

HP A5120 EI Switch Series IRF

Configuration Guide

Abstract

This document describes the software features for the HP A Series products and guides you through the software configuration procedures. These configuration guides also provide configuration examples to help you apply software features to different network scenarios.

This documentation is intended for network planners, field technical support and servicing engineers, and network administrators working with the HP A Series products.

Part number: 5998-1789
Software version: Release 2208
Document version: 5W100-20110530



Legal and notice information

© Copyright 2011 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 configuration	1
IRF overview	1
Introduction	1
Benefits	1
Application scenario	1
IRF topologies	2
Basic concepts	3
Establishment, operation, and maintenance of an IRF fabric	4
Connecting the IRF member switches	4
Topology collection	6
Master election	6
IRF fabric management and maintenance	7
IRF multi-active detection	9
IRF fabric configuration task list	10
Configuring an IRF fabric	11
Specifying a domain ID for an IRF fabric	11
Changing the IRF member ID of a switch	12
Configuring IRF ports	13
Specifying a priority for a member switch	14
Configuring a description for a member switch	14
Configuring load sharing criteria for IRF links	14
Specifying the preservation time of bridge MAC address	15
Enabling automatic system software updating	16
Setting the IRF link down report delay	17
Configuring MAD detection	17
Configuring LACP MAD	18
Configuring ARP MAD	20
Accessing an IRF fabric	24
Accessing the master	24
Accessing a slave switch	24
Displaying and maintaining an IRF fabric	25
IRF fabric configuration examples	25
LACP MAD detection-enabled IRF configuration example	25
ARP MAD detection-enabled IRF configuration example	28
Support and other resources	31
Contacting HP	31
Subscription service	31
Related information	31
Documents	31
Websites	31
Conventions	32
Index	34

IRF configuration

NOTE:

In the HP A5120 EI Switch Series, only the following switch models can form an IRF fabric, and they must have expansion interface cards:

- HP A5120-24G EI Switch with 2 Interface Slots
 - HP A5120-24G EI TAA Switch with 2 Interface Slots
 - HP A5120-48G EI Switch with 2 Interface Slots
 - HP A5120-48G EI TAA Switch with 2 Interface Slots
 - HP A5120-24G-PoE+ EI Switch with 2 Interface Slots
 - HP A5120-24G-PoE+ EI TAA Switch with 2 Interface Slots
 - HP A5120-48G-PoE+ EI Switch with 2 Interface Slots
 - HP A5120-48G-PoE+ EI TAA Switch with 2 Interface Slots
-

IRF overview

Introduction

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

Benefits


IRF delivers the following benefits:

- Simplified topology and streamlined management. An IRF fabric appears as one node on the network. You can log in at any member switch to manage all members of the IRF fabric.
- High availability and reliability. The member switches in an IRF fabric work in 1:N redundancy. One member switch works as the master to manage and maintain the entire IRF fabric, and all other member switches process services as well as back up the master. As soon as the master fails, all other member switches elect a new master among them to prevent service interruption. In addition, you can perform link aggregation not only for IRF links but also for physical links between the IRF fabric and its upper or lower layer devices for link redundancy.
- Network scalability and resiliency. You can increase ports, bandwidth, and processing capability of an IRF fabric simply by adding member switches.

Application scenario

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

Figure 1 IRF application scenario




IRF topologies

You can create an IRF fabric in daisy chain topology, or more reliably, ring topology, as shown in [Figure 2](#).

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 affecting network services.

Figure 2 IRF connections



NOTE:

You can use at most four A5120 EI switches to form an IRF fabric.

Basic concepts

IRF member switch roles

IRF uses two member switch roles: master and slave.

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 avoid service interruption. For more information about master election, see “[Master election](#).”

IRF port

An IRF port is a logical interface for the internal connection between IRF member switches. Each IRF member switch has two IRF ports: IRF-port 1 and IRF-port 2. An IRF port is activated when you bind a physical port to it.

Physical IRF port


Physical IRF ports are physical ports bound to an IRF port. They connect IRF member switches and forward IRF protocol packets and data traffic between IRF member switches.

For more information about ports that can be used as IRF physical ports on the A5120 EI switches, see “[Connecting the IRF member switches](#).”

IRF partition

IRF partition occurs when an IRF fabric splits into two or more IRF fabrics because of IRF link failures, as shown in [Figure 3](#). The partitioned IRF fabrics operate with the same IP address and cause routing and forwarding problems on the network.


Figure 3 IRF partition



IRF merge

IRF merge occurs when two partitioned IRF fabrics re-unite 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 role that a member switch can play in an IRF fabric. A member with a higher priority is more likely to be a master. The member priority of a switch is user configurable, and defaults to 1. You can modify the priority at the command line interface (CLI).

Establishment, operation, and maintenance of an IRF fabric

IRF fabric management involves these stages: [Connecting the IRF member switches](#), [Topology collection](#), [Master election](#), [IRF fabric management and maintenance](#), and [IRF multi-active detection](#).

Connecting the IRF member switches

Prerequisites

To use the IRF feature, your switch must have at least one of the expansion interface modules in [Table 1](#). Only the ports on these modules can work as physical IRF ports.

Table 1 Physical IRF ports requirements

Interface slot	Interface modules	Cabling requirements
The expansion interface slots on the rear panel	<ul style="list-style-type: none">HP A5500/A5120-EI 1-port 10-GbE XFP Module	Use fibers or CX4/SFP+ dedicated cables.
	<ul style="list-style-type: none">HP A5500/A5120-EI 2-port 10-GbE XFP Module	
	<ul style="list-style-type: none">HP A5500/A5120-EI 2-port 10-GbE CX4 Module	Fibers cover longer reach, but CX4/SFP+ dedicated cables provide higher reliability and performance.
	<ul style="list-style-type: none">HP A5500/A5120-EI 2p 10-GbE SFP+ Module	

NOTE:


For more information about an interface module, see the card manual that came with the interface module.

Binding physical ports to IRF ports at the CLI

Bind one physical port, or for link redundancy, multiple physical ports, to an IRF port (see [“Configuring IRF ports”](#)) at the CLI. You can bind up to two physical ports to an IRF port on the A5120 EI switches.

As shown in [Figure 5](#), you must always connect the physical ports of IRF-Port1 on one switch to the physical ports of IRF-Port2 on its neighbor switch. This section uses the topology in this figure to describe the binding procedures.

