

HP 6125 Switch Series: Configuring IRF



Table of contents

Introduction.....	2
Background information	2
Fundamentals.....	2
IRF member roles	2
IRF member ID	2
IRF port	2
Physical IRF port	3
IRF domain ID.....	3
IRF split	3
IRF merge	3
Member priority	4
Topology.....	4
HP c-Class enclosures	5
Configuration highlights.....	9
Equipment used	10
Network diagram	11
Configuration	12
Accessing the HP 6125XLG Blade Switch	12
Configure IRF for the switch in bay 1	12
Configure IRF for the switch in bay 2	13
Configure VLANs and LACP	13
Verify.....	15
Configure and add additional IRF members	16
Configure switch bay 1—enclosure 1	17
Configure switch bay 1—enclosure 2	17
Configure switch bay 2—enclosure 2	18
Troubleshoot	19
Appendix: Sample configuration	19
Additional links	22

Introduction

This configuration guide describes how to configure Intelligent Resilient Framework (IRF) on the HP 6125 Blade Switch Series. The intended audiences are HP Solution Architects, HP Technical Consultants, HP partners, and HP customers.

Background information

IRF is a network virtualization technology that allows you to connect multiple devices through physical IRF ports to combine them into a single logical virtual device. This virtualization technology enables the cooperation, unified management, and non-stop maintenance of multiple devices.

IRF provides the following benefits:

- **Simplified network configuration management and increased operational efficiency:** Multiple linked physical devices look like one logical device and provide a single point of management. Only a single IP address and configuration file needs to be maintained.
- **Scalable performance:** IRF and Link Aggregation Control Protocol (LACP) used together can boost performance by bundling several parallel links between devices, allowing scalable “on-demand” performance and capacity to support critical business applications.
- **IP address configuration is simplified:** Each VLAN requires a single gateway IP address, eliminating the need to create identical configurations on all devices. Additionally, the resulting logical device is viewed as a single entity in the network management system, significantly simplifying network management.
- **Redundancy protocol and loop prevention:** Because multiple devices are virtualized into one logical device, loop prevention, reliability, and redundancy protocols—such as VRRP, STP, RSTP, and MSTP—can be eliminated. This simplifies network configuration and maintenance and eliminates design complexity, while enabling significantly decreased convergence times.
- **Guaranteed system reliability:** A device failure does not impact applications that rely on network state information. Additionally, LACP allows higher performance while eliminating single points of failure in the system. L2/L3 protocols do not need to re-converge when there’s a link failure within a LAG group.
- **Expanding bandwidth capacity:** A virtualized system provides an effective load balancing mechanism between member devices, thus fully utilizing available bandwidth.
- **Much more than just stacking:** IRF provides all the traditional functions of a stacking technology and more. For example, IRF is a multiprotocol (L2, L3 IPv4, L3 IPv6, MPLS, VPLS, Unicast, Multicast) technology, which will remove the need for technologies like VRRP, providing not only the virtual IP gateway but also allowing full active/active L3 forwarding.

Fundamentals

IRF member roles

The devices that form an IRF fabric are called IRF members. Each of them plays either of the following two roles:

- **Master:** Manages the entire IRF fabric.
- **Slave:** All member devices operating as the backups of the master are called slaves. When the master fails, the system automatically elects a new master from among the slaves.

IRF member ID

An IRF fabric uses member IDs to uniquely identify and manage its members. This member ID information is included as the first part of interface numbers and file paths to uniquely identify interfaces and files in an IRF fabric.

If two switches have the same IRF member ID, they cannot form an IRF fabric.

IRF port

An IRF port is a logical interface for the connection between IRF member devices. Every IRF-capable device supports two IRF ports. The IRF ports are named IRF-port n/1 and IRF-port n/2, where n is the member ID of the switch.

To use an IRF port, you must bind at least one physical port to it. The physical ports assigned to an IRF port automatically forms an aggregate IRF link. An IRF port goes down only if all its physical IRF ports are down.

For two neighboring devices, their IRF physical links must be bound to IRF-port 1 on one device and to IRF-port 2 on the other.


Physical IRF port

Physical IRF ports connect IRF member devices and must be bound to an IRF port. They forward IRF protocol packets between IRF member devices and data packets that must travel across IRF member devices.

