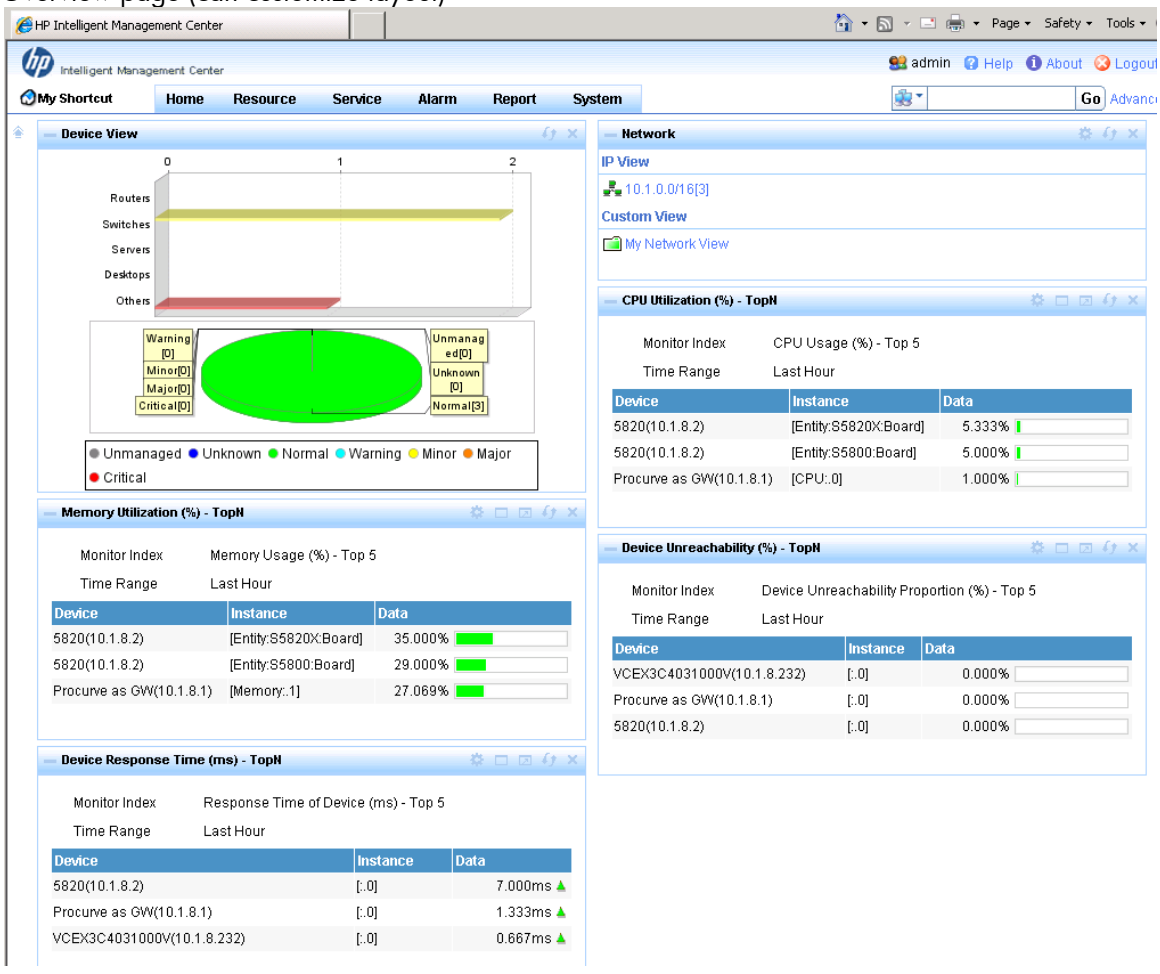

Network topology

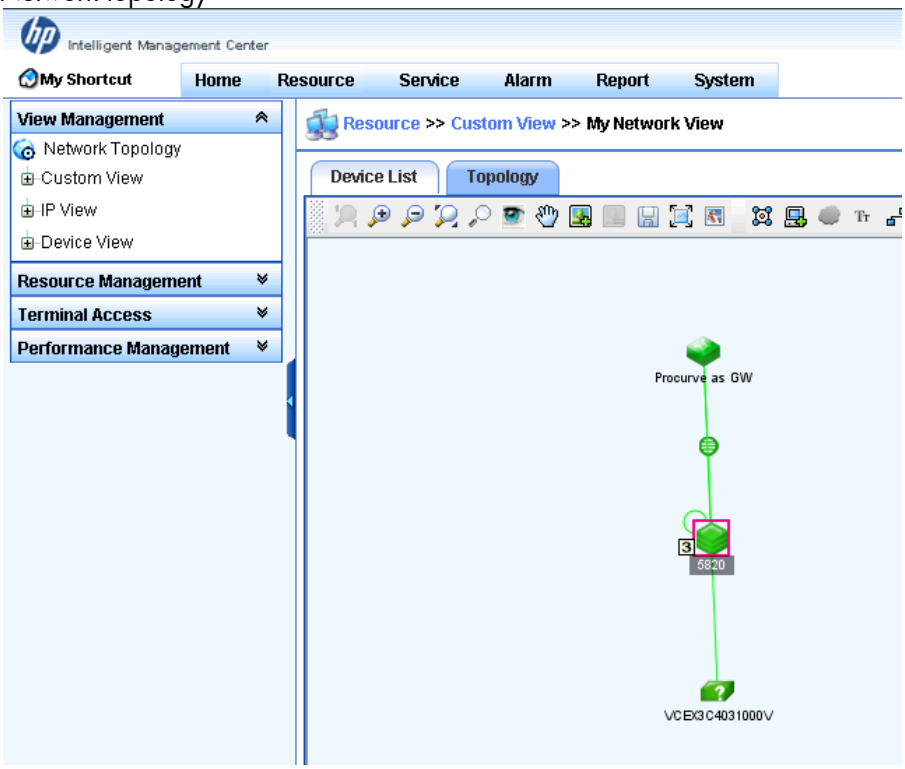

5820 IRF logical switch

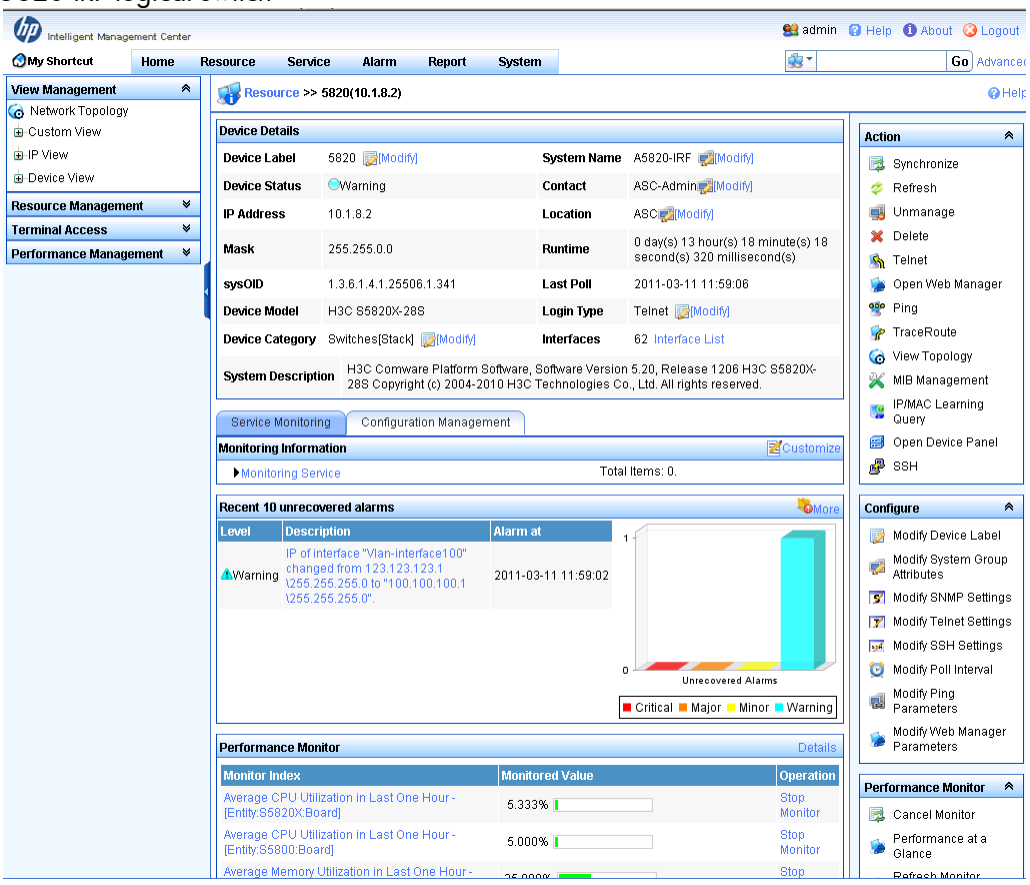

Virtual Connect interface list view

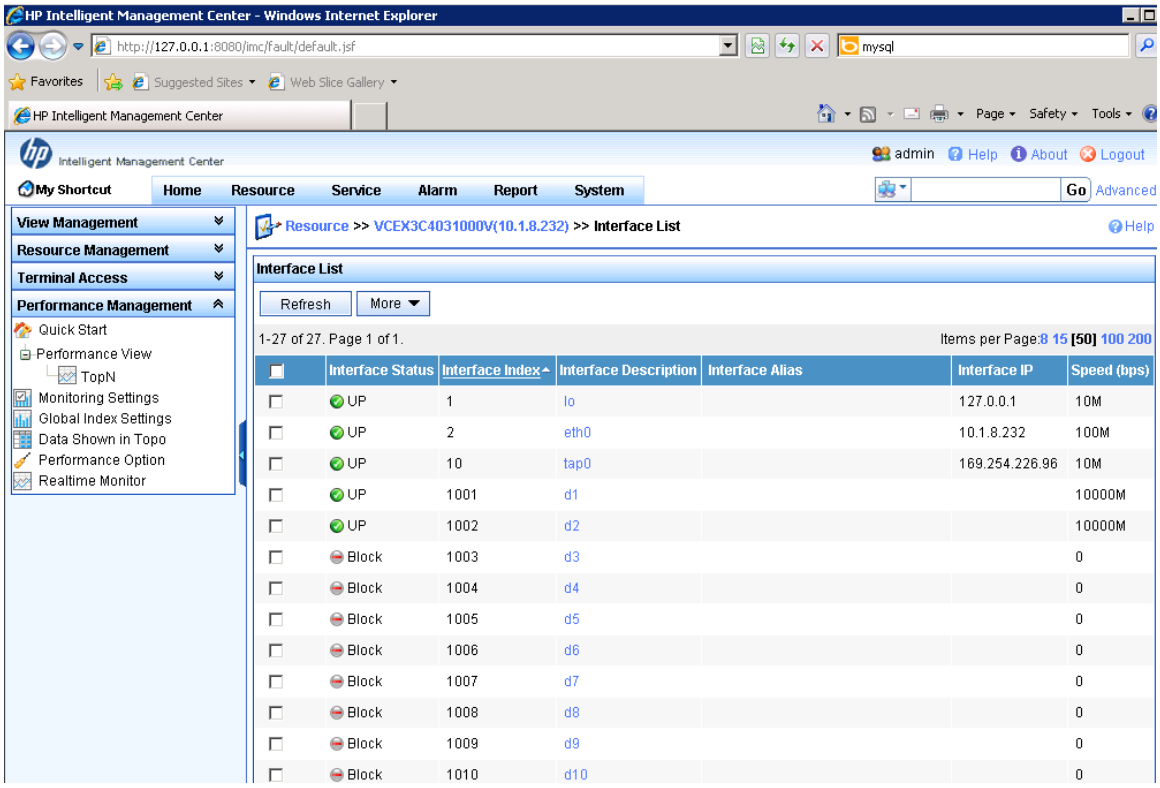

Interface traffic rate realtime monitoring

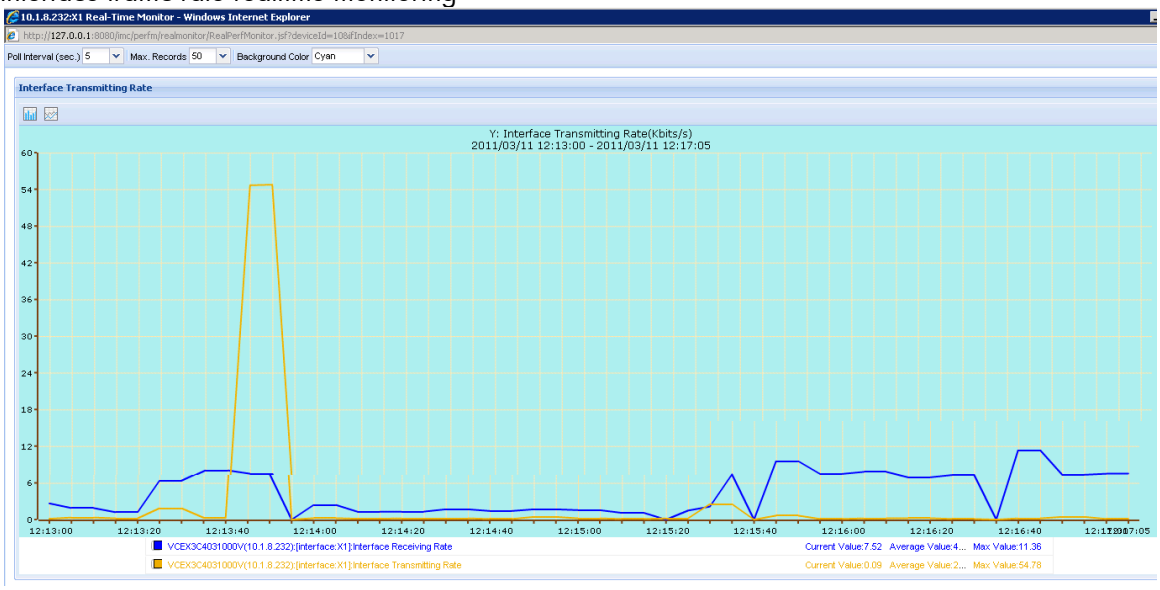

Appendix 1: A5820 logical switch IRF configuration

```
[A5820-IRF]dis current-configuration
#
 version 5.20, Release 1206
#
 sysname A5820-IRF
#
 irf mac-address persistent timer
 irf auto-update enable
 undo irf link-delay
 irf member 1 priority 10
#
 domain default enable system
#
 telnet server enable
#
vlan 1
#
vlan 2
#
vlan 100
#
radius scheme system
 server-type extended
 primary authentication 127.0.0.1 1645
 primary accounting 127.0.0.1 1646
 user-name-format without-domain
#
domain system
 access-limit disable
 state active
 idle-cut disable
 self-service-url disable
#
user-group system
#
 stp mode rstp
 stp enable
#
interface Bridge-Aggregation1
 port link-type trunk
 port trunk permit vlan 1 to 2
 link-aggregation mode dynamic
#
interface Bridge-Aggregation2
 port link-type trunk
 port trunk permit vlan 1 to 2
 link-aggregation mode dynamic
