

HP 5500 EI & 5500 SI Switch Series

IRF

Configuration Guide

Part number: 5998-1714

Software version: Release 2220

Document version: 6W100-20130810



Legal and notice information

© Copyright 2013 Hewlett-Packard Development Company, L.P.

No part of this documentation may be reproduced or transmitted in any form or by any means without prior written consent of Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice.

HEWLETT-PACKARD COMPANY MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

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.

Contents

IRF overview	1
Hardware compatibility	1
IRF benefits	1
Application scenario	1
Basic concepts	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	4
Member priority	4
Interface naming conventions	4
File system naming conventions	5
Configuration synchronization mechanism	6
Master election	6
IRF multi-active detection	6
Multi-active handling procedure	6
LACP MAD	7
BFD MAD	8
ARP MAD	9
Configuring IRF	11
General restrictions and configuration guidelines	11
Software requirements	11
IRF link redundancy	11
IRF physical port restrictions and cabling requirements	11
IRF port binding restrictions	11
MAD	11
FIPS mode requirement	12
Other configuration guidelines	12
Setup and configuration task list	12
Planning the IRF fabric setup	14
Assigning a member ID to each IRF member switch	14
Specifying a priority for each member switch	15
Connecting physical IRF ports	15
Binding physical ports to IRF ports	16
Accessing the IRF fabric	18
Accessing the CLI of the master switch	18
Accessing the CLI of a subordinate switch	19
Assigning an IRF domain ID to the IRF fabric	19
Configuring a member switch description	19
Configuring IRF link load sharing mode	20
Configuring the global load sharing mode	20
Configuring a port-specific load sharing mode	20
Configuring IRF bridge MAC persistence	21
Enabling software auto-update for system software image synchronization	21
Setting the IRF link down report delay	22

Configuring MAD.....	23
Configuring LACP MAD.....	24
Configuring BFD MAD.....	25
Configuring ARP MAD.....	26
Excluding a port from the shutdown action upon detection of multi-active collision.....	27
Recovering an IRF fabric.....	28
Displaying and maintaining an IRF fabric.....	29
Configuration examples	30
LACP MAD-enabled IRF configuration example.....	30
BFD MAD-enabled IRF configuration example.....	32
ARP MAD-enabled IRF configuration example.....	34
Support and other resources	37
Contacting HP	37
Subscription service	37
Related information	37
Documents.....	37
Websites.....	37
Conventions	38
Index	40

IRF overview

The HP Intelligent Resilient Framework (IRF) technology creates a large IRF fabric from multiple switches to provide data center class availability and scalability. IRF virtualization technology offers processing power, interaction, unified management, and uninterrupted maintenance of multiple switches.

This book describes IRF concepts and guides you through the IRF setup procedure.

Hardware compatibility

HP 5500 switches in an IRF fabric must be the same sub-series. For example, an HP 5500 EI and an HP 5500 SI cannot form an IRF fabric.

IRF benefits

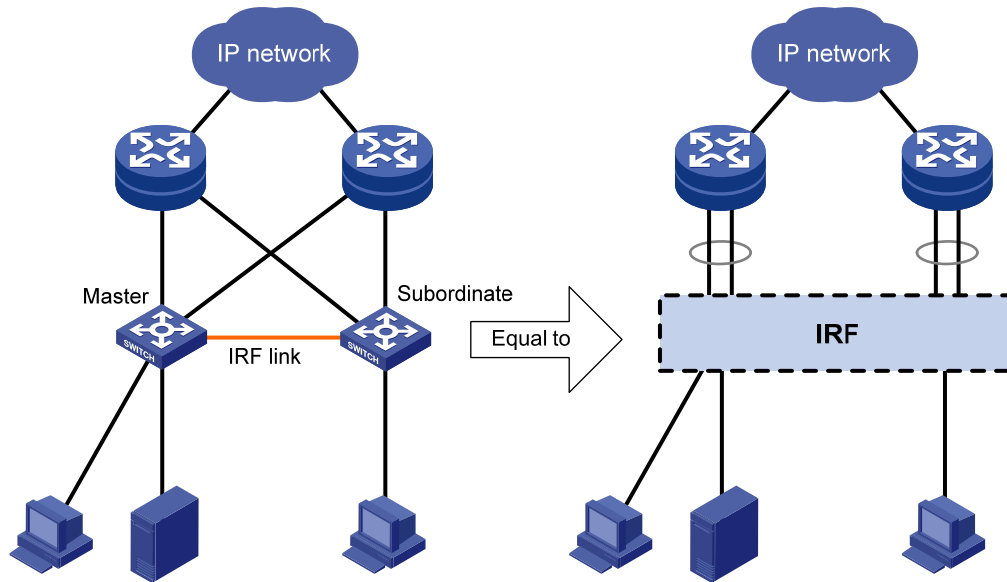
IRF provides the following benefits:

- **Simplified topology and easy management**—An IRF fabric appears as one node and is accessible at a single IP address on the network. You can use this IP address to log in at any member device to manage all the members of the IRF fabric. In addition, you do not need to run the spanning tree feature among the IRF members.
- **1:N redundancy**—In an IRF fabric, one member works as the master to manage and control the entire IRF fabric, and all the other members process services while backing up the master. When the master fails, all the other member devices elect a new master from among them to take over without interrupting services.
- **IRF link aggregation**—You can assign several physical links between neighboring members to their IRF ports to create a load-balanced aggregate IRF connection with redundancy.
- **Multiple-chassis link aggregation**—You can use the Ethernet link aggregation feature to aggregate the physical links between the IRF fabric and its upstream or downstream devices across the IRF members.
- **Network scalability and resiliency**—Processing capacity of an IRF fabric equals the total processing capacities of all the members. You can increase ports, network bandwidth, and processing capacity of an IRF fabric simply by adding member devices without changing the network topology.

Application scenario

Figure 1 shows an IRF fabric that is formed by two switches, which appear as a single node to the upper and lower layer devices.

Figure 1 IRF application scenario



Basic concepts

This section describes the basic concepts that you might encounter when working with IRF.

IRF member roles

IRF uses two member roles: master and slave (called "subordinate" throughout the documentation).

When switches form an IRF fabric, they elect a master to manage the IRF fabric, and all other switches back up the master. When the master switch fails, the other switches automatically elect a new master from among them to take over. For more information about master election, see "[Master election](#)."

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. For more information about interface and file path naming, see "[Interface naming conventions](#)" and "[File system naming conventions](#)."

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. The two IRF ports are referred to as "IRF-port 1" and "IRF-port 2" in this book for simplicity.

To use an IRF port, you must bind at least one physical port to it. The physical ports assigned to an IRF port automatically form 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.

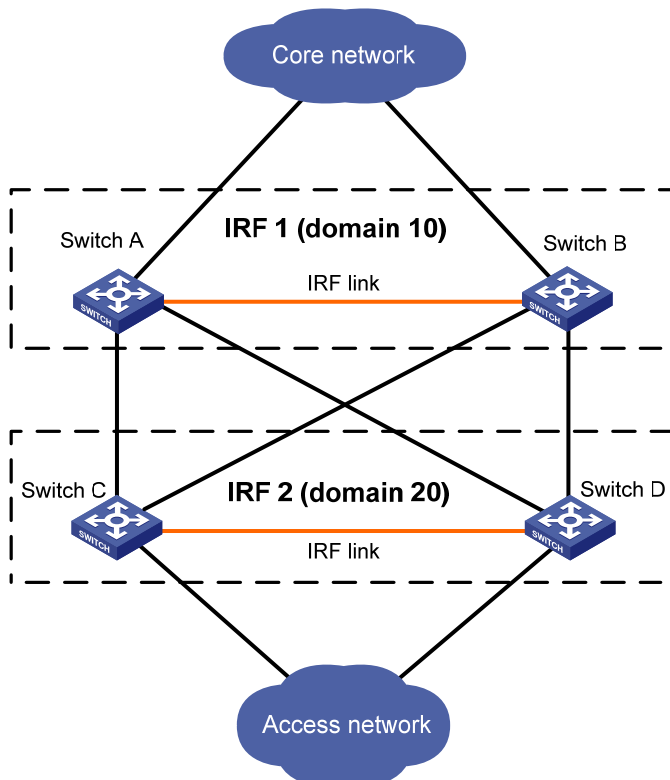
For more information about physical ports that can be used for IRF links, see "[General restrictions and configuration guidelines](#)."

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 2](#), Switch A and Switch B form IRF fabric 1, and Switch C and Switch D form IRF fabric 2. The fabrics have LACP MAD detection links between them. When a member switch in one 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 2 A network that comprises two IRF domains

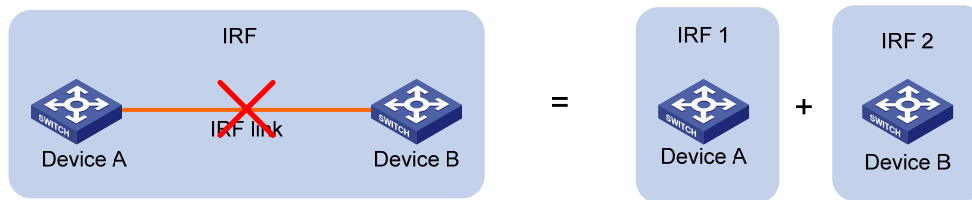


IRF split

IRF split occurs when an IRF fabric breaks up into two or more IRF fabrics because of IRF link failures, as shown in [Figure 3](#). The split IRF fabrics operate with the same IP address and cause routing and

forwarding problems on the network. To quickly detect a multi-active collision, configure at least one MAD mechanisms (see "[IRF multi-active detection](#)").

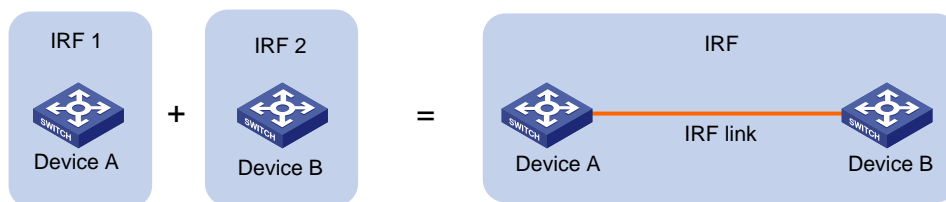
Figure 3 IRF split



IRF merge

IRF merge occurs when two split IRF fabrics reunite or when you configure and connect two independent IRF fabrics to be one IRF fabric, as shown in [Figure 4](#).

Figure 4 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.

Interface naming conventions

An interface is named in the format of *chassis-id/slot-number/port-index*, where:

- *chassis-id*—IRF member ID of the switch. This argument defaults to 1.
- *slot-number*—Represents the slot number of the interface card. This argument takes 0 for the fixed ports on the front panel, and takes 1 and 2 for the two interface cards (from left to right) on the rear panel.
- *port-index*—Port index depends on the number of ports available on the switch. To identify the index of a port, look at its port index mark on the chassis.

