HP FlexFabric 5700 Switch Series

IRF

Configuration Guide

Contents

Setting up an IRF fabric

Overview

HP Intelligent Resilient Framework (IRF) technology virtualizes multiple physical devices at the same layer into one virtual fabric to provide data center class availability and scalability. IRF virtualization technology offers processing power, interaction, unified management, and uninterrupted maintenance of multiple devices.

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

Figure 1 IRF application scenario
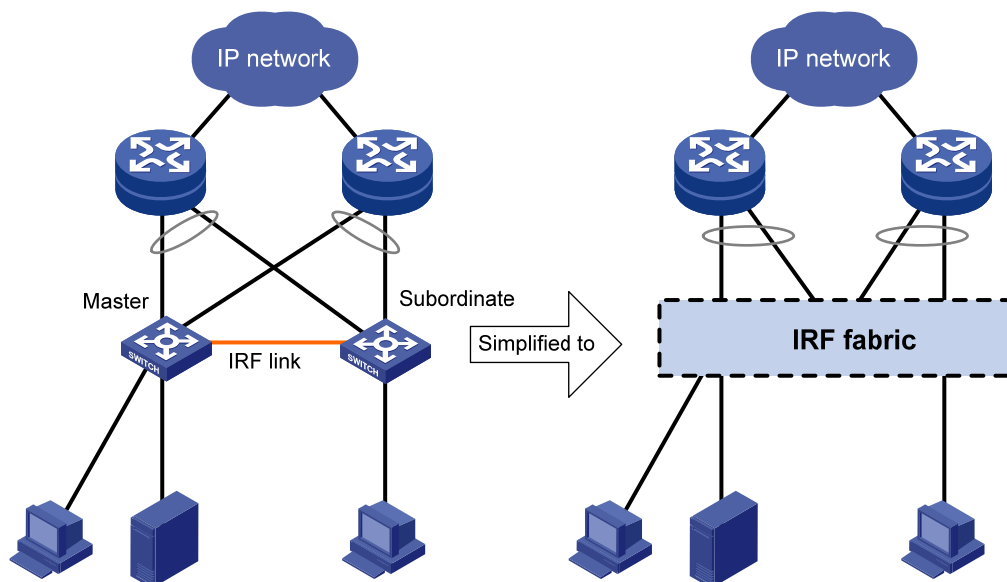
IRF provides the following benefits:

- **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 acts as the master to manage and control the entire IRF fabric. 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.

- **Multichassis 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.

Network topology

An IRF fabric can use a daisy-chain or ring topology. IRF does not support the full mesh topology. For information about connecting IRF member devices, see "Connecting IRF physical interfaces."

Basic concepts

IRF member roles

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

When devices form an IRF fabric, they elect a master to manage and control the IRF fabric, and all the other devices back up the master. When the master device fails, the other devices automatically elect a new master. 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 devices have the same IRF member ID, they cannot form an IRF fabric. If the IRF member ID of a device has been used in an IRF fabric, the device cannot join the fabric.

IRF port

An IRF port is a logical interface that connects 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 device. The two IRF ports are referred to as IRF-port 1 and IRF-port 2 in this book.

To use an IRF port, you must bind a minimum of one physical interface to it. The physical interfaces assigned to an IRF port automatically form an aggregate IRF link. An IRF port goes down only if all its IRF physical interfaces are down.

IRF physical interface

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

For more information about physical interfaces that can be used for IRF links, see "IRF physical interface requirements."

MAD

An IRF link failure causes an IRF fabric to split in two IRF fabrics operating with the same Layer 3 settings, 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.

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, IRF fabric 1 contains Device A and Device B, and IRF fabric 2 contains Device C and Device D. Both fabrics use the LACP aggregate links between them for MAD. When a member device receives an extended LACPDU for MAD, it checks the domain ID to see whether the packet is from the local IRF fabric. Then, the device can handle the packet correctly.

Figure 2 A network that contains two IRF domains
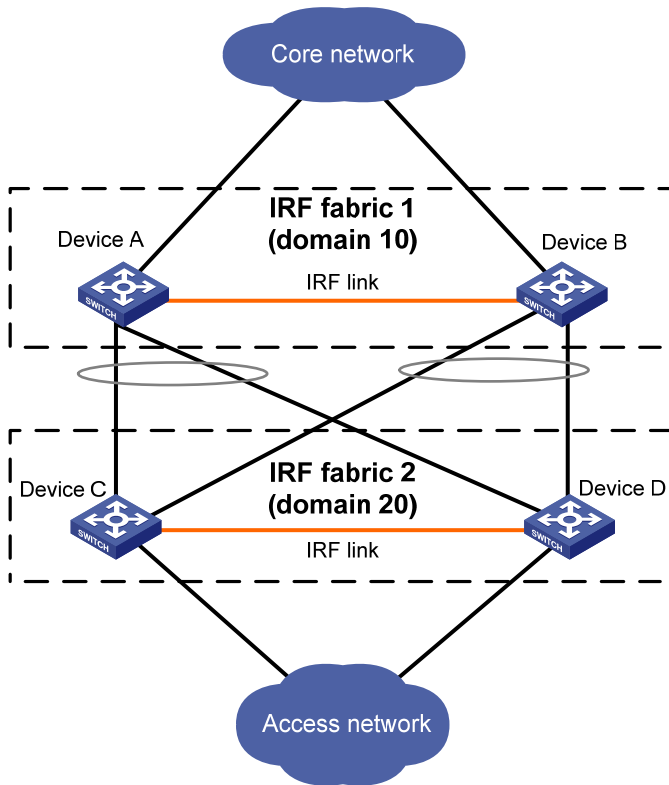
IRF split

IRF split occurs when an IRF fabric breaks up into multiple IRF fabrics because of IRF link failures, as shown in Figure 3. The split IRF fabrics operate with the same IP address. IRF split causes routing and forwarding problems on the network. To quickly detect a multi-active collision, configure a minimum of one MAD mechanism (see "Configuring MAD").