Figure 5 Connect physical IRF ports



! **IMPORTANT:**

- The physical port bound to IRF port 1 must have a smaller port number than the physical port bound to IRF port 2.
 - If you bind two physical ports to one IRF port, they must be located on one expansion interface module.
-

1. The switch has only one single-port expansion interface module

The switch can work only as Switch B or Switch C at either end of a daisy chain topology. You bind the physical port to IRF port 1, if the remote port is IRF port 2, and bind the physical port to IRF port 2, if the remote port is IRF port 1.


2. The switch has one dual-port expansion module

NOTE:

This procedure uses a HP A5500/A5120-EI 2-port 10-GbE CX4 Module as an example.

Bind physical port 1 to IRF-port 1, and physical IRF port 2 to IRF-port 2, as shown in [Figure 6](#).

Figure 6 Bind physical ports to IRF ports on a dual-port interface module




3. The switch has two expansion interface modules

Always make sure that the physical port bound to IRF port 1 is on the left side of the physical port bound to IRF port 2, as you face the rear panel of the switch. This binding method ensures that the physical port bound to IRF 1 has a smaller port number than the physical port bound to IRF 2.

The physical ports bound to IRF port 1 and IRF port 2 can be located on different modules.


For example, bind physical port 1 on the left-side module to IRF port 1, and bind physical port 1 on the right-side module to IRF port 2, as shown in [Figure 7](#).

Figure 7 Bind one physical port to each IRF port



If you need to bind two physical ports to each IRF port, bind the two ports on the left-side expansion module to IRF port 1, and the two ports on the right-side expansion module to IRF port 2, as you face the rear panel of the switch, as shown in [Figure 8](#).

Figure 8 Bind two physical ports to each IRF port




NOTE:

You can also bind one physical port to one IRF port and two physical ports to the other IRF port.

Connecting the neighbor switches

Connect the physical ports of IRF port 1 on one switch to the physical ports of IRF port 2 on its neighbor switch, as shown in [Figure 9](#). Make sure the connection is the same as the binding scheme.

Figure 9 IRF fabric physical connection



Topology collection

Each member switch exchanges IRF hello packets with its directly connected neighbors to collect the topology data, including IRF port connection states, member IDs, priorities, and bridge MAC addresses.

Each member switch has a local topology database. At startup, an IRF member switch has only local topology data. When an IRF port goes up, the member switch sends its topology data out of the port periodically. The neighbor switch then updates its topology database with the received topology data.

The topology collection lasts until all members eventually get complete topology information (topology convergence), the IRF fabric enters the next stage: master election.

Master election

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

The master is elected based on the following rules in descending order:

1. The current master, even if a new member has a higher priority. (When an IRF fabric is being formed, all member switches consider themselves as the master, this rule is skipped)
2. The member with a higher priority.
3. The member with the longest system up-time. (The member switches exchange system up-time in the IRF hello packets.)
4. The member with the lowest bridge MAC address

The IRF fabric is formed on election of the master.

NOTE:

- The precision of the system up-time is 10 minutes. If two switches with the same priority reboot one after another within 10 minutes, they will have the same system up-time and the last master election rule will be followed, and the one with the lowest bridge MAC address wins.
 - During an IRF merge, the switches of the IRF fabric that fails the master election automatically reboot to join the IRF fabric that wins the election.
 - After a master election, all slave member switches initialize and reboot with the configuration on the master, and their original configuration, even if has been saved, will be lost.
-

IRF fabric management and maintenance

After the IRF fabric is established, you can access the master from any member switch to manage all the resources of the member switches.

Member ID

An IRF fabric uses member IDs to uniquely identify its members. Member IDs are also included in interface names and file system names for interface and file system identification. To guarantee the operation of the IRF fabric, you must assign each member switch a unique member ID.

Interface naming conventions

The interfaces are named in the format of *member ID/subslot number/interface serial number*, where:

- The member ID identifies the IRF member switch on which the interface resides. If the switch is standalone, the member ID defaults to 1. If the standalone switch was once an IRF member switch, it uses the same member ID as it was in the IRF fabric.
- The subslot number is the number of the slot in which the interface card resides. On the A5120 EI switches, the subslot for the fixed ports on the front panel is numbered 0, and the subslots for the two expansion slots on the rear panel are numbered 1 and 2 from left to right, as you face the rear panel.
- The interface serial number depends on the number of interfaces provided by the switch. Look at the number on the silkscreen on the interface card for the number of supported interfaces.

For 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:

```
<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:

```
<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 the *Fundamentals Configuration Guide*.

On an IRF fabric, you can also 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 following format: *Member-ID#Storage-device-name*. For example:

1. To access the **test** folder under the Flash root directory on the master switch, perform the following steps:

```
<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)
```

2. To create and access the **test** folder under the Flash root directory on member switch 3, perform the following steps:

```
<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.
```