For one example, on the standalone switch Sysname, GigabitEthernet 1/0/1 represents the first fixed port on the front panel. Set its link type to trunk, as follows:

```
<Sysname> system-view
[Sysname] interface gigabitethernet 1/0/1
[Sysname-GigabitEthernet1/0/1] port link-type trunk
```

For another example, on the IRF fabric Master, GigabitEthernet 3/0/1 represents the first fixed port on the front panel of member switch 3. Set its link type to trunk, as follows:


```
<Master> system-view
[Master] interface gigabitethernet 3/0/1
[Master-GigabitEthernet3/0/1] port link-type trunk
```

File system naming conventions

On a standalone switch, you can use the name of storage device to access its file system. For more information about storage device naming conventions, see *Fundamentals Configuration Guide*.

On an IRF fabric, you can use the name of storage device to access the file system of the master. To access the file system of any other member switch, use the name in the **slotmember-ID#storage-device-name** format. For example:

To access the **test** folder under the root directory of the Flash on the master switch:

```
<Master> mkdir test
...
%Created dir flash:/test.
<Master> dir
Directory of flash:/
  0   -rw-  10105088  Apr 26 2000 13:44:57  test.app
  1   -rw-    2445    Apr 26 2000 15:18:19  config.cfg
  2   drw-      -    Jul 14 2008 15:20:35  test
30861 KB total (20961 KB free)
```

To create and access the **test** folder under the root directory of the Flash on member switch 3:

```
<Master> mkdir slot3#flash:/test
%Created dir slot3#flash:/test.
<Master> cd slot3#flash:/test
<Master> pwd
slot3#flash:/test
```

Or:

```
<Master> cd slot3#flash:/
<Master> mkdir test
%Created dir slot3#flash:/test.
```

To copy the file **test.app** on the master to the root directory of the Flash on member switch 3:

Display the current working path. In this example, the current working path is the root directory of the Flash on member switch 3.

```
<Master> pwd
slot3#flash:
```

Change the current working path to the root directory of the Flash on the master switch.

```
<Master> cd flash:/
<Master> pwd
flash:
```

Copy the file to member switch 3.

```
<Master> copy test.app slot3#flash:/
Copy flash:/test.app to slot3#flash:/test.app?[Y/N]:y
%Copy file flash:/test.app to slot3#flash:/test.app...Done.
```

Configuration synchronization mechanism

IRF uses a strict running-configuration synchronization mechanism so all chassis in an IRF fabric can work as a single node, and after the master fails, other members can operate normally.

In an IRF fabric, all chassis get and run the running configuration of the master. Any configuration you have made is propagated to all members.

When you execute the **save [safely] [backup | main] [force]** command or the **save file-url all** command, the system saves the running configuration, as follows:

- If the configuration auto-update function (the **slave auto-update config** command) is enabled, saves the configuration as the startup configuration on all member switches for the next startup.
- If the configuration auto-update function is disabled, saves the configuration as the startup configuration on the master for the next startup.

By default, configuration auto-update is enabled.

For more information about configuration management, see *Fundamentals Configuration Guide*.

Master election

Master election is held each time the IRF fabric topology changes, for example, when the IRF fabric is established, a new member device is plugged in, the master device fails or is removed, the IRF fabric splits, or IRF fabrics merge.

Master election uses the following rules in descending order:

1. Current master, even if a new member has higher priority.
When an IRF fabric is being formed, all member switches consider themselves as the master, and this rule is skipped
2. Member with higher priority.
3. Member with the longest system uptime.
4. Member with the lowest bridge MAC address.

The IRF fabric is formed on election of the master.

During an IRF merge, the switches of the IRF fabric that fails the master election must reboot to re-join the IRF fabric that wins the election.

After a master election, all subordinate switches reboot with the configuration on the master. Their configuration files do not take effect.

IRF multi-active detection

An IRF link failure causes an IRF fabric to split in two IRF fabrics operating with the same Layer 3 configurations, including the same IP address. To avoid IP address collision and network problems, IRF uses multi-active detection (MAD) mechanisms to detect the presence of multiple identical IRF fabrics, handle collisions, and recover from faults.

Multi-active handling procedure

The multi-active handling procedure includes detection, collision handling and failure recovery.

Detection

The MAD implementation of the switch detects active IRF fabrics with the same Layer 3 global configuration by extending the LACP, BFD, or gratuitous ARP protocol.

These MAD mechanisms identify each IRF fabric with a domain ID and an active ID (the member ID of the master). If multiple active IDs are detected in a domain, MAD determines that an IRF collision or split has occurred.

You can use at least one of these mechanisms in an IRF fabric, depending on your network topology. For a comparison of these MAD mechanisms, see "[Configuring MAD](#)."

Collision handling

When multiple identical active IRF fabrics are detected, MAD compares the member IDs of their masters. If the master in one IRF fabric has the lowest member ID among all the masters, the members in the fabric continue to operate in Active state and forward traffic. MAD sets all the other IRF fabrics in Recovery (disabled) state and shuts down all their physical ports except the console ports, physical IRF ports, and any ports you have specified with the **mad exclude interface** command.

Failure recovery

To merge two split IRF fabrics, first repair the failed IRF link and remove the IRF link failure.

If the IRF fabric in Recovery state fails before the failure is recovered, repair the failed IRF fabric and the failed IRF link.

If the IRF fabric in Active state fails before the failure is recovered, first enable the IRF fabric in Recovery state to take over the active IRF fabric and protect the services from being affected. After that, recover the MAD failure.

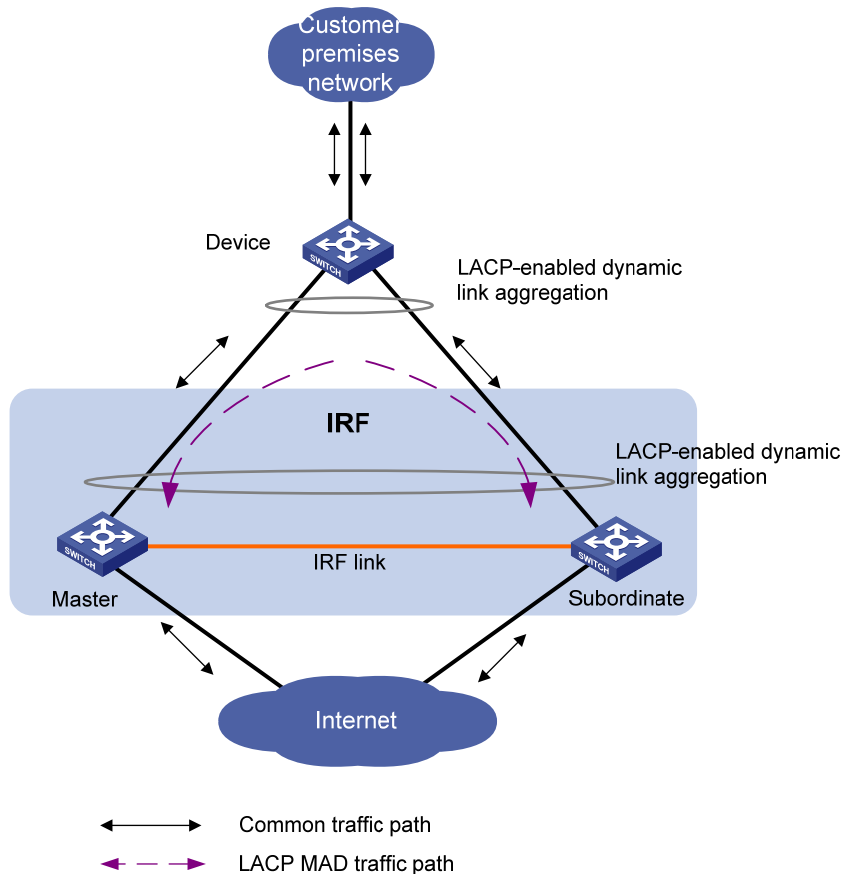
LACP MAD

LACP MAD requires that every IRF member have a link with an intermediate device, and all these links form a dynamic link aggregation group, as shown in [Figure 5](#). In addition, the intermediate device must be an HP device that supports extended LACP for MAD.

The IRF member switches send extended LACPDU s with TLVs that convey the domain ID and the active ID of the IRF fabric. The intermediate device transparently forwards the extended LACPDU s received from one member switch to all the other member switches:

- If the domain IDs and the active IDs in the extended LACPDU s sent by all the member devices are the same, the IRF fabric is integrated.
- If the extended LACPDU s convey the same domain ID but different active IDs, a split has occurred. To handle this situation, LACP MAD sets the IRF fabric with higher active ID in Recovery state, and shuts down all its physical ports but the console port, IRF ports, and any ports you have specified with the **mad exclude interface** command. The IRF fabric with lower active ID is still in Active state and forwards traffic.

Figure 5 LACP MAD application scenario



BFD MAD

BFD MAD can work with or without intermediate devices. Figure 6 shows a typical BFD MAD application scenario.

To use BFD MAD:

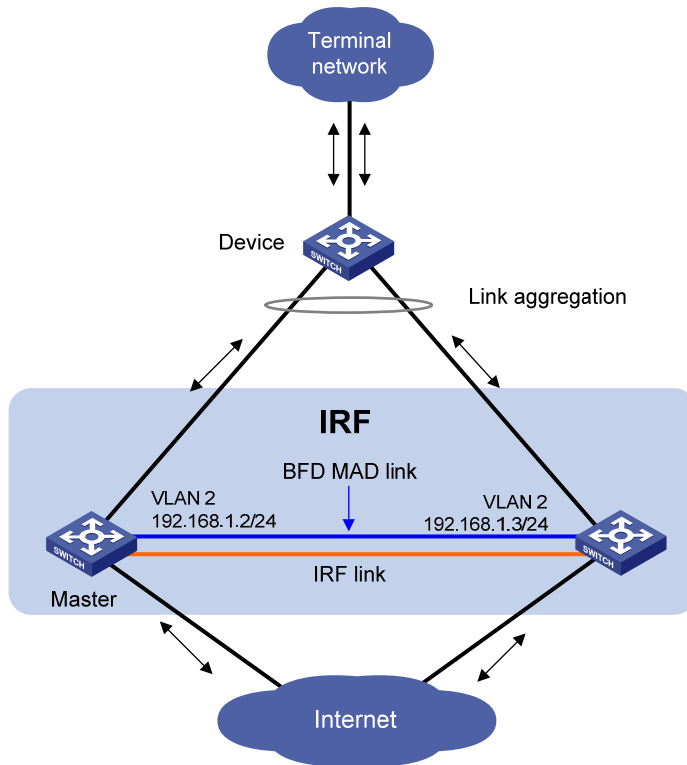
- Set up dedicated BFD MAD link between each pair of IRF members or between each IRF member and the intermediate device. Do not use the BFD MAD links for any other purpose.
- Assign the ports connected by BFD MAD links to the same VLAN, create a VLAN interface for the VLAN, and assign a MAD IP address to each member on the VLAN interface.