Figure 3 IRF split
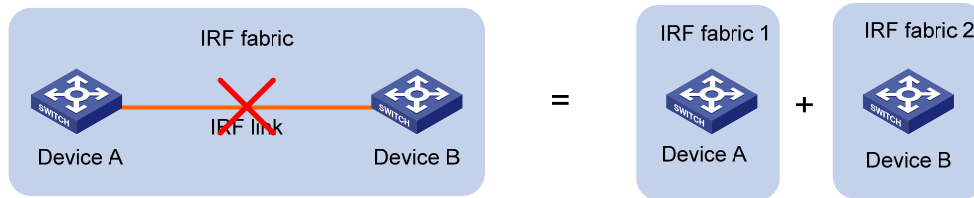
IRF merge

IRF merge occurs when two split IRF fabrics reunite or when two independent IRF fabrics are united, as shown in Figure 4.

Figure 4 IRF merge
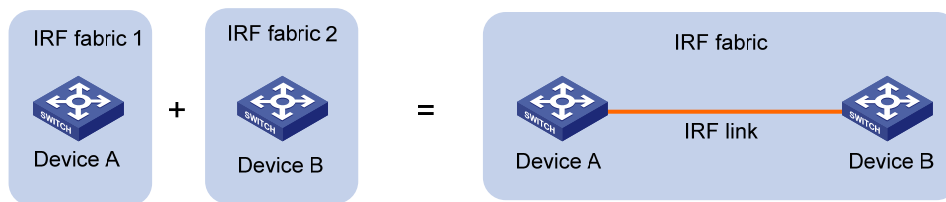
Member priority

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.

Interface naming conventions

An interface is named in the *chassis-number/slot-number/port-index* format.

- *chassis-number*—IRF member ID of the switch. This argument defaults to 1. The IRF member ID always takes effect, whether or not the device has formed an IRF fabric with other devices. If the device is alone, the device is considered to be a one-chassis IRF fabric.
- *slot-number*—Slot number of the front panel. This argument is fixed at 0.
- *port-index*—Index of the port on the device. Port index depends on the number of ports available on the device. To identify the index of a port, examine its port index mark on the chassis.

For example:

- On the single-chassis IRF fabric **Sysname**, Ten-GigabitEthernet 1/0/1 represents the first port on the device. Set its link type to trunk, as follows:
  ```
  <Sysname> system-view
  [Sysname] interface ten-gigabitethernet 1/0/1
  [Sysname-Ten-GigabitEthernet1/0/1] port link-type trunk
  ```
- On the multi-chassis IRF fabric **Master**, Ten-GigabitEthernet 3/0/1 represents the first fixed port on member device 3. Set its link type to trunk, as follows:
  ```
  <Master> system-view
  [Master] interface ten-gigabitethernet 3/0/1
  ```

```
[Master-Ten-GigabitEthernet3/0/1] port link-type trunk
```

File system naming conventions

On a single-chassis fabric, you can use its storage device name to access its file system.

On a multichassis IRF fabric, you can use the storage device name to access the file system of the master. To access the file system of any other member device, use the name in the **slot**_member-ID#storage-device-name_ format.

For example:

To access the **test** folder under the root directory of the flash memory on the master device:

```
<Master> mkdir test
Creating directory flash:/test... Done.
<Master> dir
Directory of flash:
   0 -rw-    43548660 Jan 01 2011 08:21:29   system.ipe
   1 drw-           - Jan 01 2011 00:00:30   diagfile
   2 -rw-         567 Jan 02 2011 01:41:54   dsakey
   3 -rw-         735 Jan 02 2011 01:42:03   hostkey
   4 -rw-          36 Jan 01 2011 00:07:52   ifindex.dat
   5 -rw-           0 Jan 01 2011 00:53:09   lauth.dat
   6 drw-           - Jan 01 2011 06:33:55   log
   7 drw-           - Jan 02 2000 00:00:07   logfile
   8 -rw-    23724032 Jan 01 2011 00:49:47   switch-cmw710-system.bin
   9 drw-           - Jan 01 2000 00:00:07   seclog
  10 -rw-         591 Jan 02 2011 01:42:03   serverkey
  11 -rw-        4609 Jan 01 2011 00:07:53   startup.cfg
  12 -rw-        3626 Jan 01 2011 01:51:56   startup.cfg_bak
  13 -rw-       78833 Jan 01 2011 00:07:53   startup.mdb
  14 drw-           - Jan 01 2011 00:15:48   test
  25 drw-           - Jan 01 2011 04:16:53   versionInfo


524288 KB total (365292 KB free)
```

To create and access the **test** folder under the root directory of the flash memory on member device 3:

```
<Master> mkdir slot3#flash:/test
Creating directory slot3#flash:/test... Done.
<Master> cd slot3#flash:/test
<Master> pwd
slot3#flash:/test
```

Or:

```
<Master> cd slot3#flash:/
<Master> mkdir test
Creating directory slot3#flash:/test... Done.
```

To copy the file **test.ipe** on the master to the root directory of the flash memory on member device 3:

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

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

\# Change the current working path to the root directory of the flash memory on the master device.

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

\# Copy the file to member device 3.

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

For more information about storage device naming conventions, see *Fundamentals Configuration Guide*.

Configuration synchronization

IRF uses a strict running-configuration synchronization mechanism. In an IRF fabric, all devices obtain and run the running configuration of the master. Any configuration change is automatically propagated from the master to the remaining devices. The configuration files of these devices are retained, but the files do not take effect. The devices use their own startup configuration files only after they are removed from the IRF fabric.