IRF domain ID

One IRF fabric forms one IRF domain. IRF uses IRF domain IDs to uniquely identify IRF fabrics and prevent IRF fabrics from interfering with one another. As shown in figure 1, Switch A and Switch B form IRF fabric 1, and Switch C and Switch D form IRF fabric 2. The fabrics have LACP Multi-Active Detection (MAD) detection links between them. When a member switch in one of the IRF fabric receives an extended LACP packet for MAD detection, it looks at the domain ID in the packet to see whether the packet is from the local IRF fabric or from a different IRF fabric. Then, the switch can handle the packet correctly.


Figure 1. Two IRF domains



IRF split


An IRF split, also known as “split brain,” occurs when an IRF fabric breaks up into two or more IRF fabrics because of IRF link failures, as shown in figure 2. The split IRF fabrics operate with the same IP address and cause routing and forwarding problems on the network. To quickly detect a multiactive collision, configure at least one MAD mechanisms (see “[IRF configuration guide](#) for details on various MAD mechanisms”).

Figure 2. IRF split



IRF merge


IRF merge occurs when two split IRF fabrics re-unite or when you configure and connect two independent IRF fabrics to be one IRF fabric, as shown in figure 3.

Figure 3. IRF merge**Member priority**



Member priority determines the possibility of a member device to be elected the master. A member with higher priority is more likely to be elected the master. The default member priority is 1. You can change the member priority of a member device to affect the master election result.

Topology

The typical HP 6125 Blade Switch deployment will utilize at least two 6125 blade switches installed in adjacent bays in a c-Class enclosure. These two devices will form an IRF fabric using the internal IRF cross-connect port/s, and essentially would constitute a daisy chain topology.

Figure 4. IRF using a pair of blade switches in same enclosure

Additional blade switches can be added to the IRF fabric. The additional blade switches can be located in the same c-Class enclosure, or in separate enclosures. The primary requirement to be aware of is that blades in adjacent bays will use the internal cross-connect ports for the IRF connection while blades in non-adjacent bays will use front facing ports.

Figure 5. IRF using four blade switches in across multiple enclosures**Figure 6.** IRF using four blade switches in same enclosures


HP c-Class enclosures

HP BladeSystem c3000 and c7000 enclosures are able to consolidate and scale DC server deployments by providing an enclosure that can provide all the power, cooling, and I/O infrastructure needed to support multiple servers, interconnects, and storage components.

The c7000 enclosure is 10U high and holds up to 16 server and/or storage blades plus optional redundant network and storage interconnect modules.

The c3000 enclosure is 6U high and holds up to 8 server and/or storage blades plus optional redundant network and storage interconnect modules.


Figure 7. c7000 enclosure rear view




Although not a requirement, HP recommends installing interconnect blade switches in pairs for redundancy. Pairs of redundant interconnect blade switches should be installed in adjacent bays (i.e., bay 1 and bay 2, bay 3 and bay 4, etc.).

Understanding how the NICs and mezzanine cards on the servers connect to each interconnect bay is important for proper configurations.

As shown in figure 8, the built in LAN on Motherboard (LOM) NICs connect to interconnect bays 1 and 2. Specifically LOM NIC1 connects to interconnect bay 1 and LOM NIC2 connects to interconnect bay 2. The diagram also shows how additional added mezzanine cards connect to the interconnect bays.

Figure 8. c7000 half-height server/interconnect port mapping

Each HP 6125 Blade Switch supports 16 downlink ports (shown in CLI as ports 1/0/1 to 1/0/16 [for device with an IRF member 1 ID]). Understanding these concepts helps administrators determine which switch ports are connected to which server NIC. Figure 9 shows another example diagram showing Blade Server 1 LOM connectivity to Interconnect bays 1 and 2.

Figure 9. c7000 interconnect bay 1 and 2 to server 1 mapping

Additional diagrams here detail the port mapping for full height servers as well as for c3000 deployments.