The MAD addresses identify the member switches and must belong to the same subnet.

With BFD MAD, the master tries to establish BFD sessions with other member switches by using its MAD IP address as the source IP address:

- If the IRF fabric is integrated, only the MAD IP address of the master is effective, and the master cannot establish a BFD session with any other member. If you execute the **display bfd session** command, the state of the BFD sessions is **Down**.
- When the IRF fabric splits, the IP addresses of the masters in the split IRF fabrics take effect, and the two masters can establish a BFD session. If you execute the **display bfd session** command, the state of the BFD session between the two devices is **Up**.

Figure 6 BFD MAD application scenario

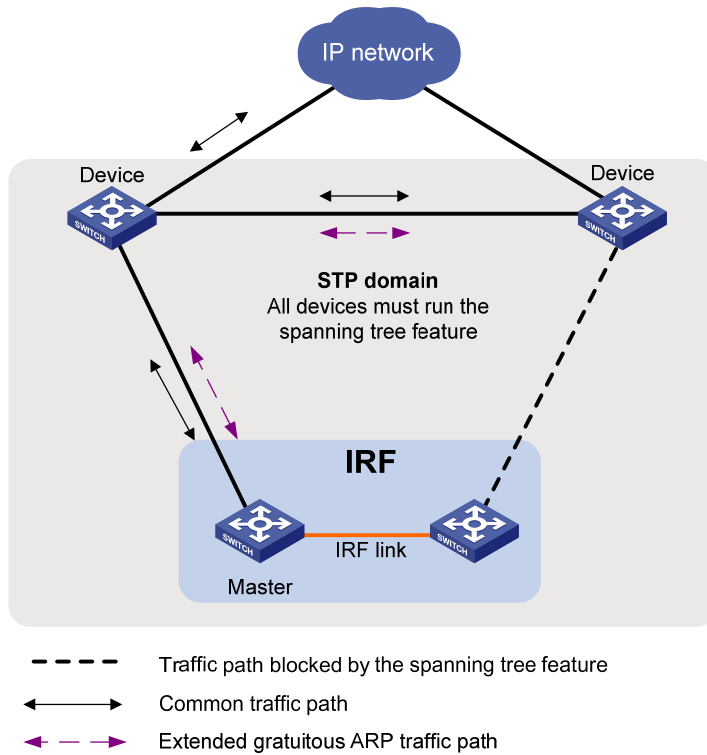


ARP MAD

ARP MAD detects multi-active collisions by using extended gratuitous ARP packets that convey the IRF domain ID and the active ID.

You can set up ARP MAD links between neighbor IRF member devices, or between each IRF member device and an intermediate device (see [Figure 7](#)). If an intermediate device is used, you must also run the spanning tree feature between the IRF fabric and the intermediate device.

Figure 7 ARP MAD application scenario



Each IRF member compares the domain ID and the active ID in incoming extended gratuitous ARP packets with its domain ID and active ID:

- If the domain IDs are different, the extended gratuitous ARP packet is from a different IRF fabric, and the device does not continue to process the packet with the MAD mechanism.
- If the domain IDs are the same, the device compares the active IDs:
 - If the active IDs are different, the IRF fabric has split.
 - If the active IDs are the same, the IRF fabric is integrated.

Configuring IRF

Read the configuration restrictions and guidelines carefully when you connect and set up an IRF fabric.

General restrictions and configuration guidelines

This section describes the restrictions and configuration guidelines you must follow.

Software requirements

All IRF member switches must run the same system software image version.

IRF link redundancy

The switch supports up to two physical ports for an IRF port, and these two ports must be located on the same interface card.

IRF physical port restrictions and cabling requirements

Candidate IRF physical ports include 10-GE ports on expansion interface cards. To use IRF, purchase at least one interface card. The following are the interface cards available for IRF connections:

- LSPM1XP1P (JD361B)
- LSPM1XP2P (JD359B)
- LSPM1CX2P (JD360B)
- LSPM2SP2P (JD368B)
- LSPM1XGT2P (JG535A)

For long-distance IRF connections, use XFP or SFP+ transceiver modules and fibers. For short-distance IRF connections, use 10 GE Ethernet twisted pairs, CX4 or SFP+ cables. For more information about transceiver modules available for an interface card, see the interface card user guide.

For more information about transceiver modules, see *HP A-Series Switches Transceiver Modules User Guide*.

The SFP+ modules and SFP+ cables available for the switch are subject to change over time. For the most up-to-date list of SFP+ modules and cables, contact HP technical support or marketing staff.

IRF port binding restrictions

Facing the rear panel, make sure the physical ports bound to IRF port 1 is on the left of the physical ports bound to IRF port 2.

Physical ports of an IRF port must be on the same interface card.

MAD

- Configure at least one MAD mechanism for prompt IRF split detection and IRF fabric recovery.

- If LACP MAD or ARP MAD runs between two IRF fabrics, assign each fabric a unique IRF domain ID. For BFD MAD, this task is optional.
- BFD MAD is available only on 5500 EI switches.
- To exclude a port from the shutdown action that is executed when an IRF fabric transits to the Recovery state, use the **mad exclude interface** command. To bring up a port after the IRF fabric transits to the Recovery state, you must use the **mad restore** command to activate the entire IRF fabric, and not the **undo shutdown** command.

FIPS mode requirement

To form an IRF fabric, all member devices must use the same FIPS mode setting (configurable with the **fips mode enable** command). For more information about FIPS mode, see *Security Configuration Guide*.

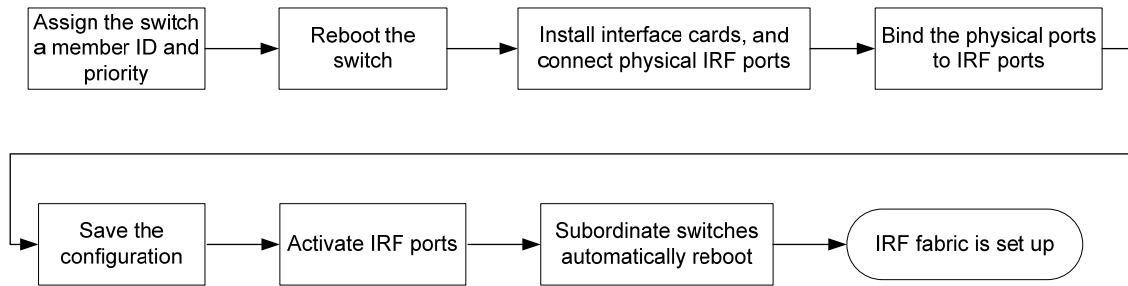
Other configuration guidelines

- Strictly follow the IRF fabric setup procedure described in "[Setup and configuration task list](#)" to plan the IRF fabric, identify IRF physical ports, connect IRF member switches, and configure basic settings.
- Assign each member a unique IRF member ID to make sure they can merge. 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.
- Before removing an interface card that has physical IRF ports in an IRF fabric, remove the IRF connection cables, or use the **shutdown** command to shut down the IRF physical ports.
- If a subordinate switch uses the same next-startup configuration file name as the master switch, the file might be overwritten depending on your configuration file management settings. To continue to use the configuration file after removing the switch from the IRF fabric, back up the file before setting up the IRF fabric.
- Save any configuration you have made to the startup configuration file before rebooting the IRF member devices.
- The Layer 3 Ethernet port in this book refers to an Ethernet port that can perform IP routing and inter-VLAN routing. You can set an Ethernet port as a Layer 3 Ethernet interface by using the **port link-mode route** command (see *Layer 2—LAN Switching Configuration Guide*).
- You can configure an Ethernet port as a Layer 3 Ethernet port only on the 5500 EI switch.

Setup and configuration task list

HP recommends the basic IRF setup procedure in [Figure 8](#). Perform the tasks in this figure on each member switch. After the IRF fabric is set up, you can access the IRF fabric to manage its member switches as if they were one switch.

Figure 8 Basic IRF setup flow chart



HP recommends the following IRF fabric setup and configuration procedure:

Task	Remarks
1. Planning the IRF fabric setup	Required.
2. Assigning a member ID to each IRF member switch	Required. Perform this task on each member switch.
3. Specifying a priority for each member switch	Required. Perform this task on each member switch.
4. Connecting physical IRF ports	Required.
5. Binding physical ports to IRF ports	Required. Perform this task on each member switch.
6. Accessing the IRF fabric: <ul style="list-style-type: none"> o Accessing the CLI of the master switch o Accessing the CLI of a subordinate switch 	Login to the master's CLI is required. You configure all member switches at the master's CLI. From the master's CLI, you can log in to any other member switch's CLI to execute a limited set of maintenance commands.
7. Assigning an IRF domain ID to the IRF fabric	This task is required for ARP MAD and LACP MAD.
8. Configuring a member switch description	Optional.
9. Configuring IRF link load sharing mode: <ul style="list-style-type: none"> o Configuring the global load sharing mode o Configuring a port-specific load sharing mode 	Optional.
10. Configuring IRF bridge MAC persistence	Optional.
11. Enabling software auto-update for system software image synchronization	Optional. HP recommends enabling software auto-update to make sure system software image synchronization
12. Setting the IRF link down report delay	Optional.

Task	Remarks
13. Configuring MAD: <ul style="list-style-type: none"> ○ Configuring LACP MAD ○ Configuring BFD MAD ○ Configuring ARP MAD ○ Excluding a port from the shutdown action upon detection of multi-active collision ○ Recovering an IRF fabric 	Required. MAD mechanisms are independent of one another. You can configure at least one MAD mechanism for an IRF fabric.

Planning the IRF fabric setup

Consider the following items when you plan an IRF fabric:

- Hardware compatibility and restrictions
- IRF fabric size
- Master switch
- IRF physical ports
- Member ID and priority assignment scheme
- Fabric topology and cabling scheme

For more information about hardware and cabling, see the switch installation guide.

Assigning a member ID to each IRF member switch

CAUTION:

In an IRF fabric, changing IRF member IDs might cause undesirable configuration changes and even data loss. Before you do that, back up the configuration and make sure you fully understand the impact on your network. For example, all member switches in an IRF fabric are the same model. If you swapped the IDs of any two members, their interface settings would also be swapped.

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.

Perform this task before the IRF fabric is formed. To prevent any undesirable configuration change or data loss, avoid changing member IDs after the IRF fabric is formed.