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

Master election

Master election occurs each time the IRF fabric topology changes in the following situations:

- The IRF fabric is established.
- The master device fails or is removed.
- The IRF fabric splits.
- Independent IRF fabrics merge.

NOTE:

Master election does not occur when two split IRF fabrics merge.

Master election selects a master in descending order:

1. Current master, even if a new member has higher priority.

    When an IRF fabric is being formed, all members consider themselves as the master. This rule is skipped.

2. Member with higher priority. If all members have the same priority, this rule is skipped.
3. Member with the longest system uptime.

    Two members are considered to start up at the same time if the difference between their startup times is equal to or less than 10 minutes. For these members, the next tiebreaker applies.

4. Member with the lowest CPU MAC address.

For the setup of a new IRF fabric, the subordinate devices must reboot to complete the setup after the master election.

For an IRF merge, devices must reboot if they are in the IRF fabric that fails the master election.

Multi-active handling procedure

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

Detection

MAD identifies 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.

For more information about the MAD mechanisms and their application scenarios, see "MAD mechanisms."

Collision handling

When MAD detects a multi-active collision, it sets all IRF fabrics except one to the Recovery state. The fabric that is not placed in Recovery state can continue to forward traffic. The Recovery-state IRF fabrics are inactive and cannot forward traffic.

LACP MAD uses the following process to handle a multi-active collision:

1. Compares the number of members in each fabric.
2. Sets all fabrics to the Recovery state except the one that has the most members.
3. Compares the member IDs of the masters if all IRF fabrics have the same number of members.
4. Sets all fabrics to the Recovery state except the one that has the lowest numbered master.
5. Shuts down all physical network ports in the Recovery-state fabrics except for the following ports:
   o IRF physical interfaces.
   o Ports you have specified with the **mad exclude interface** command.

In contrast, ARP MAD and ND MAD do not compare the number of members in fabrics. These MAD mechanisms use the following process to hand a multi-active collision:

1. Compare the member IDs of the masters in the IRF fabrics.
2. Set all fabrics to the Recovery state except the one that has the lowest numbered master.
3. Take the same action on the network ports in Recovery-state fabrics as LACP MAD.

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 active IRF fabric fails before the failure is recovered, enable the inactive IRF fabric to take over the active IRF fabric. Then, recover the MAD failure.

MAD mechanisms

IRF provides MAD mechanisms by extending LACP, ARP, and IPv6 ND.

> **(!) IMPORTANT:**
> You can configure ARP MAD and ND MAD together in an IRF fabric for prompt IRF split detection. However, do not configure any of these mechanisms together with LACP MAD, because they handle collisions differently.

Table 1 compares the MAD mechanisms and their application scenarios.

Table 1 Comparison of MAD mechanisms

MAD mechanism	Advantages	Disadvantages	Application scenario
LACP MAD	- Detection speed is fast. - Does not require MAD-dedicated physical links or Layer 3 interfaces.	Requires an intermediate device that supports extended LACP for MAD.	Link aggregation is used between the IRF fabric and its upstream or downstream device. For information about LACP, see *Layer 2—LAN Switching Configuration Guide.*
ARP MAD	- No intermediate device is required. - Intermediate device, if used, can come from any vendor. - Does not require MAD dedicated ports.	- Detection speed is slower than LACP MAD. - The spanning tree feature must be enabled.	Spanning tree-enabled non-link aggregation IPv4 network scenario. For information about ARP, see *Layer 3—IP Services Configuration Guide.*
ND MAD	- No intermediate device is required. - Intermediate device, if used, can come from any vendor. - Does not require MAD dedicated ports.	- Detection speed is slower than LACP MAD. - The spanning tree feature must be enabled.	Spanning tree-enabled non-link aggregation IPv6 network scenario.

LACP MAD

As shown in Figure 5, LACP MAD has the following requirements:

- Every IRF member must have a link with an intermediate device.
- All the links form a dynamic link aggregation group.
- The intermediate device must be a device that supports extended LACP for MAD.

The IRF member devices send extended LACPDUs that convey a domain ID and an active ID. The intermediate device transparently forwards the extended LACPDUs received from one member device to all the other member devices.

- If the domain IDs and active IDs sent by all the member devices are the same, the IRF fabric is integrated.
- If the extended LACPDUs convey the same domain ID but different active IDs, a split has occurred. LACP MAD handles this situation as described in "Collision handling."