3. To copy the **test.app** file on the master to the Flash root directory on member switch 3, perform the following steps:

```
<Master> pwd
slot3#flash:

//The current working path is the Flash root directory on slave 3.
<Master> cd flash:/
<Master> pwd
flash:

//The current working path is the Flash root directory on the master.
<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 file synchronization

IRF uses a strict configuration file synchronization mechanism to ensure that all switches in an IRF fabric can work as a single node on the network, and to ensure that after the master fails, the other switches can operate normally.

- When a slave switch starts up, it automatically gets and runs the master's configuration file. If all switches in an IRF fabric start up simultaneously, the slave switches get and run the master's startup configuration file.
- Any configuration you made on the IRF fabric is stored on the master and synchronized in real time to each member switch. When you save the current configuration to the startup configuration file of the master by using the **save** command, all slave switches execute the same saving operation.

This real-time configuration synchronization ensures that all the IRF member switches keep the same configuration file. If the master fails, all the other switches can still operate with the same configuration file.

IRF fabric topology maintenance

As soon as a member switch is down or an IRF link is down, its neighbor switches broadcast the leaving of the switch to other members. When a member switch receives the leave message, it looks up its IRF topology database to determine whether the leaving switch is the master. If yes, the member switch starts a master election and updates its IRF topology database. If the leaving switch is not a master, the member switch directly updates its IRF topology database.

NOTE:

An IRF port goes down only when all its physical IRF ports are down.

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, such as the same IP address. To avoid IP address collision and network problems, IRF uses the multi-active detection (MAD) mechanism to detect the presence of multiple identical IRF fabrics and handle collisions. MAD provides the following functions:

1. Detection

MAD detects active IRF devices with the same Layer 3 global configuration by extending the Link Aggregation Control Protocol (LACP) or the gratuitous address resolution (ARP) protocol. For more information, see "[Configuring MAD detection.](#)"

2. Collision handling

If multiple identical active IRF fabrics are detected, only the one that has the lowest master ID can operate in active state and forward traffic normally. MAD sets all other IRF fabrics in the recovery state (disabled) and shuts down all physical ports but the IRF ports and any other ports you have specified with the **mad exclude interface** command.

3. Failure recovery

An IRF link failure triggers IRF fabric partition and causes multi-active collision. In this case, repair the failed IRF link to make the collided IRF fabrics merge into one and recover the failure. If the IRF fabric in the recovery state fails before the failure is recovered, repair both the failed IRF fabric and the failed IRF link, and then the collided IRF fabrics can merge into one and the failure is recovered. If the IRF fabric in the active state fails before the failure is recovered, enable the IRF fabric in the recovery state at the CLI to make it take over the active IRF fabric and protect the services from being affected. Then, recover the MAD failure.

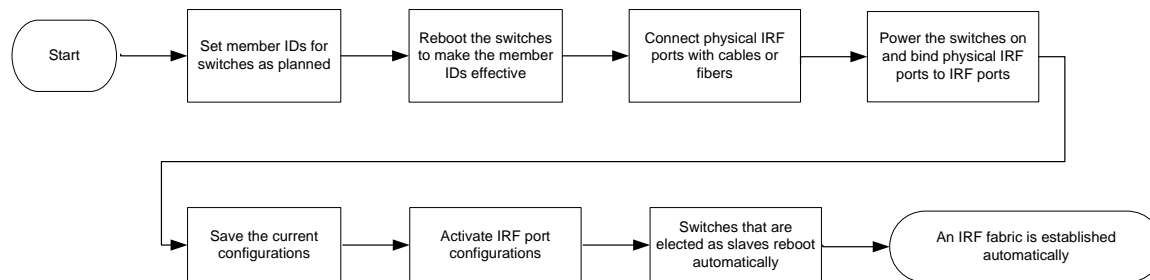
NOTE:

For more information about LACP, see the *Layer 2—LAN Switching Configuration Guide*; for information about gratuitous ARP, see the *Layer 3—IP Services Configuration Guide*.

IRF fabric configuration task list

Before configuring an IRF fabric, plan the roles and functions of all member switches. HP recommends the configuration procedure in [Figure 10](#).

Figure 10 IRF configuration flow chart



You can connect physical IRF ports with CX4/SFP+ cables or fibers after activating IRF port configurations. After the device detects that the IRF ports are connected normally, master election is started immediately, and then the elected slave switches reboot automatically.

After an IRF fabric is formed, you can configure and manage the IRF fabric by logging in to any device in the IRF.

Complete the following tasks to configure an IRF fabric:

Task	Remarks
Specifying a domain ID for an	Optional
Changing the IRF member ID of a switch	Required
Configuring IRF ports	Required
Specifying a priority for a member switch	Optional
Configuring a description for a member switch	Optional
Configuring load sharing criteria for IRF links	Optional
Specifying the preservation time of bridge MAC address	Optional
Enabling automatic system software updating	Optional
Setting the IRF link down report delay	Optional
Connect the physical IRF ports of devices and make sure that the physical IRF ports are interconnected (a ring connection is recommended).	
	Configuring LACP MAD
	Optional
	Use one of the approaches
	Configuring ARP MAD
	Optional
	Configure the MAD detection after an IRF fabric is established.
Configuring MAD detection	
	Excluding a port from the shut down action on detection of multi-active collision
	Optional
	Manually recovering an
	Optional
Accessing an	
	Accessing the master
	Required
	Accessing a slave switch
	Optional

Configuring an IRF fabric

Specifying a domain ID for an IRF fabric


Introduction to IRF domain

To differentiate IRF fabrics, each IRF fabric is assigned a domain ID.

As shown in Figure 11, Switch A and Switch B form IRF fabric 1, and Switch C and Switch D form IRF fabric 2. If there is a MAD detection link between the two IRF fabrics, they send MAD detection packets to each other through the detection link. The system statuses and operations of both IRF fabrics are affected. To solve this problem, specify different domain IDs for the two IRF fabrics.

After assigning a domain ID to an IRF fabric, the extended LACPDUs sent by the member switches carry the IRF domain information to distinguish the LACP detection packets from different IRF fabrics.

Figure 11 Network diagram for multiple domains



Assigning a domain ID to an IRF fabric

NOTE:

- If LACP MAD detection is enabled on multiple IRF fabrics, and LACP MAD detection links exist among the IRF fabrics, assign different domain IDs to the IRF fabrics.
 - If there is no LACP MAD detection link among IRF fabrics, or ARP MAD is used, you do not need to assign domain IDs to them.
-

Follow these steps to assign a domain ID to an IRF fabric:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Assign a domain ID to the IRF fabric	irf domain <i>domain-id</i>	Required if LACP MAD is adopted. By default, the domain ID of an IRF fabric is 0.