Figure 10. c7000 full-height server/interconnect port mapping




Figure 11. c3000 half-height server/interconnect port mapping





Figure 12. c3000 full-height server/interconnect port mapping



Configuration highlights

When deploying an HP 6125 IRF Solution consider the following design guidelines:

- All IRF member switches must run the same system software image version.
- The HP 6125G Blade Switch supports up to two physical ports per IRF port.
- The HP 6125G/XG and 6125XLG Blade Switch supports up to four physical ports per IRF port.
- The HP 6125G and 6125G/XG can create an IRF fabric together.
- The HP 6125XLG Blade Switch can form an IRF fabric with other HP 6125XLG Switches.
- The HP 6125XLG is supported in the c7000 enclosure only.
- To connect blade switches in different chassis, or in non-adjacent bays, into an IRF fabric, you must use the SFP+ and QSFP+ ports on the switch front panel.
- To connect blade switches in adjacent bays in the same chassis into an IRF fabric, you must use the internal cross-connect 10GbE ports. These ports are invisible to users and do not require physical cabling.
- On the HP 6125G Blade Switch, candidate IRF physical ports are the two front panel IRF/SFP ports, and one internal 10GbE cross-connect port.
- On the HP 6125G/XG Blade Switch, candidate IRF physical ports are the four front panel SFP+ ports, and one internal 10GbE cross-connect port.
- On the HP 6125XLG Blade Switch, candidate IRF physical ports are the eight front panel SFP+ ports and four QSFP+ ports, and four internal 10GbE cross-connect ports.
 - When configuring IRF using a QSFP+ 40GbE to 4x10GbE splitter cable, you must use all or none of the four 10GbE interfaces as IRF physical ports. The four interfaces can be bound to different IRF ports.
 - Before you bind a 10GbE interface to an IRF port or remove it from the IRF port, you must shut down all the 10GbE interfaces of the 40GbE port. If any of the interface is in the “up” state, the bind or remove action will fail.

- The front panel SFP+ ports are grouped by port index into two 4-port groups. One port group contains ports 5, 6, 9, and 10. The other port group contains ports 7, 8, 11, and 12. If you use one port in a group for an IRF connection, you must also use all the other ports in the group for an IRF connection. However, you can bind them to different IRF ports.
- Before you bind an SFP+ port to an IRF port or remove it from the IRF port, you must shut down all the SFP+ ports in the same group.
- When you use the internal 10GbE cross-connect ports on the switch rear panel as IRF physical ports, follow these guidelines:
 - If you use one of the downlink 10GbE ports for an IRF connection, you must also use all the other downlink 10GbE ports for an IRF connection. However, you can bind them to different IRF ports.
 - Before you bind a downlink 10GbE port to an IRF port or remove it from the IRF port, you must shut down all the other downlink 10GbE ports on the switch rear panel.
- When you connect two neighboring IRF members, connect IRF-port 1 on one member to IRF-port 2 on the other.
 - Suppose you have four chassis: A, B, C, and D. IRF-port 1 and IRF-port 2 are represented by A1 and A2 on chassis A, represented by B1 and B2 on chassis B, and so on. To connect the four chassis into a ring topology of A-B-C-D(A), the IRF link cabling scheme must be one of the following:
 - A. A1-B2, B1-C2, C1-D2, and D1-A2.
 - B. A2-B1, B2-C1, C2-D1, and D2-A1.
- Configure at least one MAD mechanism for prompt IRF split detection and IRF fabric recovery.
- If LACP MAD runs between two IRF fabrics, assign each fabric a unique IRF domain ID.
- By default, the member IDs of all switches are 1. To create an IRF fabric, you must assign a unique IRF member ID to each switch. You must reboot the members to validate the IRF member ID settings.
- Assign the highest member priority to the device you want to use as the master.
- Save any configuration you have made to the startup configuration file before rebooting the IRF member devices.
- No intermediate devices are allowed between neighboring IRF members.

Table 1. HP 6125 IRF specifications

	HP 6125G and 6125G/XG Blade Switch	HP 6125XLG Blade Switch
Stack interface	10GbE	10GbE/40GbE
Stack bandwidth	6125G = 2x10GbE 6125G/XG = 4x10GbE	4x40GbE/4x10GbE
Max devices in IRF fabric	10	8
Internal 10GbE cross-connect IRF ports	1	4
Stack with different model	Yes 6125G and 6125G/XG can form IRF	No
IRF split detection (i.e., MAD)	Yes	Yes

Equipment used

The following hardware is used in the configuration example below:

- HP 6125XLG Blade Switch (2)
- HP 12508 switches. The pair of 12508 switches represent a core IRF device (2)
- HP c7000 Enclosure G2 (1)
 - HP ProLiant BL460c G7 Servers (4)
- Optic cables to make the connections (as needed)


The following software was used:

- HP 6125XLG—HP Comware Software, version 7.1.045, release 2403
- HP 12508—HP Comware Software, version 7.1.034, release 7129
- c-Class enclosure firmware—2013.02.0

Network diagram

The network diagram below illustrates the connectivity for this configuration.