Figure 5 LACP MAD scenario
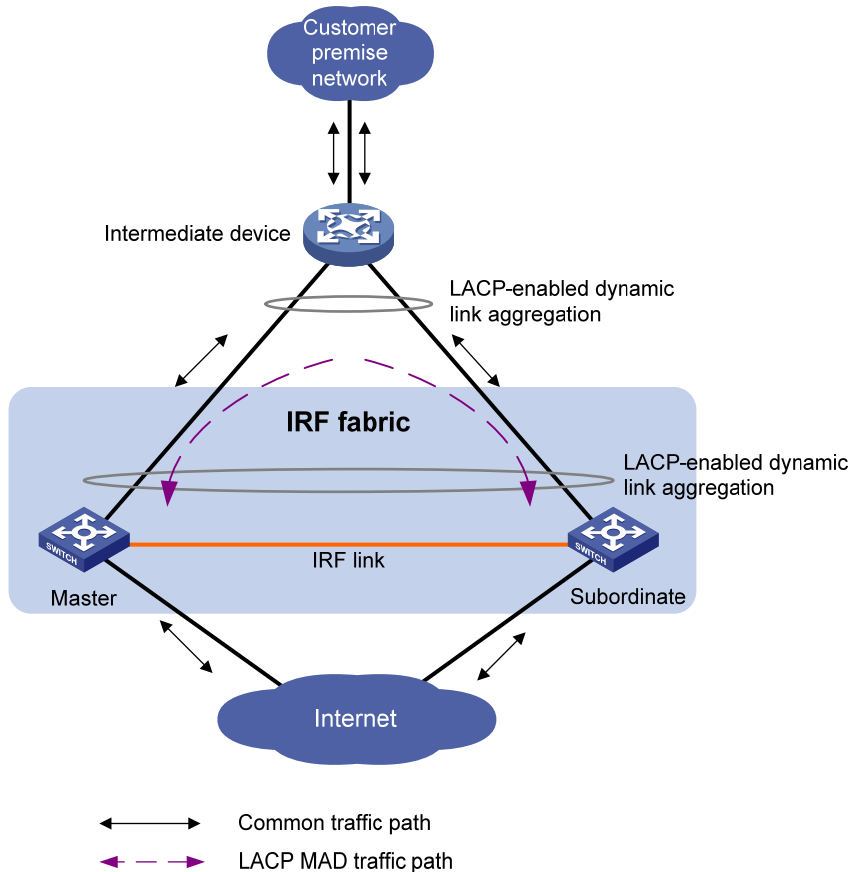
ARP MAD

ARP MAD detects multi-active collisions by using extended ARP packets that convey the IRF domain ID and the active ID.
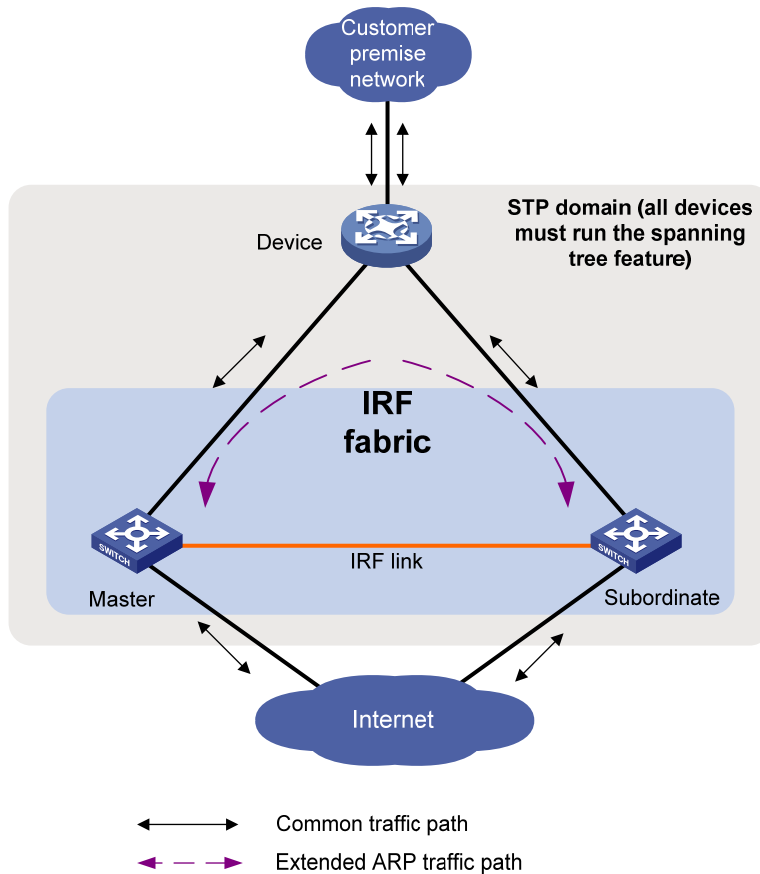
You can set up ARP MAD links between neighbor IRF member devices, or between each IRF member device and an intermediate device (see Figure 6). If an intermediate device is used, you must also run the spanning tree feature between the IRF fabric and the intermediate device.

Figure 6 ARP MAD scenario

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

- If the domain IDs are different, the extended ARP packet is from a different IRF fabric. 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.

ND MAD

ND MAD detects multi-active collisions by using NS packets to transmit the IRF domain ID and the active ID.

You can set up ND 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 protocol between the IRF fabric and the intermediate device.

Figure 7 ND MAD scenario

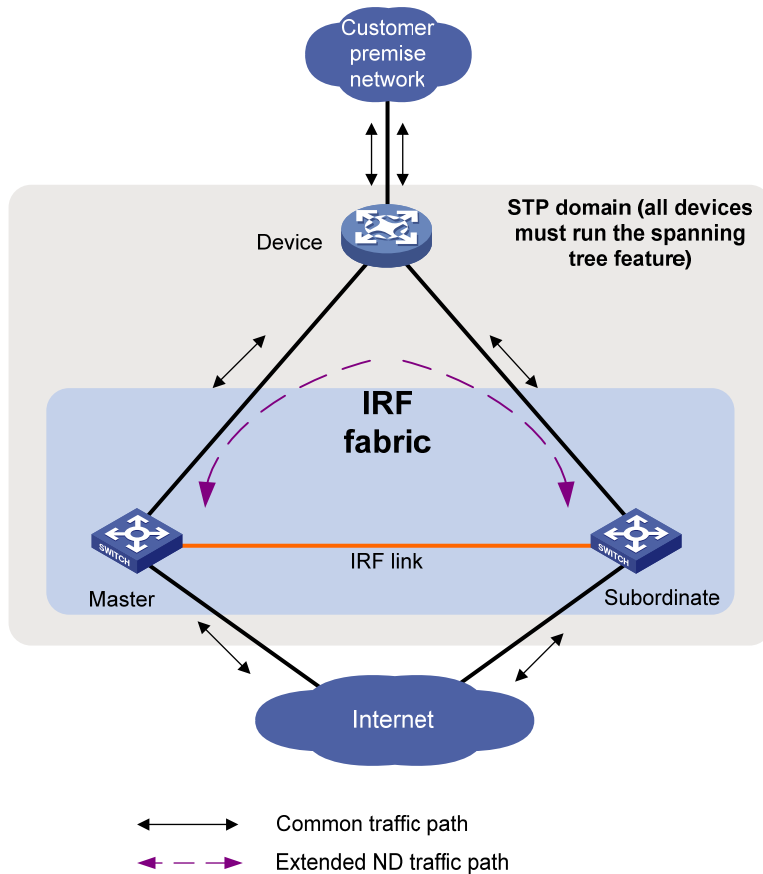Each IRF member device compares the domain ID and the active ID in incoming NS packets with its domain ID and active ID.