NOTE:

- You must assign a domain ID for an IRF fabric before enabling LACP MAD detection.
- Although switches with different domain IDs can form an IRF fabric, HP recommends that you assign the same domain ID to the members of the same IRF fabric; otherwise, the LACP MAD detection function cannot function properly.
- To display the domain IDs and verify your configuration, execute the **display irf** command in any view.

Changing the IRF member ID of a switch

An IRF fabric uses member IDs to uniquely identify its members. After you change the member ID of a switch, you must reboot the switch to validate the setting.

- If you do not reboot the switch, the original member ID still takes effect and all physical resources are identified by the original member ID. In the configuration file, only the IRF port numbers, configurations on IRF ports, and priority of the device change with the member ID, other configurations do not change.
- If you save the current configuration and reboot the switch, the new member ID takes effect and all physical resources are identified by the new member ID. In the configuration file, only the IRF port numbers, configurations on IRF ports, and priority of the device still take effect, other configurations (such as configuration for physical IRF ports) no longer take effect and you will need to configure them again.

You can change the IRF member ID of a switch when it is standalone or after it joins an IRF fabric. If the switch is standalone, make sure that the member ID of the switch does not conflict with the member ID of any other switch, so the change does not affect the operation of the IRF fabric. After changing the member ID, save the current configuration, power off the switch, connect the switch to its neighbor switch, power it on, and configure the IRF port to enable IRF on the switch.

Follow these steps to change the IRF member ID of a switch:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Change the IRF member ID of the switch	irf member <i>member-id</i> renumber <i>new-member-id</i>	Optional The member ID of a switch defaults to 1.

You can verify the IRF member ID of a switch with the **display irf configuration** command.

△ CAUTION:

Change member ID for the switches in an IRF fabric with caution. The change might cause configuration change and even data loss. Consider an IRF fabric that comprises three member switches of the same model with member IDs 1, 2, and 3. If you change the member ID of switch 2 to 3 and that of switch 3 to 2, then switch 2 will use the original port configurations of switch 3, and switch 3 will use those of switch 2 after they are rebooted.

Configuring IRF ports

To bring the IRF function into work, you must connect the IRF member switches, assign the connected physical ports to the appropriate IRF port on each member switch, and activate the IRF port configuration. After the IRF port configuration is activated, the IRF ports go up, a master election is held, and the switches that has failed in the election automatically reboot to join the IRF fabric as slave switches.

When binding a physical port, check that its link state is **DIS** or **DOWN** by using the **display irf topology** command.

Follow these steps to configure an IRF port:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Enter the view of the port you are binding to the IRF port	interface <i>interface-type interface-number</i>	—
Shut down the port	shutdown	Required
Return to system view	quit	—
Create the IRF port and enter IRF port view	irf-port <i>member-id/port-number</i>	—
Bind the physical port to the IRF port	port group interface <i>interface-type interface-number</i> [mode { enhanced normal }]	Required By default, no physical port is bound to any IRF port. The enhanced keyword is not supported now. You can bind up to two physical ports to an IRF port. ⚠ IMPORTANT: Make sure that the two ends of the IRF link are both using the normal mode.
Return to system view	quit	—
Enter physical IRF port view	interface <i>interface-type interface-number</i>	—
Bring up the physical port	undo shutdown	Required
Return to system view	quit	—
Save the current configuration	save	Required
Activate the IRF port configuration	irf-port-configuration active	Required

! **IMPORTANT:**

- Before you create or remove an IRF port binding, always shut down the physical IRF port. After you are finished, perform the **undo shutdown** command to bring up the port.
 - Before unplugging an interface card that contains any IRF physical port, unplug the cable of the port or shut down the port by using the **shutdown** command in IRF physical port view.
-

NOTE:

You can perform only the **shutdown**, **description** and **flow-interval** commands on the physical port bound to an IRF port. For more information about the **shutdown**, **description**, and **flow-interval** commands, see the *Layer 2—LAN Switching Command Reference*.

Specifying a priority for a member switch

The greater the priority value, the higher the priority. A member with a higher priority is more likely to be a master.

Follow these steps to specify a priority for a member switch:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Specify a priority for a member of an IRF fabric	irf member <i>member-id</i> priority <i>priority</i>	Optional The priority of a member defaults to 1

NOTE:

The priority setting takes effect immediately after configuration without the need to reboot the switch.

Configuring a description for a member switch

You can configure a description for a member switch to identify its physical location, or for any other management purpose.

Follow these steps to configure a description for a member switch:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Configure a description for a member switch	irf member <i>member-id</i> description <i>text</i>	Optional Not configured by default.

Configuring load sharing criteria for IRF links

You can bind multiple physical ports to an IRF port for link redundancy and load sharing. You can also configure the switch to distribute traffic across the physical ports of an IRF port based on one of the following criteria:

- Source IP address
- Destination IP address
- Source MAC address

- Destination MAC address
- The combination of source and destination IP addresses
- The combination of source and destination MAC addresses

You can configure global or IRF port specific load sharing criteria. The switch preferentially uses the port-specific load sharing criteria. If no port-specific load sharing criteria is available, it uses the global load sharing criteria.

Configuring global load sharing criteria