Figure 13. Network design



The above configuration is an IRF configuration example using a pair of HP 6125XLG in a c7000 enclosure.

The HP 6125XLG Switches utilize four internal cross-connects for the IRF link, two redundant 10GbE downlinks to Blade Server 1 through 4, and four redundant 10GbE uplinks to the core HP 12500 Switches using LACP and LACP MAD detection. The core pair of HP 12500 Switches are enabled with IRF technology, however, configuration of the core is out of scope of this document.

For complete, detailed IRF configuration options review the [HP 6125 IRF configuration guide](#).

In this example:

- Two HP 6125XLG Blade Switches form a single IRF fabric using four internal cross-connect links
- Two redundant 10GbE downlinks are used to connect to blade servers 1 to 4
- Four redundant 10GbE uplinks are used to connect to the core HP 12500 Switches (configuration of core pair of HP 12508 Switches is out of scope of this document)
- To prevent a split brain scenario, the HP 6125XLG Blade Switches utilize LACP for the MAD mechanism (LACP MAD requires an HP Comware-based ToR switch. Refer the [HP 6125 IRF configuration guide](#) for more MAD options)

Note:

1. Configuring NIC teaming on the server side will vary based on the platform/OS used. Review documentation for chosen server solution to determine method needed.
2. The solution shown below has server blade 1 and 2 connecting to eSXI host running vSphere 5.1 Standard vSwitches. This means the link-aggregation configuration on the switch side cannot use dynamic mode, and the vSwitch on the eSXI host needs to use IP hash load balancing.
3. Server blades 3 and 4 are connected to Windows® 2008 Host with teaming configured, so dynamic mode is supported.

Configuration

Accessing the HP 6125XLG Blade Switch

When the HP 6125XLG first starts up there are two serial connection methods that can be used to access the CL—no username or password required. (After the first login, HP strongly recommends that users configure password or scheme authentication mode to improve the devices security):

- **con0:** External front panel console port
- **aux0:** Internal console port, which can be accessed through the HP Blade Enclosure Onboard Administrator (OA)

After logging into the switch, administrators can change the console or AUX login parameters or configure other access methods, including Telnet, SSH, and SNMP.

To access the device using the aux port, users can first connect to the OA using Telnet. Once in the OAs CLI, the 6125 aux port can be accessed by using the following command (replace X with the bay number of the switch):

```
connect interconnect X
```

The serial connection number (con0/aux0) comes from the IRF member ID (ID-1). A switch with an IRF member ID of 2 would see the numbers change to con1 and aux1.

Note:

The default aux0 port will allow full network-admin access for configuration, while default config for aux1 does not permit full admin management. If an admin re-numbers the switch, for example, from 1 to 2, the aux port will now change to aux1 and the admin will lose full network-admin access. This issue can be resolved in the initial configuration by changing the user-role of the aux line class using the following commands:

```
line class aux
 authentication-mode none
 user-role network-admin
```

Configure IRF for the switch in bay 1

1. Change member 1 priority to 32 to ensure device is master (optional)

```
#switch bay 1
system-view
 irf member 1 priority 32
```

2. Shut down the physical ports used for IRF

```
system-view
 interface range Ten-GigabitEthernet 1/0/17 to Ten-GigabitEthernet
 1/0/20
 shutdown
 quit
```

3. Create IRF port 1/1, bind the physical ports to it, and undo shutdown on all physical IRF ports

```
irf-port 1/1
 port group interface Ten-GigabitEthernet1/0/17
 port group interface Ten-GigabitEthernet1/0/18
 port group interface Ten-GigabitEthernet1/0/19
 port group interface Ten-GigabitEthernet1/0/20
 quit
 interface range Ten-GigabitEthernet 1/0/17 to Ten-GigabitEthernet
 1/0/20
 undo shutdown
 quit
 save
```

4. Activate the IRF port

```
irf-port-configuration active
```

Configure IRF for the switch in bay 2

1. Change the member ID to 2, save and reboot

```
#switch bay 2
system-view
  irf member 1 renumber 2
  save
quit
reboot
```

2. Shut down the physical ports used for IRF

```
system-view
  interface range Ten-GigabitEthernet 2/0/17 to Ten-GigabitEthernet
  2/0/20
  shutdown
quit
```