- If the domain IDs are different, the NS packet is from a different IRF fabric. 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.

Hardware compatibility

An HP FlexFabric 5700 switch can form an IRF fabric only with devices in the same series.

General restrictions and configuration guidelines

For a successful IRF setup, follow the restrictions and guidelines in this section and the setup procedure in "Setup and configuration task list."

Software requirements

All IRF member devices must run the same software image version. Make sure the software auto-update feature is enabled on all member devices.

IRF physical interface requirements

Use 10-GE, SFP+, or QSFP+ ports for IRF links.

Selecting transceiver modules and cables

When you select transceiver modules and cables, follow these restrictions and guidelines:

- Use Category 6A (or above) twisted-pair cables to connect 10-GE ports in a short distance.
- Use SFP+/QSFP+ DAC cables to connect SFP+/QSFP+ ports in a short distance.
- Use SFP+/QSFP+ transceiver modules and fibers to connect SFP+/QSFP+ ports in a long distance.
- The transceiver modules at the two ends of an IRF link must be the same type.

For more information about the SFP+/QSFP+ transceiver modules, see *HP Comware-Based Devices Transceiver Modules User Guide*.

NOTE:

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

10-GE port restrictions for the HP FlexFabric 5700-32XGT-8XG-2QSFP+ Switch and HP FlexFabric 5700-32XGT-8XG-2QSFP+ TAA-compliant Switch

When you use the 10-GE ports on the HP FlexFabric 5700-32XGT-8XG-2QSFP+ Switch (JG898A) and HP FlexFabric 5700-32XGT-8XG-2QSFP+ TAA-compliant Switch (JG899A) as IRF physical interfaces, follow these restrictions and guidelines:

- The 10-GE ports are grouped by port number in order, starting from 1. Each group contains four ports. If you use one port in a group for IRF links, you must also use all the other ports in the group for IRF links. However, you can bind the ports to different IRF ports.
- Before you bind a 10-GE port to an IRF port or remove it from the IRF port, you must shut down all the 10-GE ports in the same group.
- Bring up the ports after you complete the operation.

QSFP+ port restrictions for the HP FlexFabric 5700 switch

You can use a QSFP+ port as an IRF physical interface, or use the **using tengige** command to split a QSFP+ port into four 10-GE breakout interfaces.

When you use the breakout interfaces of a 40-GE port for IRF links, follow these restrictions and guidelines:

- You must use all or none of the four 10-GE breakout interfaces as IRF physical interfaces. The four breakout interfaces can be bound to different IRF ports.
- Before you bind one 10-GE breakout interface to an IRF port or remove it from the IRF port, you must shut down all the other 10-GE breakout interfaces. If any of the breakout interfaces are in up state, the bind or remove action will fail.

- Bring up the breakout interfaces after you complete the operation.

SFP+ port restrictions for the HP FlexFabric 5700-40XG-2QSFP+ Switch and HP FlexFabric 5700-40XG-2QSFP+ TAA-compliant Switch

When you use the SFP+ ports on the HP FlexFabric 5700-40XG-2QSFP+ Switch (JG896A) and HP FlexFabric 5700-40XG-2QSFP+ TAA-compliant Switch (JG897A) as IRF physical interfaces, follow these restrictions and guidelines:

- The SFP+ ports are grouped by port number in order, starting from 1. Each group contains four ports. If you use one port in a group for IRF links, you must also use all the other ports in the group for IRF links. However, you can bind the ports to different IRF ports.
- Before you bind an SFP+ port to an IRF port or remove it from the IRF port, you must shut down all the SFP+ ports in the same group.
- Bring up the ports after you complete the operation.

SFP+ port restrictions for the HP FlexFabric 5700-48G-4XG-2QSFP+ Switch and HP FlexFabric 5700-48G-4XG-2QSFP+ TAA-compliant Switch

When you use the SFP+ ports on the HP FlexFabric 5700-48G-4XG-2QSFP+ Switch (JG894A) and HP FlexFabric 5700-48G-4XG-2QSFP+ TAA-compliant Switch (JG895A) as IRF physical interfaces, follow these restrictions and guidelines:

- If you use one SFP+ port for IRF links, you must also use all the other SFP+ ports for IRF links. However, you can bind the ports to different IRF ports.
- Before you bind an SFP+ port to an IRF port or remove it from the IRF port, you must shut down all the SFP+ ports.
- Bring up the ports after you complete the operation.

Connecting IRF ports

When you connect two neighboring IRF members, follow these restrictions and guidelines:

- You must connect the physical interfaces of IRF-port 1 on one member to the physical interfaces of IRF-port 2 on the other.
- Do not connect physical interfaces of both IRF ports on one member device to the physical interfaces of both IRF ports on the other device.

Feature compatibility and configuration restrictions

Make sure the feature settings in Table 2 are the same across member devices.