The new member ID takes effect at a reboot. After the switch reboots, the settings on all member-ID related physical resources (including common physical network ports) are removed and require reconfiguration, regardless of whether you have saved the configuration.

To set a member ID for a switch:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Assign an IRF member ID to the switch.	irf member <i>member-id</i> renumber <i>new-member-id</i>	The default IRF member ID is 1.

Step	Command	Remarks
3. Save the configuration.	<code>save [safely] [backup main] [force]</code>	Optional. If you have bound physical ports to IRF ports or assigned member priority, save the configuration before rebooting the switch so these settings can continue to take effect after the reboot.
4. Reboot the switch.	<code>reboot [slot slot-number]</code>	N/A

Specifying a priority for each member switch

IRF member priority represents the possibility for a device to be elected the master in an IRF fabric. The higher the priority, the higher the possibility.

A member priority change affects the election result at the next master election, but does not cause immediate master re-election.

To specify a priority for the switch:

Step	Command	Remarks
1. Enter system view.	<code>system-view</code>	N/A
2. Specify a priority for the switch.	<code>irf member member-id priority priority</code>	The default IRF member priority is 1.

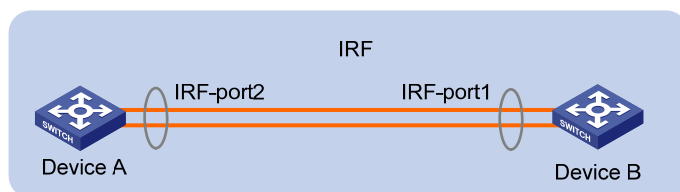
Connecting physical IRF ports

When you connect two neighboring IRF members, connect the physical ports of IRF-port 1 on one member to the physical ports of IRF-port 2 on the other, as shown in [Figure 9](#).

! IMPORTANT:

No intermediate devices are allowed between neighboring members.

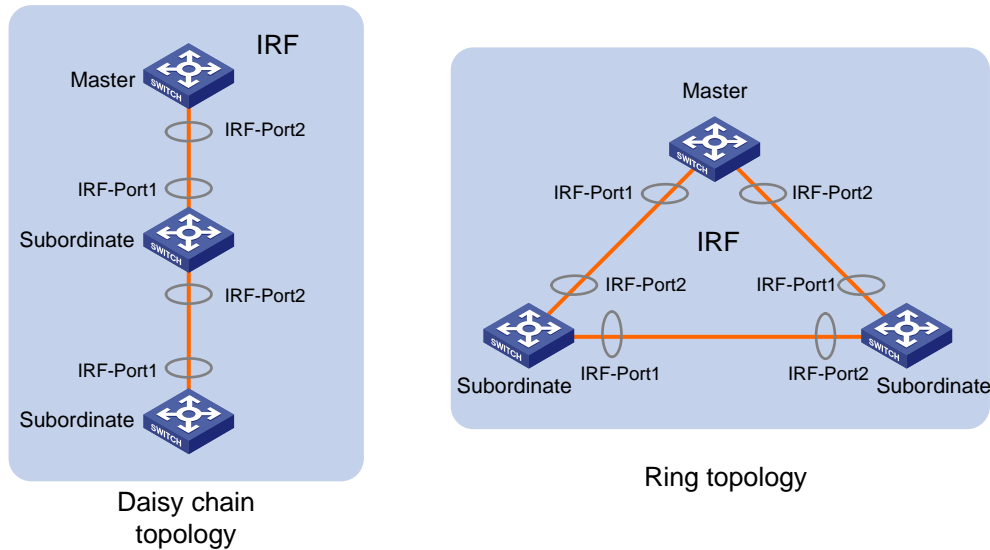
Figure 9 Connecting IRF physical ports



Connect the switches into a daisy chain topology or a ring topology, which is more reliable (see [Figure 10](#)). In ring topology, the failure of one IRF link does not cause the IRF fabric to split as in daisy chain topology. Rather, the IRF fabric changes to a daisy chain topology without interrupting network services.

To use the ring topology, you must have at least three member switches.

Figure 10 Daisy chain topology versus ring topology



Binding physical ports to IRF ports

To establish an IRF connection between two devices, you must bind at least one physical port to IRF-port 1 on one device and to IRF-port 2 on the other. For link redundancy and load sharing, bind multiple physical ports to one IRF port.

When you bind physical ports to IRF ports, follow the restrictions in "[IRF port binding restrictions](#)."

On a physical port that has been bound to an IRF port, you can only use the **cfd**, **default**, **shutdown**, **description**, and **flow-interval** commands. For more information about these commands, see *Layer 2—LAN Switching Command Reference*.

To bind physical ports to IRF ports:

Step	Command	Remarks
1. Enter system view.	system-view	N/A

Step	Command	Remarks
2. Enter Ethernet interface view or interface range view.	<ul style="list-style-type: none"> • Enter interface range view: <ul style="list-style-type: none"> ○ Method 1: interface range { <i>interface-type</i> <i>interface-number</i> [to <i>interface-type</i> <i>interface-number</i>] } &<1-5> ○ Method 2: interface range name <i>name</i> [interface { <i>interface-type</i> <i>interface-number</i> [to <i>interface-type</i> <i>interface-number</i>] } &<1-5>] • Enter interface view: interface <i>interface-type</i> <i>interface-number</i> 	<p>To shut down a range of physical IRF ports, enter interface range view.</p> <p>To shut down one physical IRF port, enter its interface view.</p>
3. Shut down the port or ports.	shutdown	<p>Always shut down a physical port before binding it to an IRF port or removing the binding.</p> <p>Start the shutdown operation on the master and then the switch that has the fewest number of hops from the master.</p>
4. Return to system view.	quit	N/A
5. Enter IRF port view.	irf-port <i>member-id/port-number</i>	N/A
6. Bind each physical port to the IRF port.	port group interface <i>interface-type</i> <i>interface-number</i> [mode { enhanced normal }]	<p>By default, no physical port is bound to any IRF port.</p> <p>The switch does not support the enhanced keyword.</p>
7. Return to system view.	quit	N/A

Step	Command	Remarks
8. Enter Ethernet interface view or interface range view.	<ul style="list-style-type: none"> • Enter interface range view: <ul style="list-style-type: none"> ○ Method 1: <code>interface range</code> <code>{ interface-type</code> <code>interface-number [to</code> <code>interface-type</code> <code>interface-number] }</code> <code>&<1-5></code> ○ Method 2: <code>interface range name</code> <code>name [interface</code> <code>{ interface-type</code> <code>interface-number [to</code> <code>interface-type</code> <code>interface-number] }</code> <code>&<1-5>]</code> • Enter interface view: <code>interface interface-type</code> <code>interface-number</code> 	N/A
9. Bring up the port or ports.	<code>undo shutdown</code>	N/A
10. Return to system view.	<code>quit</code>	N/A
11. Save the running configuration.	<code>save</code>	N/A
12. Activate the IRF port configuration.	<code>irf-port-configuration active</code>	<p>After this step is performed, the state of the IRF port changes to UP, the member switches automatically elect a master, and the subordinate switch automatically reboots.</p> <p>After the IRF fabric is formed, you can add more physical ports to an IRF port (in UP state) without performing this step.</p>

Accessing the IRF fabric

The IRF fabric appears as one device after it is formed. You configure and manage all IRF members at the CLI of the master. All settings you have made are automatically propagated to the IRF members.

When you log in to an IRF fabric, you are placed at the CLI of the master, regardless of at which member switch you are logged in. After that, you can access the CLI of a subordinate switch to execute a limited set of maintenance commands.

The IRF fabric supports up to 16 concurrent VTY users. The maximum number of concurrent console users equals the total number of member switches in the IRF fabric.

Accessing the CLI of the master switch

Access an IRF fabric in one of the following ways:

- **Local login**—Log in through the console port of any member switch.

- **Remote login**—Remotely log in at a Layer 3 Ethernet interface on any member switch by using a methods including Telnet, Web, and SNMP.

For more information, see the chapter on login in *Fundamentals Configuration Guide*.

Accessing the CLI of a subordinate switch

You can log in to the CLI of a subordinate switch for maintenance or debugging. At the CLI of a subordinate switch, you are placed in user view, and the command prompt changes to `<Sysname-Slave#member-ID/slot-number>`, for example, `<Sysname-Slave#2>`. You can use the following commands at a subordinate switch's CLI:

- **display**
- **quit**
- **return**
- **system-view**
- **debugging**
- **terminal debugging**
- **terminal logging**
- **terminal monitor**
- **terminal trapping**

Perform the following task in user view:

Task	Command	Remarks
Log in to a subordinate switch.	irf switch-to <i>member-id</i>	By default, you are placed at the master's CLI.

To return to the master's CLI, use the **quit** command.

Assigning an IRF domain ID to the IRF fabric

This task is required for running LACP MAD or ARP MAD between two IRF fabrics. For BFD MAD, this task is optional.

To assign a domain ID to an IRF fabric:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Assign a domain ID to the IRF fabric.	irf domain <i>domain-id</i>	By default, the domain ID of an IRF fabric is 0.

Configuring a member switch description

You can configure a description to describe the location or purpose of a member switch for easy management.

To configure a description for a member switch:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the description of a member.	irf member <i>member-id</i> description <i>text</i>	By default, no member switch description is configured.

Configuring IRF link load sharing mode

On an IRF port that has multiple links, traffic is balanced across its physical links. You can configure the IRF port to distribute traffic based on certain criteria, including source IP address, destination IP address, source MAC address, destination MAC address, or any combination of the criteria. If a criteria combination is not supported, the system displays an error message.

Configure the IRF link load sharing mode for IRF links in system view or IRF port view.

- In system view, the configuration is global and takes effect on all IRF ports.
- In IRF port view, the configuration is port specific and takes effect only on the specific IRF port.

An IRF port preferentially uses the port-specific load sharing mode. If no port-specific load sharing mode is available, an IRF port uses the global load sharing mode.

Configuring the global load sharing mode

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the global IRF link load sharing mode.	irf-port load-sharing mode { destination-ip destination-mac source-ip source-mac } *	By default, the switch uses source and destination MAC addresses for non-IP packets, and uses source and destination IP addresses for IP packets.

Configuring a port-specific load sharing mode

Before you configure a port-specific load sharing mode, make sure you have bound at least two physical ports to the IRF port.

To configure a port-specific load sharing mode for an IRF port:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter IRF port view.	irf-port <i>member-id/port-number</i>	N/A
3. Configure the port-specific load sharing mode.	irf-port load-sharing mode { destination-ip destination-mac source-ip source-mac } *	By default, the switch uses source and destination MAC addresses for non-IP packets, and uses source and destination IP addresses for IP packets.