 stp edged-port enable
#
interface Bridge-Aggregation3
 port link-type trunk
 port trunk permit vlan 1 to 2
 link-aggregation mode dynamic
 stp edged-port enable
#
interface NULL0
#
interface Vlan-interface100
```

```
 mad bfd enable
 mad ip address 100.100.100.1 255.255.255.0 member 1
 mad ip address 100.100.100.2 255.255.255.0 member 2
#
interface GigabitEthernet1/0/25
 port link-mode bridge
 port access vlan 100
 stp disable
#
interface GigabitEthernet1/0/26
 port link-mode bridge
#
interface GigabitEthernet1/0/27
 port link-mode bridge
#
interface GigabitEthernet1/0/28
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 to 2
 port link-aggregation group 1
#
interface GigabitEthernet2/0/1
 port link-mode bridge
#
interface GigabitEthernet2/0/2
 port link-mode bridge
#
interface GigabitEthernet2/0/3
 port link-mode bridge
 port access vlan 100
 stp disable
#
interface GigabitEthernet2/0/4
 port link-mode bridge
#
interface GigabitEthernet2/0/5
 port link-mode bridge
#
interface GigabitEthernet2/0/6
 port link-mode bridge
#
interface GigabitEthernet2/0/7
 port link-mode bridge
#
interface GigabitEthernet2/0/8
 port link-mode bridge
#
interface GigabitEthernet2/0/9
 port link-mode bridge