Follow these steps to configure the global IRF link load sharing criteria:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Configure the global IRF link load sharing criteria	irf-port load-sharing mode { destination-ip destination-mac source-ip source-mac } *	Required By default, the A5120 EI switches use the combination of the source and destination MAC addresses as the load sharing criteria for Layer 2 packets, and the combination of the source and destination IP addresses for Layer 3 packets.

Configuring port-specific load sharing criteria

Follow these steps to configure the port-specific load sharing criteria:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Enter IRF port view	irf-port <i>member-id/port-number</i>	—
Configure the port-specific load sharing criteria	irf-port load-sharing mode { destination-ip destination-mac source-ip source-mac } *	Required By default, the A5120 EI switches use the combination of the source and destination MAC addresses as the load sharing criteria for Layer 2 packets, and the combination of the source and destination IP addresses for Layer 3 packets.

Specifying the preservation time of bridge MAC address

An IRF fabric uses the bridge MAC address of the master switch as its bridge MAC address. The IRF fabric uses bridge MAC address to identify the IRF fabric by Layer 2 protocols such as MSTP and LACP. It uses bridge MAC address for generating an MAC address for a Layer 3 interface. As with any other node on a switched LAN, this bridge MAC address must be unique 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 temporary service interruption. Depending on your network condition, you can enable the IRF fabric to preserve or change its bridge MAC address after the master leaves. The following lists available options:

- **irf mac-address persistent timer**—Preserves the bridge MAC address for 6 minutes after the master leaves. If the master has not come back before the timer expires, the IRF fabric uses the bridge MAC address of the newly elected master as its bridge MAC address. This option avoids unnecessary switching of bridge MAC address due to a device reboot or transient link failure.
- **irf mac-address persistent always**—Keeps the bridge MAC address even after the master leaves.
- **undo irf mac-address persistent**—Uses the bridge MAC address of the newly elected master to replace the original one as soon as the master leaves.

Follow these steps to specify the preservation time of the bridge MAC address of an IRF fabric:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Enable the IRF fabric to preserve its bridge MAC address permanently even after the master leaves	irf mac-address persistent always	Optional
Enable the IRF fabric to preserve its bridge MAC address for six minutes after the master leaves	irf mac-address persistent timer	By default, the IRF fabric preserves its bridge MAC address for 6 minutes after the master leaves.
Enable the IRF fabric to change its bridge MAC address as soon as the master leaves	undo irf mac-address persistent	

CAUTION:

- Bridge MAC address change can cause transient traffic interruption.
- When deploying the ARP MAD with MSTP solution, you must enable the IRF fabric to change its bridge MAC address as soon as the master leaves.

Enabling automatic system software updating

- When you add a switch to the IRF fabric, the automatic system software updating function compares the 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 file as the system software for the next reboot, and automatically reboots with the new system software to re-join the IRF fabric.
- If this function is enabled, as soon as a switch is added into an IRF fabric, the IRF fabric compares its software version with that of the master. If the versions are not consistent, the switch automatically downloads the system software image from the master, reboots with the new system software image, and joins the IRF fabric again. If the downloaded system software image and the local system software image have duplicate file names, the local file is overwritten.

Follow these steps to enable an IRF fabric to automatically synchronize the system software of the master to the switch you are adding to the IRF fabric:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Enable automatic system software updating	irf auto-update enable	Optional Enabled by default.

△ CAUTION:

- After automatically loading the master's system software, a slave switch configures the file as the system software to be used at the next boot and reboots automatically.
 - To ensure a successful auto upgrade, check that the storage device of the slave switch has sufficient space.
-

△ CAUTION:

If you want to upgrade Boot ROM (see the release notes for detailed requirements) when you upgrade the system software, follow these steps to upgrade Boot ROM:

1. Upload the software version file to be used to the master, and then use the **bootrom upgrade** command to upgrade Boot ROM for the master.
 2. Use the **boot-loader** command with the **slot all** keywords to specify the software version file as the system software to be used at the next reboot and apply this configuration on all member switches.
 3. Reboot the all member switches in the IRF fabric to complete the software upgrade process.
-

Setting the IRF link down report delay

You can avoid link flapping causing frequent IRF splits and merges during device reconfiguration by configuring the IRF ports to delay reporting link down events. With a report delay specified, an IRF port works as follows:

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

Follow these steps to set the IRF link down report delay:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Set the IRF link down report delay	irf link-delay <i>interval</i>	Optional The default is 250 milliseconds.

△ CAUTION:

A long delay can prevent the IRF fabric from detecting IRF topology changes in time and affect service recovery.

Configuring MAD detection

You have the following MAD mechanisms for detecting multi-active collisions in different network scenarios:

- LACP MAD
- ARP MAD.

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

Configuring LACP MAD

1. LACP MAD detection mechanism

With LACP MAD, an IRF member switch sends extended LACP data units (LACPDU) with a type length value (TLV) that conveys the domain ID and active ID of the IRF fabric for detecting an IRF split. The domain ID uniquely identifies an IRF device in the network, and the active ID is identical to the member ID of the master switch in the IRF fabric.

An IRF member switch compares the domain ID and the active ID in each received extended LACPDU with its domain ID and active ID:

- If the domain IDs are different, the extended LACPDU is from a different IRF fabric, and the switch does not continue to process the extended LACPDU with the MAD mechanism.