Configuring IRF bridge MAC persistence

An IRF fabric by default uses the bridge MAC address of the master switch as its bridge MAC address. This bridge MAC address is used by Layer 2 protocols, for example, LACP, to identify the IRF fabric, and must be unique on a switched LAN for proper communication.

To avoid duplicate bridge MAC addresses, an IRF fabric can automatically change its bridge MAC address after its master leaves, but the change can cause transient traffic disruption.

Depending on your network condition, enable the IRF fabric to retain or change its bridge MAC address after the master leaves. Available options include:

- **irf mac-address persistent timer**—Bridge MAC address of the IRF fabric is retained for six minutes after the master leaves. If the master does not return before the timer expires, the IRF fabric uses the bridge MAC address of the new master as its bridge MAC address. This option avoids unnecessary bridge MAC address change due to a device reboot, transient link failure, or purposeful link disconnection.
- **irf mac-address persistent always**—Bridge MAC address of the IRF fabric does not change after the master leaves.
- **undo irf mac-address persistent**—Bridge MAC address of the new master replaces the original one as soon as the old master leaves.

⚠ IMPORTANT:

If ARP MAD is used, configure the **undo irf mac-address persistent** command to enable immediate bridge MAC address change after a master leaves.

If two IRF fabrics have the same bridge MAC address, they cannot merge.

To configure the IRF bridge MAC persistence setting:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure IRF bridge MAC persistence.	<ul style="list-style-type: none">• Retain the bridge MAC address even if the master has changed: irf mac-address persistent always• Preserve the bridge MAC address for six minutes after the master leaves: irf mac-address persistent timer• Change the bridge MAC address as soon as the master leaves: undo irf mac-address persistent	By default, the IRF fabric's bridge MAC address is retained permanently even after the master leaves.

Enabling software auto-update for system software image synchronization

To join an IRF fabric, a switch must use the same system software image as the master in the fabric.

The software auto-update function automatically propagates the system software image of the master to all members in the IRF fabric. If software auto-update is disabled, you must manually update the switch with the system software image of the master.

When you add a switch to the IRF fabric, the software auto-update function compares the system software versions of the switch and the IRF master. If the versions are different, the switch automatically downloads the system software image from the master, sets the downloaded image as the system software for the next startup, and automatically reboots with the new system software image to re-join the IRF fabric.

Before you use the software auto-update function, make sure:

- The switch you are adding to the IRF fabric is compatible with the software version running on the master. If not, the automatic system software upgrading function cannot correctly work.
- The switch you are adding to the IRF fabric has sufficient space for the new system software image.

To enable the IRF fabric to automatically synchronize the system software of the master to the switch you are adding to the IRF fabric:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable the software auto-update function.	irf auto-update enable	By default, this function is disabled.

In an IRF fabric enabled with software auto-update, if a software upgrade requires upgrading the Boot ROM image, use the following upgrading procedure:

1. Download the new system software image to the master device.
2. Use the **bootrom update** command to upgrade the Boot ROM image on the master.
This step guarantees that the master can complete startup prior to other member switches.
3. Use the **boot-loader file file-url slot slot-number main** command to specify the system software image as the startup image for the master.
4. Reboot the entire IRF fabric to complete upgrading software.

For the system software image and Boot ROM compatibility, see the release notes for the new software release.

Setting the IRF link down report delay

You can avoid IRF link flapping causing frequent IRF splits and merges during a short time by configuring the IRF ports to delay reporting link down events. An IRF port works as follows:

- When the IRF link changes from up to down, the port does not immediately report the change to the IRF fabric. If the IRF link state is still down when the delay time is reached, the port reports the change to the IRF fabric.
- When the IRF link changes from down to up, the link layer immediately reports the event to the IRF fabric.

To set the IRF link down report delay:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Set the IRF link down report delay.	irf link-delay <i>interval</i>	The default IRF link down report delay is 4 seconds. If your IRF fabric requires a quick master/subordinate or IRF link switchover or has deployed the BFD or GR feature, HP recommends setting the delay to 0 seconds.

Configuring MAD

The following MAD mechanisms are available for detecting multi-active collisions in different network scenarios:

- LACP MAD
- BFD MAD
- ARP MAD

These MAD detection mechanisms operate independently. You can configure all of them for an IRF fabric.

Table 1 provides a reference for you to make a MAD mechanism selection decision.

Table 1 A comparison of the MAD mechanisms

MAD mechanism	Advantages	Disadvantages	Application scenario
LACP MAD	<ul style="list-style-type: none"> • Detection speed is fast. • Requires no MAD-dedicated physical ports or interfaces. 	Requires an intermediate HP device that supports LACP MAD packets.	<p>Link aggregation is used between the IRF fabric and its upstream or downstream device.</p> <p>For information about LACP, see <i>Layer 2—LAN Switching Configuration Guide</i>.</p>
BFD MAD	<ul style="list-style-type: none"> • Detection speed is fast. • No intermediate device is required. • Intermediate device, if used, can come from any vendor. 	<ul style="list-style-type: none"> • Requires MAD dedicated physical ports and Layer 3 interfaces, which cannot be used for transmitting user traffic. • If no intermediate device is used, the IRF members must be fully meshed. • If an intermediate device is used, every IRF member must connect to the intermediate device. 	<ul style="list-style-type: none"> • Suitable for various network scenarios. • If no intermediate device is used, this mechanism is only suitable for IRF fabrics that have a small number of members that are geographically close to one another. <p>For information about BFD, see <i>High Availability Configuration Guide</i>.</p>

MAD mechanism	Advantages	Disadvantages	Application scenario
ARP MAD	<ul style="list-style-type: none"> No intermediate device is required. Intermediate device, if used, can come from any vendor. Requires no MAD dedicated ports. 	<ul style="list-style-type: none"> Detection speed is slower than BFD MAD and LACP MAD. Spanning tree feature must be enabled. 	<p>Spanning tree-enabled non-link aggregation IPv4 network scenario.</p> <p>For information about ARP, see <i>Layer 3—IP Services Configuration Guide</i>.</p>

Configuring LACP MAD

When you use LACP MAD, follow these guidelines:

- The intermediate device must be an HP device that support extended LACP for MAD.
- If the intermediate device is in an IRF fabric, assign this fabric a different domain ID than the LACP MAD-enabled fabric to avoid false detection of IRF split.
- Use dynamic link aggregation mode. MAD is LACP dependent. Even though LACP MAD can be configured on both static and dynamic aggregate interfaces, it takes effect only on dynamic aggregate interfaces.
- Configure link aggregation settings also on the intermediate device.

To configure LACP MAD:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Assign a domain ID to the IRF fabric.	irf domain <i>domain-id</i>	The default IRF domain ID is 0.
3. Create an aggregate interface and enter aggregate interface view.	<ul style="list-style-type: none"> Enter Layer 2 aggregate interface view: interface bridge-aggregation <i>interface-number</i> Enter Layer 3 aggregate interface view: interface route-aggregation <i>interface-number</i> 	<p>Use either command.</p> <p>Perform this step also on the intermediate device.</p> <p>The 5500 SI switch series does not support Layer 3 aggregate interfaces.</p>
4. Configure the aggregation group to operate in dynamic aggregation mode.	link-aggregation mode dynamic	<p>By default, an aggregation group operates in static aggregation mode.</p> <p>Perform this step also on the intermediate device.</p>
5. Enable LACP MAD.	mad enable	By default, LACP MAD is disabled.
6. Return to system view.	quit	N/A
7. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A

Step	Command	Remarks
8. Assign the Ethernet interface to the specified aggregation group.	port link-aggregation group <i>number</i>	Perform this step on the intermediate device as well.

Configuring BFD MAD

When you use BFD MAD, follow these guidelines:

- You cannot enable BFD MAD on VLAN-interface 1.
- Do not use the BFD MAD VLAN for any other purpose. No Layer 2 or Layer 3 features, including ARP and LACP, can work on the BFD MAD-enabled VLAN interface or any port in the VLAN. If you configure any other feature on the VLAN, neither the configured feature nor the BFD MAD function can work correctly.
- Do not enable the spanning tree feature on any port in the BFD MAD VLAN. The MAD function is mutually exclusive with the spanning tree feature.
- Do not bind a BFD MAD-enabled VLAN interface to any VPN. The MAD function is mutually exclusive with VPN.
- To avoid problems, make sure a BFD MAD-enabled VLAN interface has only IP addresses assigned with the **mad ip address** command. Do not use any other command, including the **ip address** command, to assign an IP address to the interface.
- All MAD IP addresses on the BFD MAD-enabled VLAN interface must be on the same subnet.
- If an intermediate device is used, assign the ports of BFD MAD links to the BFD MAD VLAN on the device.

To configure BFD MAD:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Create a VLAN dedicated to BFD MAD.	vlan <i>vlan-id</i>	The default VLAN on the switch is VLAN 1.
3. Return to system view.	quit	N/A
4. Enter Ethernet interface view or interface range view.	<ul style="list-style-type: none"> • Enter interface range view: <ul style="list-style-type: none"> ◦ Method 1: interface range { <i>interface-type interface-number</i> [to <i>interface-type interface-number</i>] } &<1-5> ◦ Method 2: interface range name <i>name</i> [interface { <i>interface-type interface-number</i> [to <i>interface-type interface-number</i>] } &<1-5>] • Enter Ethernet interface view: interface <i>interface-type interface-number</i> 	<p>To assign a range of ports to the BFD MAD VLAN, enter interface range view.</p> <p>To assign one port to the BFD MAD VLAN, enter Ethernet interface view.</p>

Step	Command	Remarks
5. Assign the port or the range of ports to the BFD MAD VLAN.	<ul style="list-style-type: none"> Assign the port to the VLAN as an access port: port access vlan <i>vlan-id</i> Assign the port to the VLAN as a trunk port: port trunk permit vlan <i>vlan-id</i> Assign the port to the VLAN as a hybrid port: port hybrid vlan <i>vlan-id</i> { tagged untagged } 	<p>Choose one command depending on the port type.</p> <p>BFD MAD has no requirement for the link type of the detection port.</p> <p>The default link type of a port is access.</p>
6. Return to system view.	quit	N/A
7. Enter VLAN interface view.	interface <i>vlan-interface</i> <i>interface-number</i>	N/A
8. Enable BFD MAD.	mad bfd enable	By default, BFD MAD is disabled.
9. Configure a MAD IP address for a member device on the VLAN interface.	mad ip address <i>ip-address</i> { <i>mask</i> <i>mask-length</i> } member <i>member-id</i>	<p>By default, no MAD IP address is configured on any VLAN interface.</p> <p>The MAD IP address must not be on the same subnet as any IP address configured on any member device.</p>

NOTE:

After a BFD MAD-enabled IRF fabric splits, route conflict messages (for example, %May 5 16:15:47:733 2010 HP ARP/3/ROUTECONFLICT: Slot=5;Route conflict found, IP:192.168.2.1, VrfIndex:0) might appear on the fabric part that does not have the original master, because this fabric part still keeps the forwarding entries with the original master as the destination. This message does not affect forwarding. The system stops generating it after the forwarding entries are aged out.