3. Create IRF port 2/2, bind the physical ports to it, and undo shutdown on all physical IRF ports

```
irf-port 2/2
  port group interface Ten-GigabitEthernet2/0/17
  port group interface Ten-GigabitEthernet2/0/18
  port group interface Ten-GigabitEthernet2/0/19
  port group interface Ten-GigabitEthernet2/0/20
quit
interface range Ten-GigabitEthernet 2/0/17 to Ten-GigabitEthernet
2/0/20
  undo shutdown
quit
save
```

4. Activate the IRF port. The two devices perform master election, and the one that has lost the election reboots to form an IRF fabric with the master

```
irf-port-configuration active
```

Configure VLANs and LACP

1. Configure VLANs as needed

```
#switch_IRF
system-view
  vlan 1
  quit
  vlan 1010
  quit
  interface vlan 1
  ip address 10.1.1.10 24
  quit
  interface vlan 1010
  ip address 10.10.10.74 24
  quit
```

2. Configure downlink interfaces connecting to servers (configuration leaves downlinks to servers as access ports for VLAN 1)

```
interface Bridge-Aggregation 1
  description to Server 1
  quit
interface Bridge-Aggregation 2
  description to Server 2
  quit
interface Bridge-Aggregation 3
  description to Server 3
  link-aggregation mode dynamic
  quit
interface Bridge-Aggregation 4
```

```

description to Server 4
link-aggregation mode dynamic
quit
interface Ten-GigabitEthernet 1/0/1
description to Server 1
port link-aggregation group 1
quit
interface Ten-GigabitEthernet 2/0/1
description to Server 1
port link-aggregation group 1
quit
interface Ten-GigabitEthernet 1/0/2
description to Server 2
port link-aggregation group 2
quit
interface Ten-GigabitEthernet 2/0/2
description to Server 2
port link-aggregation group 2
quit
interface Ten-GigabitEthernet 1/0/3
description to Server 3
port link-aggregation group 3
quit
interface Ten-GigabitEthernet 2/0/3
description to Server 3
port link-aggregation group 3
quit
interface Ten-GigabitEthernet 1/0/4
description to Server 4
port link-aggregation group 4
quit
interface Ten-GigabitEthernet 2/0/4
description to Server 4
port link-aggregation group 4
quit

```

3. Configure uplink interfaces connecting to core, configure domain ID, and enable MAD

```

interface Bridge-Aggregation 20
description to Core 12508
link-aggregation mode dynamic
quit
interface Ten-GigabitEthernet 1/1/5
description to Core 12508
port link-aggregation group 20
quit
interface Ten-GigabitEthernet 1/1/6
description to Core 12508
port link-aggregation group 20
quit
interface Ten-GigabitEthernet 2/1/5
description to Core 12508
port link-aggregation group 20
quit
interface Ten-GigabitEthernet 2/1/6
description to Core 12508
port link-aggregation group 20
quit
irf domain 2
interface Bridge-Aggregation 20
port link-type trunk
port trunk permit vlan 1 1010
mad enable
You need to assign a domain ID (range: 0-4294967295)
[Current domain is: 2]: 2
The assigned domain ID is: 2
MAD LACP only enable on dynamic aggregation interface.

```

Verify

1. Verify IRF

```
[Switch Bay 1]display irf
```

MemberID	Role	Priority	CPU-Mac	Description
*+1	Master	32	4431-922c-3649	---
2	Standby	1	4431-926c-7fd4	---

```
-----
* indicates the device is the master.
```

```
+ indicates the device through which the user logs in.
```

```
The Bridge MAC of the IRF is: 4431-922c-3648
```

```
Auto upgrade           : yes
```

```
Mac persistent         : 6 min
```

```
Domain ID              : 2
```

```
[Switch Bay 1]display irf configuration
```

MemberID	NewID	IRF-Port1	IRF-Port2
1	1	Ten-GigabitEthernet1/0/17	disable
		Ten-GigabitEthernet1/0/18	
		Ten-GigabitEthernet1/0/19	
		Ten-GigabitEthernet1/0/20	
2	2	disable	Ten-GigabitEthernet2/0/17
			Ten-GigabitEthernet2/0/18
			Ten-GigabitEthernet2/0/19
			Ten-GigabitEthernet2/0/20

```
[Switch Bay 1]display irf topology
```

```
Topology Info
```