#
interface GigabitEthernet2/0/10
 port link-mode bridge
#
interface GigabitEthernet2/0/11
 port link-mode bridge
#
interface GigabitEthernet2/0/12
 port link-mode bridge
#
interface GigabitEthernet2/0/13
 port link-mode bridge
```

```
#
interface GigabitEthernet2/0/14
 port link-mode bridge
#
interface GigabitEthernet2/0/15
 port link-mode bridge
#
interface GigabitEthernet2/0/16
 port link-mode bridge
#
interface GigabitEthernet2/0/17
 port link-mode bridge
#
interface GigabitEthernet2/0/18
 port link-mode bridge
#
interface GigabitEthernet2/0/19
 port link-mode bridge
#
interface GigabitEthernet2/0/20
 port link-mode bridge
#
interface GigabitEthernet2/0/21
 port link-mode bridge
#
interface GigabitEthernet2/0/22
 port link-mode bridge
#
interface GigabitEthernet2/0/23
 port link-mode bridge
#
interface GigabitEthernet2/0/24
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 to 2
 port link-aggregation group 1
#
interface M-GigabitEthernet0/0/0
 ip address 10.1.8.2 255.255.0.0
#
interface Ten-GigabitEthernet1/0/1
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 to 2
 port link-aggregation group 2
#
interface Ten-GigabitEthernet1/0/2
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 to 2