Table 2 IRF and feature compatibility

| Feature	Command	Remarks
Enhanced ECMP mode	**ecmp mode enhanced**	See *Layer 3—IP Routing Configuration Guide*.
Maximum number of ECMP routes	**max-ecmp-num**	See *Layer 3—IP Routing Configuration Guide*.
System operating mode	**system-working-mode**	See *Fundamentals Configuration Guide*.

Configuration backup

HP recommends that you back up the next-startup configuration file on a device before adding the device to an IRF fabric as a subordinate.

A subordinate device's next-startup configuration file might be overwritten if the master and the subordinate use the same file name for their next-startup configuration files. You can use the backup file to restore the original configuration after removing the subordinate from the IRF fabric.

Setup and configuration task list

HP recommends the basic IRF setup procedure in Figure 8. Perform the tasks in this figure on each member device.

Figure 8 Basic IRF setup flow chart
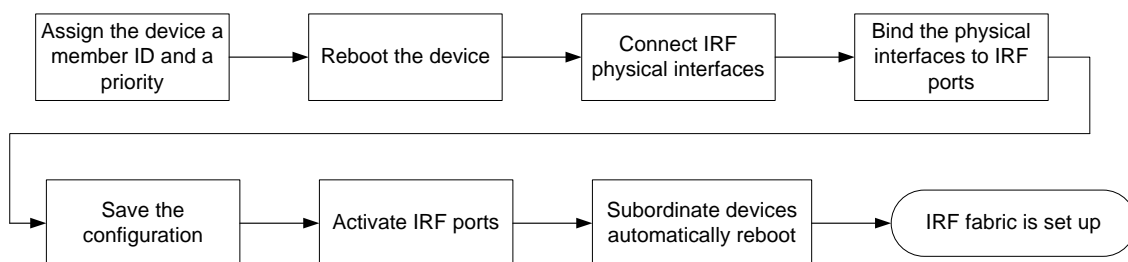
To set up an IRF fabric, perform the following tasks:

Tasks at a glance	Remarks
1. (Required.) Planning the IRF fabric setup	N/A
2. (Required.) Assigning a member ID to each IRF member device	Perform this task on each member device.
3. (Required.) Specifying a priority for each member device	Perform this task on each member device.
4. (Required.) Connecting IRF physical interfaces	N/A
5. (Required.) Binding physical interfaces to IRF ports	Perform this task on each member device. When you complete IRF port binding and activation on all IRF member devices, the IRF fabric is formed.
6. (Required.) Accessing the IRF fabric	When you log in to the IRF fabric, you are placed at the master's CLI, where you complete subsequent IRF settings and configure other features for the member devices as if they were one device.
7. (Optional.) Configuring a member device description	N/A

Tasks at a glance	Remarks

8. (Optional.) Configuring IRF link load sharing mode:
   o Configuring the global load sharing mode
   o Configuring a port-specific load sharing mode
	N/A

9. (Optional.) Configuring IRF bridge MAC persistence	N/A
10. (Optional.) Enabling software auto-update for software image synchronization	HP recommends enabling software auto-update to ensure system software image synchronization.
11. (Optional.) Setting the IRF link down report delay	N/A

12. (Required.) Configuring MAD:
   o Configuring LACP MAD
   o Configuring ARP MAD
   o Configuring ND MAD
   o Excluding a port from the shutdown action upon detection of multi-active collision
	MAD mechanisms are independent of one another. You can configure multiple MAD mechanisms for an IRF fabric.

13. (Optional.) Recovering an IRF fabric	N/A

Planning the IRF fabric setup

Consider the following items when you plan an IRF fabric:

- Hardware compatibility and restrictions.
- IRF fabric size.
- Master device.
- IRF physical interfaces.
- 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 device

△ 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 devices in an IRF fabric are the same model. If you swapped the IDs of any two members, their interface settings would also be swapped.

To create an IRF fabric, you must assign a unique IRF member ID to each member device.

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 device reboots, the settings on all member ID-related physical resources (including common physical network ports) are removed, regardless of whether you have saved the configuration.

To assign a member ID to a device:

Step		Command	Remarks
1.	Enter system view.	**system-view**	N/A
2.	Assign a member ID to a member device.	**irf member** *member-id* **renumber** *new-member-id*	The default IRF member ID is 1.
3.	(Optional.) Save the configuration.	**save**	If you have bound physical interfaces to IRF ports or assigned member priority, save the configuration before rebooting the device so these settings can continue to take effect after the reboot.
4.	Reboot the device.	**reboot** [**slot** *slot-number*] [**force**]	N/A

Specifying a priority for each member device

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 change to member priority affects the election result at the next master election. However, it does not cause an immediate master re-election.

To specify a priority for a member device:

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

Connecting IRF physical interfaces

When you connect two neighboring IRF members, connect the physical interfaces of IRF-port 1 on one member to the physical interfaces of IRF-port 2 on the other (see Figure 9).

For example, you have four chassis: A, B, C, and D. IRF-port 1 and IRF-port 2 are represented by A1 and A2 on chassis A, represented by B1 and B2 on chassis B, and so on. To connect the four chassis into a ring topology of A-B-C-D(A), the IRF link cabling scheme must be one of the following:

- A1-B2, B1-C2, C1-D2, and D1-A2.
- A2-B1, B2-C1, C2-D1, and D2-A1.

(!) IMPORTANT:

No intermediate devices are allowed between neighboring members.

Figure 9 Connecting IRF physical interfaces
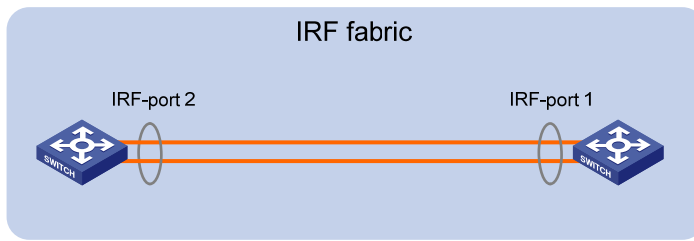
Connect the devices into a daisy-chain topology or a ring topology. A ring topology is more reliable (see Figure 10). In ring topology, the failure of one IRF link does not cause the IRF fabric to split as in daisy-chain topology. Rather, the IRF fabric changes to a daisy-chain topology without interrupting network services.