MemberID	IRF-Port1		IRF-Port2		Belong To
	Link	neighbor	Link	neighbor	
2	DIS	---	UP	1	4431-922c-3649
1	UP	2	DIS	---	4431-922c-3649

```
[Switch Bay 1]display irf link
```

```
Member 1
```

IRF Port	Interface	Status
1	Ten-GigabitEthernet1/0/17	UP
	Ten-GigabitEthernet1/0/18	UP
	Ten-GigabitEthernet1/0/19	UP
	Ten-GigabitEthernet1/0/20	UP
2	disable	--

```
Member 2
```

IRF Port	Interface	Status
1	disable	--
2	Ten-GigabitEthernet2/0/17	UP
	Ten-GigabitEthernet2/0/18	UP
	Ten-GigabitEthernet2/0/19	UP
	Ten-GigabitEthernet2/0/20	UP

```
[Switch Bay 1]display interface brief | include UP
```

InLoop0	UP	UP(s)	--	
M-E0/0/0	UP	UP	--	
NULL0	UP	UP(s)	--	
Vlan1	UP	UP	10.1.1.10	
Vlan1010	UP	UP	10.10.10.74	OOB Network
BAGG1	UP	20G(a)	F(a) A 1	to Server 1
BAGG2	UP	20G(a)	F(a) A 1	to Server 2
BAGG3	UP	20G(a)	F(a) A 1	to Server 3
BAGG4	UP	20G(a)	F(a) A 1	to Server 4
BAGG20	UP	40G(a)	F(a) T 1	to 12500 Core


XGE1/0/1	UP	10G(a)	F(a)	A	1	to Server 1
XGE1/0/2	UP	10G(a)	F(a)	A	1	to Server 2
XGE1/0/3	UP	10G(a)	F(a)	A	1	to Server 3
XGE1/0/4	UP	10G(a)	F(a)	A	1	to Server 4
XGE1/0/17	UP	10G(a)	F(a)	--	--	
XGE1/0/18	UP	10G(a)	F(a)	--	--	
XGE1/0/19	UP	10G(a)	F(a)	--	--	
XGE1/0/20	UP	10G(a)	F(a)	--	--	
XGE1/1/5	UP	10G(a)	F(a)	T	1	to 12500 Core
XGE1/1/6	UP	10G(a)	F(a)	T	1	to 12500 Core
XGE2/0/1	UP	10G(a)	F(a)	A	1	to Server 1
XGE2/0/2	UP	10G(a)	F(a)	A	1	to Server 2
XGE2/0/3	UP	10G(a)	F(a)	A	1	to Server 3
XGE2/0/4	UP	10G(a)	F(a)	A	1	to Server 4
XGE2/0/17	UP	10G(a)	F(a)	--	--	
XGE2/0/18	UP	10G(a)	F(a)	--	--	
XGE2/0/19	UP	10G(a)	F(a)	--	--	
XGE2/0/20	UP	10G(a)	F(a)	--	--	
XGE2/1/5	UP	10G(a)	F(a)	T	1	to 12500 Core
XGE2/1/6	UP	10G(a)	F(a)	T	1	to 12500 Core

```
[Switch Bay 1]display mad
MAD ARP disabled.
MAD ND disabled.
MAD LACP enabled.
MAD BFD disabled.
```

Configure and add additional IRF members

Below is an example of adding additional IRF members. In this example, a pair of 6125XLG in a second enclosure are added to the IRF fabric.

Figure 14. Example adding additional IRF members



Configure switch bay 1—enclosure 1

1. Shut down the physical ports used for new IRF port

```
#switch bay 1&2 (these switches are now in IRF)
system-view
interface range Ten-GigabitEthernet 1/1/7 Ten-GigabitEthernet
1/1/8 Ten-GigabitEthernet 1/1/11 Ten-GigabitEthernet 1/1/12 Ten- GigabitEthernet
2/1/7 Ten-GigabitEthernet 2/1/8 Ten-GigabitEthernet 2/1/11 Ten-GigabitEthernet
2/1/12
shutdown
quit
```