 port link-aggregation group 3
#
interface Ten-GigabitEthernet1/0/3
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/4
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/5
 port link-mode bridge
#
```

```
interface Ten-GigabitEthernet1/0/6
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/7
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/8
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/9
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/10
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/11
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/12
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/13
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/14
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/15
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/16
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/17
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/18
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/19
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/20
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/21
 port link-mode bridge
#
interface Ten-GigabitEthernet1/0/22
 port link-mode bridge
#
interface Ten-GigabitEthernet2/0/25
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 to 2
 port link-aggregation group 2
#
interface Ten-GigabitEthernet2/0/26
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 to 2
```

```
 port link-aggregation group 3
#
interface Ten-GigabitEthernet1/0/23
#
interface Ten-GigabitEthernet1/0/24
#
interface Ten-GigabitEthernet2/0/27
#
interface Ten-GigabitEthernet2/0/28
#
 ip route-static 0.0.0.0 0.0.0.0 10.1.0.1
#
 snmp-agent
 snmp-agent local-engineid 800063A203002389437528
 snmp-agent community read public
 snmp-agent sys-info contact ASC-Admin
 snmp-agent sys-info location ASC
 snmp-agent sys-info version all
 snmp-agent target-host trap address udp-domain 10.1.220.178 udp-port 161
params securityname public
 snmp-agent trap enable default-route
#
 load xml-configuration
#
user-interface aux 0 1
user-interface vty 0 15
 authentication-mode none
 user privilege level 3
#
irf-port 1/2
 port group interface Ten-GigabitEthernet1/0/23 mode enhanced
 port group interface Ten-GigabitEthernet1/0/24 mode enhanced
#
irf-port 2/1
 port group interface Ten-GigabitEthernet2/0/27 mode enhanced
 port group interface Ten-GigabitEthernet2/0/28 mode enhanced
#
return
[A5820-IRF]
```