Figure 10 Daisy-chain topology vs. ring topology
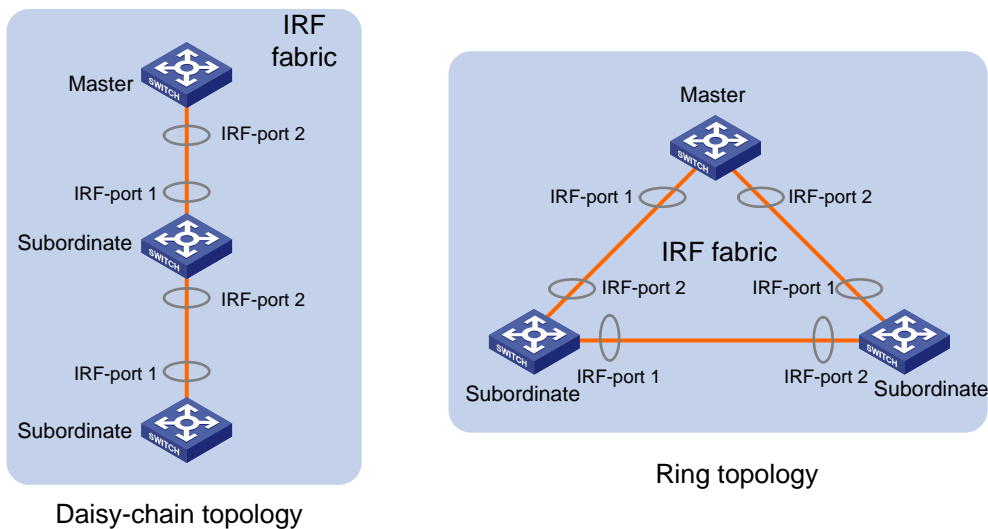
Binding physical interfaces to IRF ports

When you bind physical interfaces to IRF ports, follow these guidelines:

- Follow the restrictions in "IRF physical interface requirements."
- You must always shut down a physical interface before binding it to an IRF port or removing the binding. Start the shutdown operation on the master, and then the member device that has the fewest number of hops from the master.

On a physical interface bound to an IRF port, you can execute only the **description**, **flow-interval**, **priority-flow-control**, and **shutdown** commands. For more information about these commands, see *Layer 2—LAN Switching Command Reference*.

To bind physical interfaces to IRF ports:

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

Step	Command	Remarks
2. Enter interface view or interface range view.	• Enter interface range view: **interface range** { *interface-type interface-number* [**to** *interface-type interface-number*] } &<1-24> • Enter interface view: **interface** *interface-type interface-number*	To shut down a range of IRF physical interfaces, enter interface range view. To shut down one IRF physical interface, enter its interface view.
3. Shut down the interface or the range of interfaces.	**shutdown**	By default, all interfaces are up.
4. Return to system view.	**quit**	N/A
5. Enter IRF port view.	**irf-port** *member-id/port-number*	N/A
6. Bind each physical interface to the IRF port.	**port group interface** *interface-type interface-number*	By default, no physical interfaces are bound to any IRF port. Repeat this step to assign multiple physical interfaces to the IRF port for link redundancy. You can bind up to four physical interfaces to an IRF port.
7. Return to system view.	**quit**	N/A
8. Enter interface view or interface range view.	• Enter interface range view: **interface range** { *interface-type interface-number* [**to** *interface-type interface-number*] } &<1-24> • Enter interface view: **interface** *interface-type interface-number*	N/A
9. Bring up the interface or the range of interfaces.	**undo shutdown**	N/A
10. Return to system view.	**quit**	N/A
11. Save the configuration.	**save**	Activating IRF port settings causes IRF merge and reboot. To avoid data loss, save the running configuration to the startup configuration file before you perform the operation.
12. Activate the IRF port settings.	**irf-port-configuration active**	After this step is performed, the state of the IRF port changes to UP, the member devices elect a master automatically, and the subordinate device reboots automatically. After the IRF fabric is formed, you can add additional physical interfaces to an IRF port (in UP state) without repeating this step.

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.

The following methods are available for accessing an IRF fabric:

- **Local login**—Log in through the console port of any member device.
- **Remote login**—Log in at a Layer 3 interface on any member device by using methods including Telnet and SNMP.

When you log in to an IRF fabric, you are placed at the CLI of the master, regardless of at which member device you are logged in.

For more information, see login configuration in *Fundamentals Configuration Guide*.

Configuring a member device description

Step	Command	Remarks
1. Enter system view.	**system-view**	N/A
2. Configure a description for a member device.	**irf member** *member-id* **description** *text*	By default, no member device description is configured.

Configuring IRF link load sharing mode

An IRF port distributes traffic across its physical links.

By default, traffic is distributed automatically based on packet types, including Layer 2, IPv4, and IPv6. You can configure the IRF port to distribute traffic based on criteria including IP addresses, MAC addresses, and the combination of IP and MAC addresses. If the device does not support a criterion combination, 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 specified IRF port.

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

The IRF link load sharing mode takes effect on all types of packets, including unicast, multicast, and broadcast.