Configuring ARP MAD

When you configure ARP MAD, follow these guidelines:

- Use a VLAN dedicated to ARP MAD.
- If an intermediate device is used, you can use common data links as ARP MAD links. If no intermediate device is used, set up dedicated ARP MAD links between IRF member devices.
- If an intermediate device is used, make sure the following requirements are met:
 - Run the spanning tree feature between the IRF fabric and the intermediate device.
 - Enable the IRF fabric to change its bridge MAC address as soon as the master leaves.
 - Create an ARP MAD VLAN and assign the ports on the ARP MAD links to the VLAN.
 - If the intermediate device is in an IRF fabric, assign this fabric a different domain ID than the ARP MAD-enabled fabric to avoid false detection of IRF splits.

To configure ARP MAD:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Assign a domain ID to the IRF fabric.	irf domain <i>domain-id</i>	The default IRF domain ID is 0.
3. Create a VLAN dedicated to ARP MAD.	vlan <i>vlan-id</i>	The default VLAN on the device is VLAN 1.
4. Return to system view.	quit	N/A
5. Enter Ethernet interface view or interface range view.	<ul style="list-style-type: none"> • Enter interface range view: <ul style="list-style-type: none"> ◦ Method 1: interface range { <i>interface-type</i> <i>interface-number</i> [to <i>interface-type</i> <i>interface-number</i>] } &<1-5> ◦ Method 2: interface range name <i>name</i> [interface { <i>interface-type</i> <i>interface-number</i> [to <i>interface-type</i> <i>interface-number</i>] } &<1-5>] • Enter Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> 	<p>To assign a range of ports to the BFD MAD VLAN, enter interface range view.</p> <p>To assign one port to the BFD MAD VLAN, enter Ethernet interface view.</p>
6. Assign the port or the range of ports to the ARP MAD VLAN.	<ul style="list-style-type: none"> • Assign the port to the VLAN as an access port: port access vlan <i>vlan-id</i> • Assign the port to the VLAN as a trunk port: port trunk permit vlan <i>vlan-id</i> • Assign the port to the VLAN as a hybrid port: port hybrid vlan <i>vlan-id</i> { tagged untagged } 	<p>Choose one command depending on the port type. ARP MAD has no requirement for the link type.</p> <p>The default link type of a port is access.</p>
7. Return to system view.	quit	N/A
8. Enter VLAN interface view.	interface vlan-interface <i>interface-number</i>	N/A
9. Assign the interface an IP address.	ip address <i>ip-address</i> { <i>mask</i> <i>mask-length</i> }	By default, no IP address is assigned to any VLAN interface.
10. Enable ARP MAD.	mad arp enable	By default, ARP MAD is disabled.

Excluding a port from the shutdown action upon detection of multi-active collision

CAUTION:

Excluding a VLAN interface and its Layer 2 ports from the shutdown action introduces IP collision risks because the VLAN interface might be active on both the IRF fabric in Active state and the IRF fabric in Recovery state.

By default, all ports except the console and physical IRF ports automatically shut down when the IRF fabric transits to the Recovery state.

You can exclude a port from the shutdown action for management or other special purposes. For example:

- Exclude a port from the shutdown action, so you can Telnet to the port for managing the switch.
- Exclude a VLAN interface and its Layer 2 ports from the shutdown action, so you can log in through the VLAN interface.

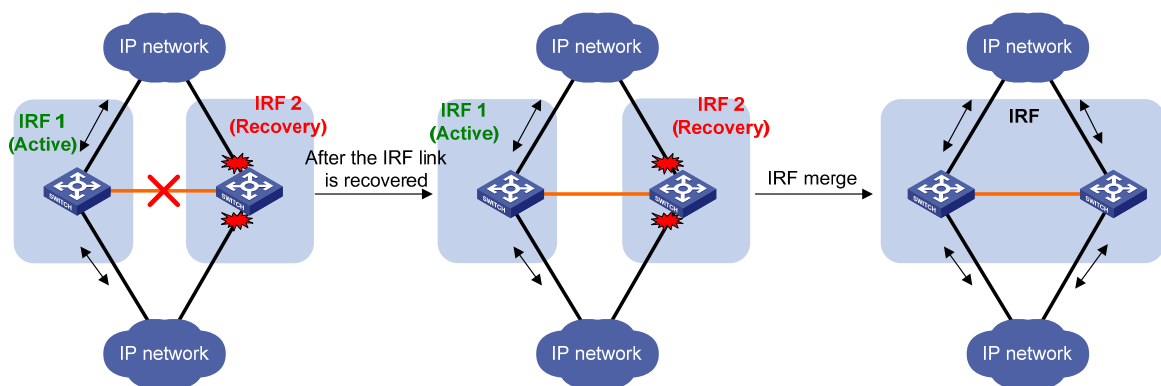
To configure a port to not shut down when the IRF fabric transits to the recovery state:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure a port to not shut down when the IRF fabric transits to the Recovery state.	mad exclude interface <i>interface-type interface-number</i>	By default, when an IRF fabric transits to the Recovery state, all its network ports but the IRF physical ports and console port are shut down.

Recovering an IRF fabric

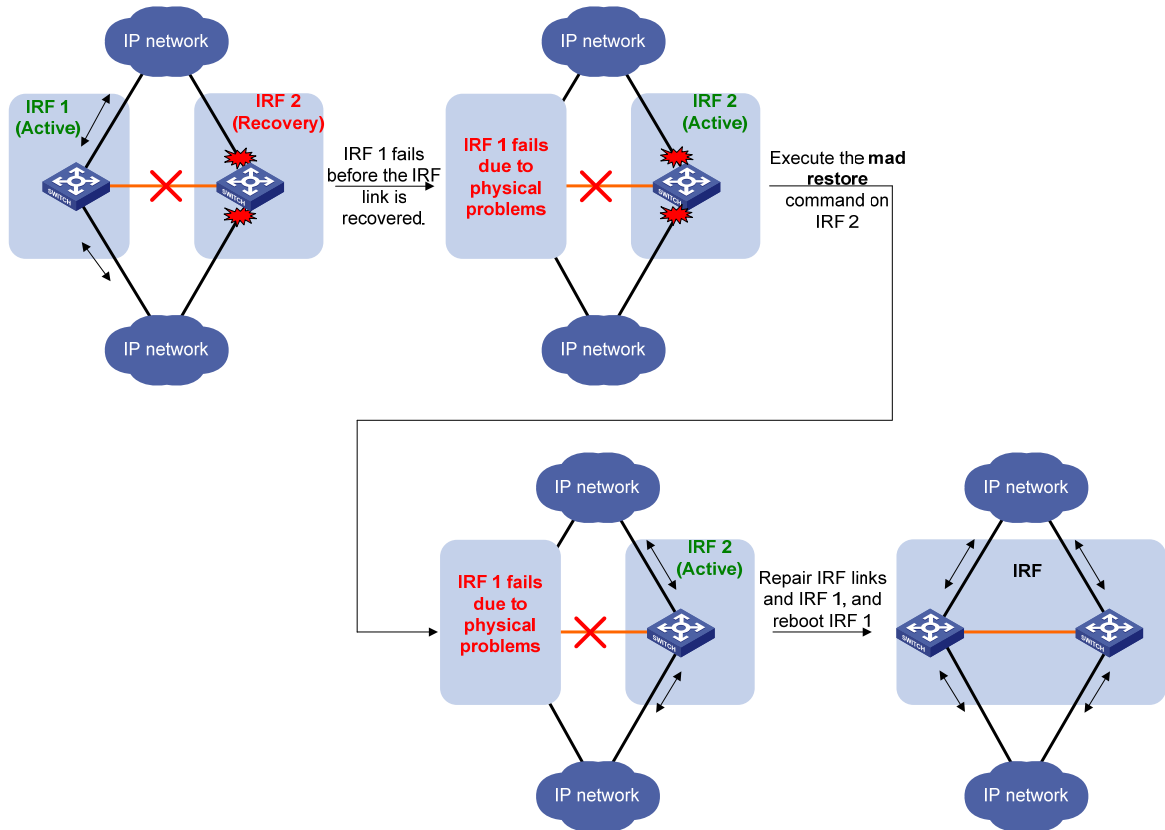
After the failed IRF link between two split IRF fabrics is recovered, log in to the Recovery-state IRF fabric, and use the **reboot** command to reboot all its members. After these member switches join the Active-state IRF fabric as subordinates, IRF merge is complete, as shown in Figure 11.

Figure 11 Recovering the IRF fabric



If the Active-state fabric has failed, for example, because of device or link failures, before the IRF link is recovered (see Figure 12), use the **mad restore** command on the Recovery-state fabric to change its state to Active for forwarding traffic. After you repair the IRF link, the two parts merge into a unified IRF fabric.

Figure 12 Active-state IRF fabric fails before the IRF link is recovered



To manually recover an IRF fabric in Recovery state:

Step	Command
1. Enter system view.	system-view
2. Change the state of the IRF fabric from Recovery to Active.	mad restore

After the IRF fabric is recovered, all ports that have been shut down by MAD automatically come up.

Displaying and maintaining an IRF fabric

Task	Command	Remarks
Display information about all IRF members.	display irf [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the IRF fabric topology.	display irf topology [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display basic IRF settings.	display irf configuration [{ begin exclude include } <i>regular-expression</i>]	Available in any view

Task	Command	Remarks
Display the load sharing mode for IRF links.	display irf-port load-sharing mode [irf-port [<i>member-id/port-number</i>]] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the master/subordinate switchover state of IRF members.	display switchover state [slot <i>member-id</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display MAD configuration.	display mad [verbose] [{ begin exclude include } <i>regular-expression</i>]	Available in any view

Configuration examples

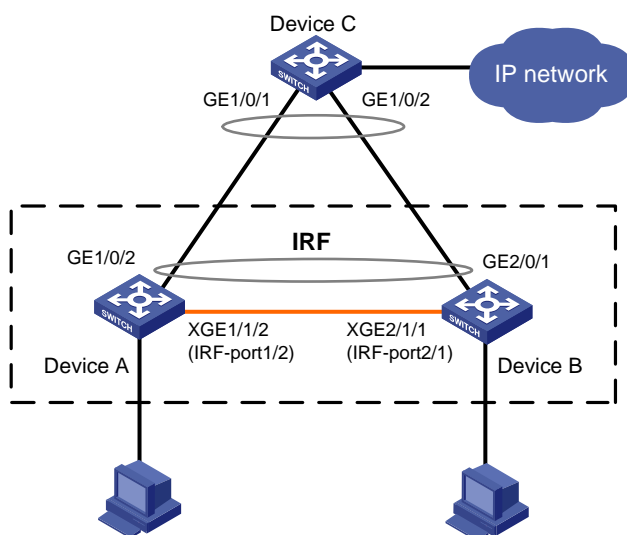
This section provides IRF configuration examples for IRF fabrics that use different MAD mechanisms.