2. Create IRF port 1/2 and 2/1, bind the physical ports to them, undo shutdown on all physical IRF ports, and save

```
irf-port 1/2
port group interface Ten-GigabitEthernet1/1/7
port group interface Ten-GigabitEthernet1/1/8
port group interface Ten-GigabitEthernet1/1/11
port group interface Ten-GigabitEthernet1/1/12
quit
irf-port 2/1
port group interface Ten-GigabitEthernet2/1/7
port group interface Ten-GigabitEthernet2/1/8
port group interface Ten-GigabitEthernet2/1/11
port group interface Ten-GigabitEthernet2/1/12
quit
interface range Ten-GigabitEthernet 1/1/7 Ten-GigabitEthernet
1/1/8 Ten-GigabitEthernet 1/1/11 Ten-GigabitEthernet 1/1/12 Ten-
GigabitEthernet 2/1/7 Ten-GigabitEthernet 2/1/8 Ten-GigabitEthernet 2/1/11 Ten-
GigabitEthernet 2/1/12
undo shutdown
quit
save
```

3. Activate the IRF port

```
irf-port-configuration active
```

Configure switch bay 1—enclosure 2

1. Connect all IRF cables between enclosures
2. Change the member ID to 3, save, and reboot

```
#switch bay 1-enclosure 2
system-view
irf member 1 renumber 4
save
quit
reboot
```

3. Shut down the physical ports used for IRF ports

```
system-view
interface range Ten-GigabitEthernet 4/0/17 to Ten-GigabitEthernet
4/0/20 Ten-GigabitEthernet 4/1/7 Ten-GigabitEthernet
4/1/8 Ten-GigabitEthernet 4/1/11 Ten-GigabitEthernet 4/1/12
shutdown
quit
```

4. Create IRF port 4/1 and 4/2, bind the physical ports to them, undo shutdown on all physical IRF ports, and save

```
irf-port 4/2
port group interface Ten-GigabitEthernet4/0/17
port group interface Ten-GigabitEthernet4/0/18
port group interface Ten-GigabitEthernet4/0/19
port group interface Ten-GigabitEthernet4/0/20
quit
```

```

irf-port 4/1
port group interface Ten-GigabitEthernet4/1/7
port group interface Ten-GigabitEthernet4/1/8
port group interface Ten-GigabitEthernet4/1/11
port group interface Ten-GigabitEthernet4/1/12
quit
interface range Ten-GigabitEthernet 4/0/17 to Ten-GigabitEthernet
4/0/20 Ten-GigabitEthernet 4/1/7 Ten-GigabitEthernet
4/1/8 Ten-GigabitEthernet 4/1/11 Ten-GigabitEthernet 4/1/12
undo shutdown
quit
save

```

5. Activate the IRF port

```
irf-port-configuration active
```

Configure switch bay 2—enclosure 2

1. Change the member ID to 3, save, and reboot

```

#switch bay 2—enclosure 3
system-view
irf member 1 renumber 3
save
quit
reboot

```

2. Shut down the physical ports used for IRF ports

```

system-view
interface range Ten-GigabitEthernet 3/0/17 to Ten-GigabitEthernet
3/0/20 Ten-GigabitEthernet 3/1/7 Ten-GigabitEthernet
3/1/8 Ten-GigabitEthernet 3/1/11 Ten-GigabitEthernet 3/1/12
shutdown
quit

```

3. Create IRF port 3/1 and 3/2, bind the physical ports to them, undo shutdown on all physical IRF ports, and save

```

irf-port 3/1
port group interface Ten-GigabitEthernet3/0/17
port group interface Ten-GigabitEthernet3/0/18
port group interface Ten-GigabitEthernet3/0/19
port group interface Ten-GigabitEthernet3/0/20
quit
irf-port 3/2
port group interface Ten-GigabitEthernet3/1/7
port group interface Ten-GigabitEthernet3/1/8
port group interface Ten-GigabitEthernet3/1/11
port group interface Ten-GigabitEthernet3/1/12
quit
interface range Ten-GigabitEthernet 3/0/17 to Ten-GigabitEthernet
3/0/20 Ten-GigabitEthernet 3/1/7 Ten-GigabitEthernet
3/1/8 Ten-GigabitEthernet 3/1/11 Ten-GigabitEthernet 3/1/12
undo shutdown
quit
save

```

4. Activate the IRF port

```
irf-port-configuration active
```

The devices perform master election, and the ones that have lost the election will reboot to form an IRF fabric with the master.

Master election uses the following rules in descending order:

- Current master stays the master, even if a new member has a higher priority

- When an IRF fabric is being formed, all members consider themselves as the master, and this rule is skipped
 - Member with higher priority
 - Member with the longest system uptime
 - Two members are considered starting up at the same time if the difference between their startup times is equal to or less than 10 minutes. For these members, the next tiebreaker applies
 - Member with the lowest bridge MAC address
5. Once IRF fabric is built, administrators will now need to modify uplinks to 12500 so each 6125XLG has a link in BAGG20, and the switches in enclosure 2 will need to configure downlink ports to servers as needed.