- If the domain IDs are the same, the switch 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 operating normally.

2. Networking requirements

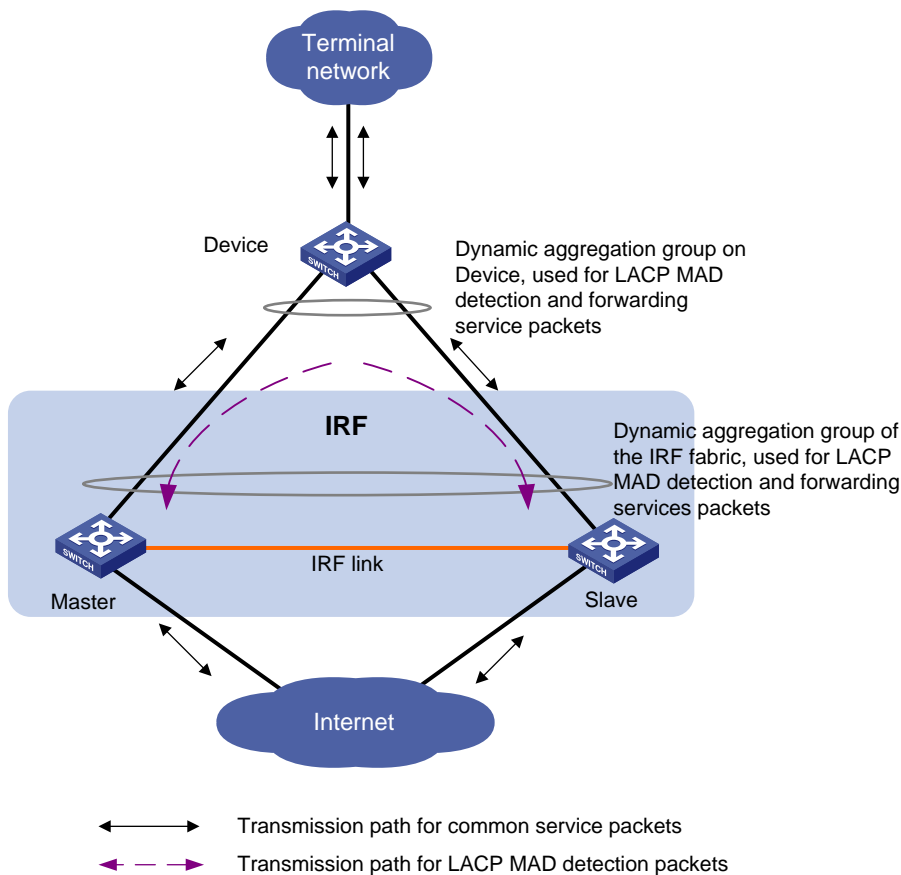
Every IRF member switch has a link with an intermediate switch, and all these links form a dynamic link aggregation group, as shown in [Figure 12](#).

The intermediate switch must be an HP switch capable of handling extended LACPDUs that carry the Active ID field. For more information about LACP and the support of the switch for extended LACPDUs, see the *Layer 2—LAN Switching Configuration Guide*.

CAUTION:

If the intermediate switch is in an IRF fabric, you must assign this virtual device a different domain ID than the LACP MAD-enabled virtual device to avoid false detection of IRF partition.

Figure 12 Network diagram for LACP MAD detection



3. Configuring LACP MAD detection

Configure LACP MAD detection by following these steps:

- Create an aggregate interface (also required on the intermediate device)
- Configure the aggregation group to work in dynamic aggregation mode (also required on the intermediate device)
- Enable LACP MAD detection on the dynamic aggregate interface (not required on the intermediate device)
- Add member ports to the aggregation group (also required on the intermediate device)

Follow these steps to configure LACP MAD detection:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Assign a domain ID to the IRF fabric	irf domain <i>domain-id</i>	Required if the intermediate switch is in an IRF fabric By default, the domain ID of an IRF fabric is 0.
Create a Layer 2 aggregate interface and enter aggregate interface view	interface bridge-aggregation <i>interface-number</i>	Required

To do...	Use the command...	Remarks
Configure the aggregation group to work in dynamic aggregation mode	link-aggregation mode dynamic	Required By default, the aggregation group works in static aggregation mode.
Enable LACP MAD detection	mad enable	Required Disabled by default. Even though this command can be configured on both static and dynamic aggregate interfaces, it takes effect only on dynamic aggregate interfaces. This is because this detection approach depends on LACP.
Return to system view	quit	—
Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
Assign the current Ethernet interface to the specified aggregation group	port link-aggregation group <i>number</i>	Required

Configuring ARP MAD

1. ARP MAD detection mechanism

ARP MAD is implemented by sending extended gratuitous ARP packets that convey MAD data. The active ID is identical to the member ID of the master of an IRF fabric and is unique to the IRF fabric.

After ARP MAD is enabled for an IRF fabric, the member switches exchange their active IDs by sending extended gratuitous packets.

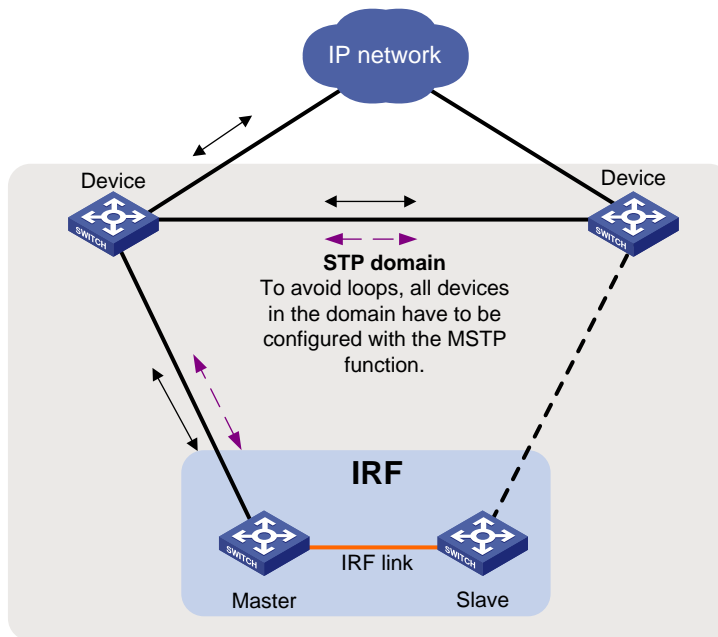
ARP MAD is applicable when an IRF fabric uses two links to connect to different upstream devices, and both the upstream devices and the IRF fabric run MSTP.

- If the IRF fabric is operating normally, the MSTP function blocks a link and the gratuitous ARP packets sent by one member switch cannot reach the other.
- When the IRF fabric splits, MSTP recalculates the topology and unblocks the link. The member switches of an IRF fabric can receive the gratuitous ARP packets from another IRF fabric. The multi-active collision is detected.

2. Network requirements

This approach can be achieved with or without intermediate devices. The commonly used networking diagram is as shown in [Figure 13](#): member switches exchange gratuitous ARP packets through two upstream devices. To avoid loops, configure the MSTP function on Device, the master and the slave switch.

Figure 13 Network diagram for ARP MAD detection



- Transmission path blocked by STP
- ====><==== Transmission path for common service packets
- - - - -><- - - - - Transmission path for gratuitous ARP packets

3. Configuring ARP MAD detection

Follow these steps to configure ARP MAD:

To do...	Use the command...	Remarks	
Enter system view	system-view	—	
Create a new VLAN dedicated for ARP MAD detection	vlan <i>vlan-id</i>	Required The default VLAN on the switch is VLAN 1.	
Return to system view	quit	—	
Enter Ethernet interface view	interface <i>interface-type</i> <i>interface-number</i>	—	
Assign the port to the VLAN dedicated for the ARP MAD detection	Access port	port access vlan <i>vlan-id</i>	Required
	Trunk port	port trunk permit vlan <i>vlan-id</i>	You can select one approach according to the port type.
	Hybrid port	port hybrid vlan <i>vlan-id</i>	ARP MAD detection has no requirement on the link type of the detection port, and you do not need to modify the current link type. By default, the port is an access port.
Return to system view	quit	—	

To do...	Use the command...	Remarks
Enter VLAN interface view	interface vlan-interface <i>interface-number</i>	—
Assign the interface an IP address	ip address <i>ip-address</i> { <i>mask</i> <i>mask-length</i> }	Required No IP address is assigned to any VLAN interface by default.
Enable ARP MAD	mad arp enable	Required By default, ARP MAD is disabled.

Excluding a port from the shut down action on detection of multi-active collision

By default all service ports of an IRF fabric except the IRF ports are shut down when the IRF fabric transits to recovery state on detection of a multi-active collision. If a port must be kept in the up state for special purposes such as telnet connection, exclude it from the shut down action.

Follow these steps to configure a port not to shut down when the IRF fabric transits to recovery state:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Configure a service port not to shut down when the IRF fabric transits to recovery state	mad exclude interface <i>interface-type</i> <i>interface-number</i>	Required When an IRF fabric transits to recovery state, all its service ports are shut down by default.

NOTE:


- Physical IRF ports are not shut down when the IRF fabric transits to recovery state.
- If a certain VLAN interface is required to go on receiving and sending packets (for example, the VLAN interface is used for remote login) after the IRF fabric transits to recovery state, you need to configure this VLAN interface and its corresponding Layer 2 Ethernet interface not to shut down when the IRF fabric transits to recovery state. However, if the VLAN interface is up in the IRF fabric in active state, IP collision will occur in your network.

Manually recovering an IRF fabric

An IRF link failure causes an IRF fabric to divide into two IRF fabrics and multi-active collision occurs. When the system detects the collision, it holds a master election between the two collided IRF fabrics. The IRF fabric whose master's member ID is smaller prevails and operates normally. The state of the other IRF fabric transits to the recovery state and temporarily cannot forward data packets. In this case, recover the IRF fabric by repairing the IRF link first (The switch tries to automatically repair the failed IRF links. If the repair fails, manually repair the failed links.)