LACP MAD-enabled IRF configuration example

Network requirements

Set up a two-member IRF fabric at the access layer of the enterprise network in [Figure 13](#). Configure LACP MAD in the IRF fabric, because the IRF fabric has a multi-chassis aggregate link to Device C, an HP device that supports extended LACP.

Figure 13 Network diagram



Configuration procedure

This example assumes that the system names of Device A, Device B and Device C are **DeviceA**, **DeviceB**, and **DeviceC** respectively before the IRF fabric is formed.

- Assign member IDs:
Keep the default member ID of Device A unchanged.

Change the member ID of Device B to 2.

```
<DeviceB> system-view
[DeviceB] irf member 1 renumber 2
Warning: Renumbering the switch number may result in configuration change or loss.
Continue? [Y/N]:y
[DeviceB]
```

2. Power off the devices, connect IRF links as shown in [Figure 13](#), and power on the two devices.
3. Configure IRF port bindings:

Bind Ten-GigabitEthernet 1/1/2 to IRF-port 1/2 on Device A and save the configuration.

```
<DeviceA> system-view
[DeviceA] interface ten-gigabitethernet 1/1/2
[DeviceA-Ten-GigabitEthernet1/1/2] shutdown
[DeviceA-Ten-GigabitEthernet1/1/2] quit
[DeviceA] irf-port 1/2
[DeviceA-irf-port1/2] port group interface ten-gigabitethernet 1/1/2
[DeviceA-irf-port1/2] quit
[DeviceA] interface ten-gigabitethernet 1/1/2
[DeviceA-Ten-GigabitEthernet1/1/2] undo shutdown
[DeviceA-Ten-GigabitEthernet1/1/2] save
```

Bind Ten-GigabitEthernet 2/1/1 to IRF-port 2/1 on Device B and save the configuration.

```
<DeviceB> system-view
[DeviceB] interface ten-gigabitethernet 2/1/1
[DeviceB-Ten-GigabitEthernet2/1/1] shutdown
[DeviceB-Ten-GigabitEthernet2/1/1] quit
[DeviceB] irf-port 2/1
[DeviceB-irf-port2/1] port group interface ten-gigabitethernet 2/1/1
[DeviceB-irf-port2/1] quit
[DeviceB] interface ten-gigabitethernet 2/1/1
[DeviceB-Ten-GigabitEthernet2/1/1] undo shutdown
[DeviceB-Ten-GigabitEthernet2/1/1] save
```

Activate IRF port configuration on Device A.

```
[DeviceA-Ten-GigabitEthernet1/1/2] quit
[DeviceA] irf-port-configuration active
```

Activate IRF port configuration on Device B.

```
[DeviceB-Ten-GigabitEthernet2/1/1] quit
[DeviceB] irf-port-configuration active
```

After the IRF port configuration is activated, the two devices automatically elect a master. In this example, Device A is the master. Device B automatically reboots and joins the Device A as a subordinate to form an IRF fabric. The system name of the IRF fabric is **DevicieA**.

4. Configure LACP MAD:

Create a dynamic aggregate interface and enable LACP MAD. Because LACP MAD is not run between IRF domains, you can use the default value 0.

```
<DeviceA> system-view
[DeviceA] interface bridge-aggregation 2
[DeviceA-Bridge-Aggregation2] link-aggregation mode dynamic
[DeviceA-Bridge-Aggregation2] mad enable
You need to assign a domain ID (range: 0-4294967295)
```

```

[Current domain is: 0]:
The assigned domain ID is: 0
Info: MAD LACP only enable on dynamic aggregation interface.
# Assign ports GigabitEthernet 1/0/1 and GigabitEthernet 2/0/1 to the aggregate interface.
[DeviceA] interface range gigabitethernet 1/0/1 gigabitethernet 2/0/1
[DeviceA-if-range] port link-aggregation group 2
[DeviceA-if-range] quit
[DeviceA] interface gigabitethernet 2/0/1
[DeviceA-GigabitEthernet2/0/1] port link-aggregation group 2

```

5. Configure Device C as the intermediate device:

```

# Create a dynamic aggregate interface.
<DeviceC> system-view
[DeviceC] interface bridge-aggregation 2
[DeviceC-Bridge-Aggregation2] link-aggregation mode dynamic
[DeviceC-Bridge-Aggregation2] quit
# Assign ports GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to the aggregate interface.
[DeviceC] interface gigabitethernet 1/0/1
[DeviceC-GigabitEthernet1/0/1] port link-aggregation group 2
[DeviceC-GigabitEthernet1/0/1] quit
[DeviceC] interface gigabitethernet 1/0/2
[DeviceC-GigabitEthernet1/0/2] port link-aggregation group 2

```

BFD MAD-enabled IRF configuration example

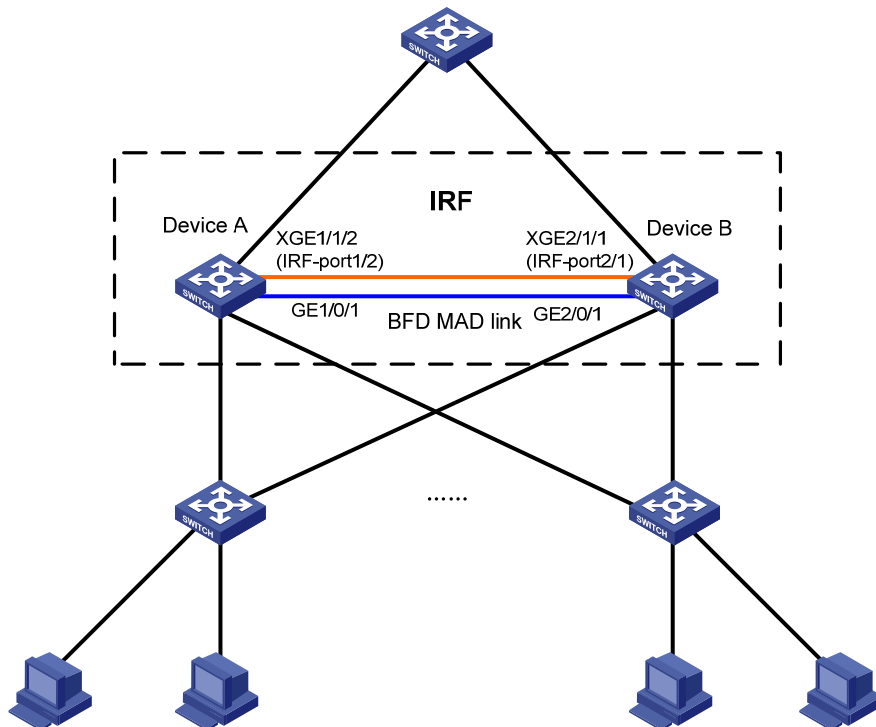
This configuration example applies only to the HP 5500 EI switches.

Network requirements

Set up an IRF fabric at the distribution layer of the network in [Figure 14](#). Configure BFD MAD in the IRF fabric and set up BFD MAD links between the member devices. Disable the spanning tree feature on the ports used for BFD MAD, because the two features conflict with each other.

Assign the highest member priority to Device A so it can be elected as the master.

Figure 14 Network diagram



Configuration procedure

This example assumes that the system names of Device A and Device B are **DeviceA** and **DeviceB** respectively before the IRF fabric is formed.

1. Assign member IDs:

Keep the default member ID of Device A unchanged.

Change the member ID of Device B to 2.

```
<DeviceB> system-view
```

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

Warning: Renumbering the switch number may result in configuration change or loss.

```
Continue? [Y/N]:y
```

```
[DeviceB]
```

2. Power off the member devices, connect IRF links as shown in Figure 14, and power on the two devices.

3. Configure IRF port bindings:

Bind Ten-GigabitEthernet 1/1/2 to IRF-port 1/2 on Device A and save the configuration.

```
<DeviceA> system-view
```

```
[DeviceA] interface ten-gigabitethernet 1/1/2
```

```
[DeviceA-Ten-GigabitEthernet1/1/2] shutdown
```

```
[DeviceA-Ten-GigabitEthernet1/1/2] quit
```

```
[DeviceA] irf-port 1/2
```

```
[DeviceA-irf-port1/2] port group interface ten-gigabitethernet 1/1/2
```

```
[DeviceA-irf-port1/2] quit
```

```
[DeviceA] interface ten-gigabitethernet 1/1/2
```

```
[DeviceA-Ten-GigabitEthernet1/1/2] undo shutdown
```

```
[DeviceA-Ten-GigabitEthernet1/1/2] save
```

Bind Ten-GigabitEthernet 2/1/1 to IRF-port 2/1 and save the configuration.

```
<DeviceB> system-view
[DeviceB] interface ten-gigabitethernet 2/1/1
[DeviceB-Ten-GigabitEthernet2/1/1] shutdown
[DeviceB-Ten-GigabitEthernet2/1/1] quit
[DeviceB] irf-port 2/1
[DeviceB-irf-port2/1] port group interface ten-gigabitethernet 2/1/1
[DeviceB-irf-port2/1] quit
[DeviceB] interface ten-gigabitethernet 2/1/1
[DeviceB-Ten-GigabitEthernet2/1/1] undo shutdown
[DeviceB-Ten-GigabitEthernet2/1/1] save
```

Activate IRF port configuration on Device A.

```
[DeviceA-Ten-GigabitEthernet1/1/2] quit
[DeviceA] irf-port-configuration active
```

Activate IRF port configuration on Device B.

```
[DeviceB-Ten-GigabitEthernet2/1/1] quit
[DeviceB] irf-port-configuration active
```

After the IRF port configuration is activated, the two devices automatically elect a master. In this example, Device A is the master. Device B automatically reboots and joins the Device A as a subordinate switch to form an IRF fabric. The system name of the IRF fabric is **DevicieA**.

4. Configure BFD MAD:

Create VLAN 3, and add port GigabitEthernet 1/0/1 on Device A (member ID 1) and port GigabitEthernet 2/0/1 on Device B (member ID 2) to VLAN 3.

```
<DeviceA> system-view
[DeviceA] vlan 3
[DeviceA-vlan3] port gigabitethernet 1/0/1 gigabitethernet 2/0/1
[DeviceA-vlan3] quit
```

Create VLAN-interface 3 and configure a MAD IP address for each member device on the interface.

```
[DeviceA] interface vlan-interface 3
[DeviceA-Vlan-interface3] mad bfd enable
[DeviceA-Vlan-interface3] mad ip address 192.168.2.1 24 member 1
[DeviceA-Vlan-interface3] mad ip address 192.168.2.2 24 member 2
[DeviceA-Vlan-interface3] quit
```

Connect the BFD MAD links as shown in [Figure 14](#).