Troubleshoot

If the configuration cannot be verified, follow these steps to troubleshoot:

- Make sure that the physical interfaces that were added to the logical IRF ports are the interfaces to which the IRF cables are connected
- Make sure that logical IRF port n/1 connects to logical IRF port n/2
- Ensure that the physical IRF interfaces are up and still not shut down
- Reboot the devices one at a time

For more information, review the [HP 6125 IRF configuration guide](#).

Appendix: Sample configuration

```
#
irf domain 2
irf mac-address persistent timer
irf auto-update enable
undo irf link-delay
irf member 1 priority 32
irf member 2 priority 1
#
lldp global enable
#
vlan 1
#
vlan 1010
description OOB Network
#
irf-port 1/1
port group interface Ten-GigabitEthernet1/0/17
port group interface Ten-GigabitEthernet1/0/18
port group interface Ten-GigabitEthernet1/0/19
port group interface Ten-GigabitEthernet1/0/20
#
irf-port 2/2
port group interface Ten-GigabitEthernet2/0/17
port group interface Ten-GigabitEthernet2/0/18
port group interface Ten-GigabitEthernet2/0/19
port group interface Ten-GigabitEthernet2/0/20
#
stp global enable
#
interface Bridge-Aggregation1
description to Server 1
#
interface Bridge-Aggregation2
description to Server 2
```

```

#
interface Bridge-Aggregation3
  description to Server 3
  link-aggregation mode dynamic
#
interface Bridge-Aggregation4
  description to Server 4
  link-aggregation mode dynamic
#
interface Bridge-Aggregation20
  description to 12500 Core
  port link-type trunk
  port trunk permit vlan 1 1010
  link-aggregation mode dynamic
  mad enable
#
interface NULL0
#
interface Vlan-interface1
  ip address 10.1.1.10 255.255.255.0
#
interface Vlan-interface1010
  description OOB Network
  ip address 10.10.10.74 255.255.255.0
#
interface Ten-GigabitEthernet1/0/1
  port link-mode bridge
  description to Server 1
  port link-aggregation group 1
#
interface Ten-GigabitEthernet1/0/2
  port link-mode bridge
  description to Server 2
  port link-aggregation group 2
#
interface Ten-GigabitEthernet1/0/3
  port link-mode bridge
  description to Server 3
  port link-aggregation group 3
#
interface Ten-GigabitEthernet1/0/4
  port link-mode bridge
  description to Server 4
  port link-aggregation group 4
#
interface Ten-GigabitEthernet1/1/5
  port link-mode bridge
  description to 12500 Core
  port link-type trunk
  port trunk permit vlan 1 1010
  port link-aggregation group 20
#
interface Ten-GigabitEthernet1/1/6
  port link-mode bridge
  description to 12500 Core
  port link-type trunk
  port trunk permit vlan 1 1010
  port link-aggregation group 20
#
interface Ten-GigabitEthernet2/0/1
  port link-mode bridge
  description to Server 1
  port link-aggregation group 1
#
interface Ten-GigabitEthernet2/0/2
  port link-mode bridge
  description to Server 2

```

```
port link-aggregation group 2
#
interface Ten-GigabitEthernet2/0/3
port link-mode bridge
description to Server 3
port link-aggregation group 3
#
interface Ten-GigabitEthernet2/0/4
port link-mode bridge
description to Server 4
port link-aggregation group 4
#
interface Ten-GigabitEthernet2/1/5
port link-mode bridge
description to 12500 Core
port link-type trunk
port trunk permit vlan 1 1010
port link-aggregation group 20
#
interface Ten-GigabitEthernet2/1/6
port link-mode bridge
description to 12500 Core
port link-type trunk
port trunk permit vlan 1 1010
port link-aggregation group 20
#
```

Additional links

HP Networking
hp.com/networking

HP Servers
hp.com/servers

HP Blade Servers and Interconnects
hp.com/go/bladesystem

Learn more
hp.com/networking/irf

Sign up for updates
hp.com/go/getupdated

   
Share with colleagues


Rate this document

© Copyright 2014 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Windows is a U.S. registered trademark of the Microsoft group of companies.

4AA5-5625ENW, November 2014