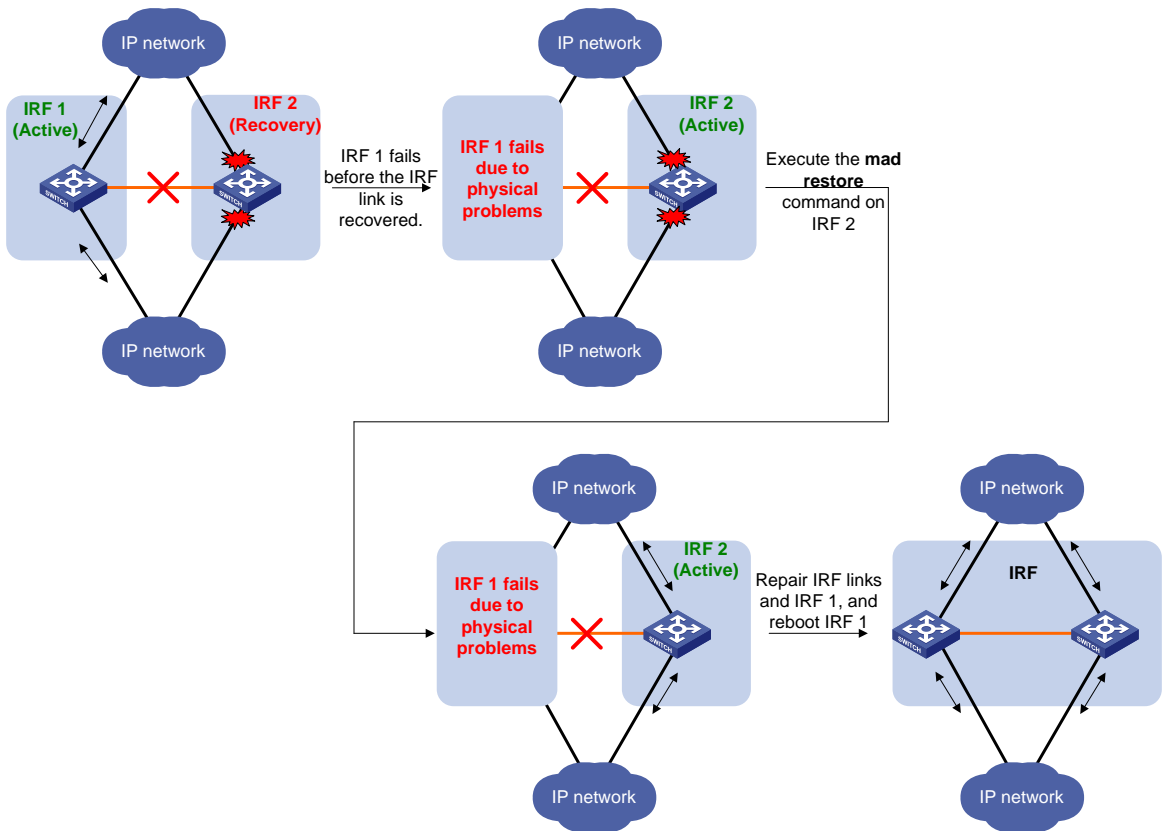
When the link is recovered, the IRF fabric in recovery state automatically reboots, and then the IRF fabrics both in active state and in recovery state automatically merge into one. Service ports that were shut down and belonged to the IRF fabric in recovery state automatically restore their original physical state, and the whole IRF fabric recovers, as shown in [Figure 14](#).

Figure 14 Recover the IRF fabric when IRF link failure occurs



If the IRF fabric in active state fails due to exceptions (for example, a member fails or link failure occurs) before the IRF link is recovered, as shown in Figure 15, enable IRF fabric 2 (in recovery state) at the CLI by executing the **mad restore** command. Then, the state of IRF fabric 2 changes from recovery to active without the need of rebooting and takes over IRF fabric 1. Repair the IRF links. When the IRF link failure is recovered, the two IRF fabrics merge. More specifically, the priorities of two masters from the two IRF fabrics are compared, and the IRF fabric whose master's priority is higher can operate normally. Members (only one in this example) of the IRF fabric whose master's priority is lower reboot themselves, and then join the other IRF fabric to complete the IRF fabric merge. After that, the original IRF fabric recovers.

Figure 15 Recover the IRF fabric when the IRF link failure occurs and the IRF fabric in active state fails



Follow these steps to manually recover an IRF fabric in recovery state:

To do...	Use the command...	Remarks
Enter system view	system-view	—
Recover an IRF fabric in recovery state	mad restore	Required

Accessing an IRF fabric

Accessing the master

Access an IRF fabric in one of the following ways:

- Local login: Log in through the console port of a member switch.
- Remote login: Configure an IP address for a Layer 3 interface of a member switch and make sure that the route is reachable. Then access the IRF fabric remotely through Telnet, web, or SNMP.

When you log in to the IRF fabric, you actually log in to the master. The master is the configuration and control center of an IRF fabric. When you configure the IRF fabric on the master, the IRF fabric synchronizes the configurations to the slave switches.

Accessing a slave switch

When you log in to an IRF fabric, you actually log in to the master. The operation interface of the access terminal displays the master console. To print the logs or debugging information of a slave switch, redirect to the specified slave switch. After that, the user access terminal displays the console of the slave switch instead of that of the master. The system enters user view of the slave switch and the command prompt is changed to <Sysname-Slave#X>, where X is the member ID of the switch, for example, <Sysname-Slave#2>. What you have input on the access terminal will be redirected to the specified slave switch for processing. You can execute the following commands on a slave switch:

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

To return to the master console, use the **quit** command. The master console is then reactivated and can output logs.

Follow these steps to log in to the specified slave switch:

To do...	Use the command...	Remarks
Enter system view	system-view	—

To do...	Use the command...	Remarks
Log in to the specified slave switch of an IRF fabric	irf switch-to <i>member-id</i>	Required By default, you actually log in to the master when you log in to the IRF fabric. Available in user view

NOTE:

An IRF fabric allows 15 concurrent VTY log-in users at most. And the maximum number of allowed console log-in users is equal to the number of IRF members.

Displaying and maintaining an IRF fabric

To do...	Use the command...	Remarks
Display related information about the IRF fabric	display irf [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display topology information about the IRF fabric	display irf topology [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display all members' configurations that take effect after switch reboots	display irf configuration [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the load sharing criteria 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/slave switchover states 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


IRF fabric configuration examples

LACP MAD detection-enabled IRF configuration example

Network requirements

The number of PCs on the enterprise network (see [Figure 16](#)) is outgrowing the number of ports available on the access switches. To accommodate business growth, the number of ports at the access layer must be increased while the present customer investments protected. In addition, the ease of management and maintenance must be ensured.

Figure 16 Network diagram for an IRF fabric that uses LACP MAD detection



Configuration considerations

- To increase the number of access ports, additional devices are needed. In this example, Device B is added.
- To address the requirements for high availability, ease of management and maintenance, use IRF2 technology to create an IRF fabric with Device A and Device B at the access layer.
- To offset the risk of IRF fabric partition, configure MAD to detect multi-active collisions. In this example, LACP MAD is adopted because the number of access devices tends to be large. In addition, for the purpose of LACP MAD, an intermediate device that supports extended LACPDU must be used.

Configuration procedure

NOTE:

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. Set member IDs

Keep the default member ID of Device A unchanged.

Set 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 two devices and connect IRF links and LACP MAD detection links according to Figure 16. Then power on the two devices.

Create IRF port 2 on Device A, and bind it to the physical IRF port Ten-GigabitEthernet 1/1/2. Then save the configuration.

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

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

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

Create IRF port 1 on Device B, and bind it to the physical IRF port Ten-GigabitEthernet 2/2/1. Then save the configuration.

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

3. Master election is held between the two devices. Master election rules are followed. Device B reboots automatically and joins the Device A as a slave switch, and the IRF fabric is formed. The system name on both devices is **DevicieA**.

4. Configure LACP MAD detection

Create a dynamic aggregation interface and enable LACP MAD detection.