Disable the spanning tree feature on GigabitEthernet 1/0/1 and GigabitEthernet 2/0/1.

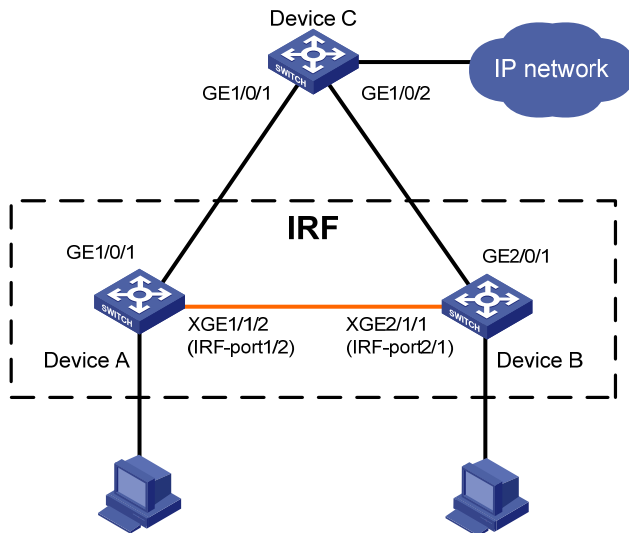
```
[DeviceA] interface range gigabitethernet 1/0/1 gigabitethernet 2/0/1
[DeviceA-if-range] undo stp enable
[DeviceA-if-range] quit
```

ARP MAD-enabled IRF configuration example

Network requirements

Set up an IRF fabric in the enterprise network in [Figure 15](#). Configure ARP MAD in the IRF fabric and use the links connected to Device C for transmitting ARP MAD packets. To prevent loops, run the spanning tree feature between Device C and the IRF fabric.

Figure 15 Network diagram



Configuration procedure

This example assumes that the system names of Device A, Device B and Device C are **DeviceA**, **DeviceB**, and **DeviceC**, respectively, before the IRF fabric is formed.

1. Assign member IDs:

Keep the default member ID of Device A unchanged.

Change the member ID of Device B to 2.

```
<DeviceB> system-view
```

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

Warning: Renumbering the switch number may result in configuration change or loss.

```
Continue? [Y/N]:y
```

```
[DeviceB]
```

2. Power off the member devices, connect IRF links as shown in Figure 15, and power on the two devices.

3. Configure IRF port bindings:

Bind Ten-GigabitEthernet 1/1/2 to IRF-port 1/2 on Device A and save the configuration.

```
<DeviceA> system-view
```

```
[DeviceA] interface ten-gigabitethernet 1/1/2
```

```
[DeviceA-Ten-GigabitEthernet1/1/2] shutdown
```

```
[DeviceA-Ten-GigabitEthernet1/1/2] quit
```

```
[DeviceA] irf-port 1/2
```

```
[DeviceA-irf-port1/2] port group interface ten-gigabitethernet 1/1/2
```

```
[DeviceA-irf-port1/2] quit
```

```
[DeviceA] interface ten-gigabitethernet 1/1/2
```

```
[DeviceA-Ten-GigabitEthernet1/1/2] undo shutdown
```

```
[DeviceA-Ten-GigabitEthernet1/1/2] save
```

Bind Ten-GigabitEthernet 2/1/1 to IRF-port 2/1 on Device B and save the configuration.

```
<DeviceB> system-view
```

```
[DeviceB] interface ten-gigabitethernet 2/1/1
```

```
[DeviceB-Ten-GigabitEthernet2/1/1] shutdown
```

```

[DeviceB-Ten-GigabitEthernet2/1/1] quit
[DeviceB] irf-port 2/1
[DeviceB-irf-port2/1] port group interface ten-gigabitethernet 2/1/1
[DeviceB-irf-port2/1] quit
[DeviceB] interface ten-gigabitethernet 2/1/1
[DeviceB-Ten-GigabitEthernet2/1/1] undo shutdown
[DeviceB-Ten-GigabitEthernet2/1/1] save
# Activate IRF port configuration on Device A.
[DeviceA-Ten-GigabitEthernet1/1/2] quit
[DeviceA] irf-port-configuration active
# Activate IRF port configuration on Device B.
[DeviceB-Ten-GigabitEthernet2/1/1] quit
[DeviceB] irf-port-configuration active

```

After the IRF port configuration is activated, the two devices automatically elect a master. In this example, Device A is the master. Device B automatically reboots and joins the Device A as a subordinate switch to form an IRF fabric. The system name of the IRF fabric is **DevicieA**.

4. Configure ARP MAD:

Enable the spanning tree feature globally on the IRF fabric to prevent loops.

```

<DeviceA> system-view
[DeviceA] stp enable

```

Connect the ARP MAD links as shown in [Figure 15](#).

Configure the IRF fabric to change its bridge MAC address as soon as the master leaves.

```

[DeviceA] undo irf mac-address persistent

```

Create VLAN 3, and add port GigabitEthernet 1/0/1 (on Device A) and port GigabitEthernet 2/0/1 (on Device B) to VLAN 3.

```

[DeviceA] vlan 3
[DeviceA-vlan3] port gigabitethernet 1/0/1 gigabitethernet 2/0/1
[DeviceA-vlan3] quit

```

Create VLAN-interface 3, assign it an IP address, and enable ARP MAD on the interface. Because ARP MAD is not run between IRF domains, you can use the default value 0.

```

[DeviceA] interface vlan-interface 3
[DeviceA-Vlan-interface3] ip address 192.168.2.1 24
[DeviceA-Vlan-interface3] mad arp enable
You need to assign a domain ID (range: 0-4294967295)
[Current domain is: 0]:
The assigned domain ID is: 0

```

5. Configure Device C:

Enable the spanning tree feature globally on Device C to prevent loops.

```

<DeviceC> system-view
[DeviceC] stp enable

```

Create VLAN 3, and add GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to VLAN 3.

```

[DeviceC] vlan 3
[DeviceC-vlan3] port gigabitethernet 1/0/1 gigabitethernet 1/0/2
[DeviceC-vlan3] quit

```

Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

<http://www.hp.com/support>

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

<http://www.hp.com/go/wwalerts>

After registering, you will receive email notification of product enhancements, new driver versions, firmware updates, and other product resources.

Related information

Documents

To find related documents, browse to the Manuals page of the HP Business Support Center website:

<http://www.hp.com/support/manuals>

- For related documentation, navigate to the Networking section, and select a networking category.
- For a complete list of acronyms and their definitions, see *HP FlexNetwork Technology Acronyms*.

Websites

- HP.com <http://www.hp.com>
- HP Networking <http://www.hp.com/go/networking>
- HP manuals <http://www.hp.com/support/manuals>
- HP download drivers and software <http://www.hp.com/support/downloads>
- HP software depot <http://www.software.hp.com>
- HP Education <http://www.hp.com/learn>

Conventions

This section describes the conventions used in this documentation set.





Command conventions

Convention	Description
Boldface	Bold text represents commands and keywords that you enter literally as shown.
<i>Italic</i>	<i>Italic</i> text represents arguments that you replace with actual values.
[]	Square brackets enclose syntax choices (keywords or arguments) that are optional.
{ x y ... }	Braces enclose a set of required syntax choices separated by vertical bars, from which you select one.
[x y ...]	Square brackets enclose a set of optional syntax choices separated by vertical bars, from which you select one or none.
{ x y ... } *	Asterisk-marked braces enclose a set of required syntax choices separated by vertical bars, from which you select at least one.
[x y ...] *	Asterisk-marked square brackets enclose optional syntax choices separated by vertical bars, from which you select one choice, multiple choices, or none.
&<1-n>	The argument or keyword and argument combination before the ampersand (&) sign can be entered 1 to n times.
#	A line that starts with a pound (#) sign is comments.








GUI conventions

Convention	Description
Boldface	Window names, button names, field names, and menu items are in bold text. For example, the New User window appears; click OK .
>	Multi-level menus are separated by angle brackets. For example, File > Create > Folder .

Symbols

Convention	Description
 WARNING	An alert that calls attention to important information that if not understood or followed can result in personal injury.
 CAUTION	An alert that calls attention to important information that if not understood or followed can result in data loss, data corruption, or damage to hardware or software.
 IMPORTANT	An alert that calls attention to essential information.
NOTE	An alert that contains additional or supplementary information.
 TIP	An alert that provides helpful information.

Network topology icons

	Represents a generic network device, such as a router, switch, or firewall.
	Represents a routing-capable device, such as a router or Layer 3 switch.
	Represents a generic switch, such as a Layer 2 or Layer 3 switch, or a router that supports Layer 2 forwarding and other Layer 2 features.
	Represents an access controller, a unified wired-WLAN module, or the switching engine on a unified wired-WLAN switch.
	Represents an access point.
	Represents a security product, such as a firewall, a UTM, or a load-balancing or security card that is installed in a device.
	Represents a security card, such as a firewall card, a load-balancing card, or a NetStream card.

Port numbering in examples

The port numbers in this document are for illustration only and might be unavailable on your device.

Index

[A](#) [B](#) [C](#) [D](#) [E](#) [F](#) [G](#) [H](#) [I](#) [M](#) [P](#) [R](#) [S](#)

A

- Accessing the IRF fabric, [18](#)
- Application scenario, [1](#)
- Assigning a member ID to each IRF member switch, [14](#)
- Assigning an IRF domain ID to the IRF fabric, [19](#)

B

- Basic concepts, [2](#)
- Binding physical ports to IRF ports, [16](#)

C

- Configuration examples, [30](#)
- Configuration synchronization mechanism, [6](#)
- Configuring a member switch description, [19](#)
- Configuring IRF bridge MAC persistence, [21](#)
- Configuring IRF link load sharing mode, [20](#)
- Configuring MAD, [23](#)
- Connecting physical IRF ports, [15](#)
- Contacting HP, [37](#)
- Conventions, [38](#)

D

- Displaying and maintaining an IRF fabric, [29](#)

E

- Enabling software auto-update for system software image synchronization, [21](#)

F

- File system naming conventions, [5](#)

G

- General restrictions and configuration guidelines, [11](#)

H

- Hardware compatibility, [1](#)

I

- Interface naming conventions, [4](#)
- IRF benefits, [1](#)
- IRF multi-active detection, [6](#)

M

- Master election, [6](#)

P

- Planning the IRF fabric setup, [14](#)

R

- Related information, [37](#)

S

- Setting the IRF link down report delay, [22](#)
- Setup and configuration task list, [12](#)
- Specifying a priority for each member switch, [15](#)