Appendix 2: Design 3 running status

Using design option three, if one port channel interface is configured on an A5820 switch and one SUS is configured on a Virtual Connect to bundle four links on both sides, the A5820 switch does not select two out of four links as active LACP links.

```
[A5820-IRF-Ten-GigabitEthernet2/0/26]dis link-aggregation ver b2

Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Port Status: S -- Selected, U -- Unselected
Flags:  A -- LACP_Activity, B -- LACP_Timeout, C -- Aggregation,
        D -- Synchronization, E -- Collecting, F -- Distributing,
        G -- Defaulted, H -- Expired

Aggregation Interface: Bridge-Aggregation2
Aggregation Mode: Dynamic
Loadsharing Type: Shar
System ID: 0x8000, 0023-8943-7524
Local:
  Port            Status  Priority Oper-Key  Flag
--------------------------------------------------------------------------------
  XGE1/0/1        S       32768    1         {ACDEF}
  XGE1/0/2        U       32768    1         {AC}
  XGE2/0/25       S       32768    1         {ACDEF}
  XGE2/0/26       U       32768    1         {AC}
Remote:
  Actor           Partner Priority Oper-Key  SystemID               Flag
--------------------------------------------------------------------------------
  XGE1/0/1        17      1        3         0x1   , d485-64ce-f015 {ABCDEF}
  XGE1/0/2        17      1        3         0x1   , d485-64ce-f033 {ABCD}
  XGE2/0/25       18      1        3         0x1   , d485-64ce-f015 {ABCDEF}
  XGE2/0/26       18      1        3         0x1   , d485-64ce-f033 {ABCD}
```

VC1 status is OK.

Bay 1 (HP VC Flex-10 Enet Module)

Manufacturer:	HP
Firmware Rev.	3.15 2010-10-09T07:18:16Z

Uplink Port Information

Label	Network (s)	Status		Connector Type	LAG ID	Connected To	Detailed statistics
Port X1	Trunk-A	OK	Linked/Active 10 Gb	SFP-SR	25	00:23:89:43:75:24(Ten-GigabitEthernet1/0/1)	Detailed statistics/information
Port X2	Trunk-A	OK	Linked/Active 10 Gb	SFP-SR	25	00:23:89:43:75:24(Ten-GigabitEthernet2/0/25)	Detailed statistics/information

VC2 status is not OK. The "LAG ID" column is empty, which means no LACP bundle is established. Both links are put into standby as individual links for this SUS.

Bay 2 (HP VC Flex-10 Enet Module)

Port X1	Trunk-A	OK	Linked/Standby 10 Gb	SFP-SR		00:23:89:43:75:24(Ten-GigabitEthernet1/0/2)	Detailed statistics/information
Port X2	Trunk-A	OK	Linked/Standby 10 Gb	SFP-SR		00:23:89:43:75:24(Ten-GigabitEthernet2/0/26)	Detailed statistics/information

Acronyms

ARP – Address Resolution Protocol
BFD – Bidirectional Forwarding Detection
BPDU - Bridge Protocol Data Unit
GW – Gateway
IC – Insight Control
IMC – Intelligent Management Center
IRF – Intelligent Resilient Framework
LACP – Link Aggression Control Protocol
LACPDU – Link Aggression Control Protocol Data Unit
LLDP – Link-Layer Discovery Protocol
MAC – Media Access Control
MAD – Multi-Active Detection
OA – Onboard Administrator
SSH – Secure Shell
STP – Spanning Tree Protocol
SUS – Shared Uplink Set
VC – Virtual Connect
VLAN – Virtual Local Area Network
VM – Virtual Machine
vPC – Virtual Port Channel
VSS – Virtual Switching System

c02843088, Created May 2011, Updated July 2011