```
<DeviceA> system-view
[DeviceA] interface bridge-aggregation 2
[DeviceA-Bridge-Aggregation2] link-aggregation mode dynamic
[DeviceA-Bridge-Aggregation2] mad enable
[DeviceA-Bridge-Aggregation2] quit
```

Add ports GigabitEthernet 1/0/1 and GigabitEthernet 2/0/1 to the aggregation interface and they are dedicated to the LACP MAD detection for Device A and Device B.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-aggregation group 2
[DeviceA-GigabitEthernet1/0/1] quit
[DeviceA] interface gigabitethernet 2/0/1
[DeviceA-GigabitEthernet2/0/1] port link-aggregation group 2
```

5. Configure Device C as the intermediate device

Acting as the intermediate device, Device C needs to support LACP to forward and process LACP protocol packets, and help Device A and Device B implement MAD detection. An LACP-supported switch is used here to save the cost.

Create a dynamic aggregation interface.

```

<DeviceC> system-view
[DeviceC] interface bridge-aggregation 2
[DeviceC-Bridge-Aggregation2] link-aggregation mode dynamic
[DeviceC-Bridge-Aggregation2] quit

# Add ports GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to the aggregation interface and they
are used for the LACP MAD detection.
[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


```

ARP MAD detection-enabled IRF configuration example

Network requirements

The network (see Figure 17) is outgrowing the forwarding capability of the existing core switch Device A. To accommodate to business growth, the network must be scaled up to extend its forwarding capability while the present network investments are protected. In addition, the ease of management and maintenance must be ensured.

Figure 17 Network diagram for an IRF fabric that uses ARP MAD detection



Configuration considerations

- Device A is located at the distribution layer of the network. To improve the forwarding capability at this layer, additional devices are needed. In this example, Device B is added.
- To address the requirements for high availability, ease of management and maintenance, use IRF2 technology to create an IRF fabric with Device A and Device B at the access layer. The IRF fabric is connected to Device C with dual links.
- To offset the risk of IRF fabric partition, configure MAD to detect multi-active collisions. In this example, ARP MAD is adopted because the number of members in the IRF fabric is small, and the ARP MAD packets are transmitted over dual links connected to Device C. Enable MSTP on the IRF fabric and Device to prevent loops.

Configuration procedure

NOTE:

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. Set member IDs

Keep the default member ID of Device A unchanged.

Set 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 two devices and connect IRF links and ARP MAD detection links according to [Figure 17](#). Then power on the two devices.

Create IRF port 2 on Device A, and bind it to the physical IRF port Ten-GigabitEthernet 1/1/2. Then save the configuration.

```
<DeviceA> system-view
[DeviceA] interface ten-gigabitethernet 1/1/2
[DeviceA-Ten-GigabitEthernet1/1/2] shutdown
[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
```

Create IRF port 1 on Device B, and bind it to the physical IRF port Ten-GigabitEthernet 2/0/26. Then save the configuration.

```
<DeviceB> system-view
[DeviceB] interface ten-gigabitethernet 2/0/26
[DeviceB-Ten-GigabitEthernet2/0/26] shutdown
[DeviceB] irf-port 2/1
[DeviceB-irf-port2/1] port group interface ten-gigabitethernet 2/0/26
[DeviceB-irf-port2/1] quit
[DeviceB] interface ten-gigabitethernet 2/0/26
[DeviceB-Ten-GigabitEthernet2/0/26] undo shutdown
[DeviceB-Ten-GigabitEthernet2/0/26] save
```

Activate IRF port configuration on Device A.

```
[DeviceA-Ten-GigabitEthernet1/0/25] quit
[DeviceA] irf-port-configuration active
```

Activate IRF port configuration on Device B.

```
[DeviceB-Ten-GigabitEthernet2/0/26] quit
[DeviceB] irf-port-configuration active
```


3. Master election is held between the two devices. As a result of the master election, Device B automatically reboots to join the IRF fabric as a slave switch. The system name on both devices is **DevicieA**.

4. Configure ARP MAD

Enable MSTP globally on the IRF fabric to prevent loops.

```
<DeviceA> system-view
[DeviceA] stp enable
```

Connect the ARP MAD detection links according to [Figure 17](#).

Configure that the bridge MAC address of the IRF fabric changes as soon as the master leaves.

```
[DeviceA] undo irf mac-address persistent
```

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

```
[DeviceA] interface vlan-interface 3
[DeviceA-Vlan-interface3] ip address 192.168.2.1 24
[DeviceA-Vlan-interface3] mad arp enable
```

5. Configure Device C

Enable MSTP globally on Device C to prevent loops.

```
<DeviceC> system-view
[DeviceC] stp enable
```

Create VLAN 3, and add port GigabitEthernet 1/0/1 and port GigabitEthernet 1/0/2 to VLAN 3 to forward ARP MAD packets.

```
[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 A-Series 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>

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.

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 E I L M S T

A

Accessing a slave switch, [24](#)
Accessing the master, [24](#)
Application scenario, [1](#)
ARP MAD detection-enabled IRF configuration example, [28](#)

B

Benefits, [1](#)

C

Changing the IRF member ID of a switch, [12](#)
Configuring a description for a member switch, [14](#)
Configuring ARP MAD, [20](#)
Configuring IRF ports, [13](#)
Configuring LACP MAD, [18](#)
Configuring load sharing criteria for IRF links, [14](#)
Connecting the IRF member switches, [4](#)

E

Enabling automatic system software updating, [16](#)

I

Introduction, [1](#)
IRF fabric management and maintenance, [7](#)
IRF multi-active detection, [9](#)
IRF topologies, [2](#)

L

LACP MAD detection-enabled IRF configuration example, [25](#)

M

Master election, [6](#)

S

Setting the IRF link down report delay, [17](#)
Specifying a domain ID for an IRF fabric, [11](#)
Specifying a priority for a member switch, [14](#)
Specifying the preservation time of bridge MAC address, [15](#)

T

Topology collection, [6](#)