Configuration restrictions and guidelines

To distribute traffic based on TCP/UDP ports, use one of the following methods:

- Use the default setting for both global and port-specific IRF link load sharing modes.
- Set the IRF link load sharing mode and the global load sharing mode for Ethernet link aggregation as follows:

- Set the IRF link load sharing mode to distribute traffic based on source IP, destination IP, or both source and destination IP addresses. The command syntax is **irf-port load-sharing mode** { **destination-ip** | **source-ip** } *.
- Set the global load sharing mode for Ethernet link aggregation to distribute traffic based on source service port, destination service port, or both source and destination service ports. The command syntax is **link-aggregation global load-sharing mode** { **destination-port** | **source-port** } *. For more information about Ethernet link aggregation load sharing, see *Layer 2—LAN Switching Configuration Guide*.

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 global load-sharing mode** { **destination-ip**	**destination-mac**	**source-ip**	**source-mac** } *	By default, packets are distributed automatically across IRF member links based on packet types. If you execute this command multiple times, the most recent configuration takes effect.

Configuring a port-specific load sharing mode

Before you configure a port-specific load sharing mode, make sure you have bound a minimum of one physical interface 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** *member-id/port-number*	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 global IRF link load sharing mode applies. If you execute this command multiple times, the most recent configuration takes effect.

Configuring IRF bridge MAC persistence

By default, an IRF fabric uses the bridge MAC address of the master device as its bridge MAC address. Layer 2 protocols, such as LACP, use this bridge MAC address to identify the IRF fabric. On a switched LAN, the bridge MAC address must be unique.

To avoid duplicate bridge MAC addresses, an IRF fabric can change its bridge MAC address automatically after its bridge MAC owner leaves. However, the change causes temporary traffic disruption.

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

- **irf mac-address persistent timer**—Bridge MAC address of the IRF fabric remains unchanged for 6 minutes after the address owner leaves. If the owner does not return before the timer expires, the IRF fabric uses the bridge MAC address of the current master as its bridge MAC address. This option avoids unnecessary bridge MAC address changes caused by 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 address owner leaves.

> NOTE:
>
> IRF fabrics that have the same bridge MAC address cannot merge. If you use the removed bridge MAC owner as the master device in a new fabric, the bridge MAC of the new fabric is the same as the old fabric. For the two IRF fabrics to merge, you must change the bridge MAC address of the old IRF fabric.

- **undo irf mac-address persistent**—Bridge MAC address of the current master replaces the original one as soon as the owner of the original bridge MAC leaves.

When you configure IRF bridge MAC persistence, follow these guidelines:

- If ARP MAD or ND MAD is used, configure the **undo irf mac-address persistent** command to enable immediate bridge MAC address change after the address owner leaves.
- Configure the **irf mac-address persistent always** command in the following situations:
  - TRILL is configured. The persistence setting ensures that other devices in the TRILL network can maintain correct network topology after the address owner leaves.
  - The IRF fabric uses a daisy-chain topology, and it has aggregate links with upstream or downstream devices. The persistence setting prevents transmission delay or packet loss after the address owner leaves.

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.	• Retain the bridge MAC address permanently even if the owner has left the IRF fabric: **irf mac-address persistent always** • Retain the bridge MAC address for 6 minutes after the owner leaves the fabric: **irf mac-address persistent timer** • Change the bridge MAC address as soon as the owner leaves the fabric: **undo irf mac-address persistent**	By default, the IRF bridge MAC address remains unchanged for 6 minutes after the address owner leaves the fabric.

Enabling software auto-update for software image synchronization

> ! IMPORTANT:
>
> To ensure a successful software auto-update in a multi-user environment, prevent anyone from rebooting member devices during the auto-update process. To inform administrators of the auto-update status, configure the information center to output the status messages to configuration terminals (see *Network Management and Monitoring Configuration Guide*).

The software auto-update feature synchronizes the current software images of the master in an IRF fabric to all its members automatically.

To join an IRF fabric, a device must use the same software images as the master in the fabric.

When you add a device to the IRF fabric, software auto-update compares the startup software images of the device with the current software images of the IRF master. If the two sets of images are different, the device automatically performs the following operations:

1. Downloads the current software images of the master.
2. Sets the downloaded images as main startup software images.
3. Reboots with the new software images to rejoin the IRF fabric.

You must manually update the new device with the software images running on the IRF fabric in the following situations:

- Software auto-update is disabled.
- Software auto-update fails to update software. This situation might occur if the IRF fabric cannot identify the software version used on the new device.

Configuration prerequisites

Make sure the device you are adding to the IRF fabric has sufficient storage space for the new software images.

If sufficient storage space is not available, the device automatically deletes the current software images. If the reclaimed space is still insufficient, the device cannot complete the auto-update. You must reboot the device, and then access the Boot menus to delete files.

Configuration procedure

To enable software image synchronization:

Step	Command	Remarks
1. Enter system view.	**system-view**	N/A
2. Enable software auto-update.	**irf auto-update enable**	By default, software auto-update is enabled.

Setting the IRF link down report delay

To prevent frequent IRF splits and merges during link flapping, configure the IRF ports to delay reporting link down events.

An IRF port does not report a link down event to the IRF fabric immediately after its link changes from up to down. If the IRF link state is still down when the delay is reached, the port reports the change to the IRF fabric.

IRF ports do not delay link up events. They report the link up event immediately after the IRF link comes up.

When you configure the IRF link down report delay, follow these restrictions and guidelines:

- Make sure the IRF link down report delay is shorter than the heartbeat or hello timeout settings of upper-layer protocols (for example, CFD and FCoE). If the report delay is longer than the timeout setting of a protocol, unnecessary recalculations might occur.

- Set the delay to 0 seconds in the following situations:

  o The IRF fabric requires a fast master/subordinate or IRF link switchover.

  o The GR feature is used.

  o You want to shut down an IRF physical interface or reboot an IRF member device. (After you complete the operation, reconfigure the delay depending on the network condition.)

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** *interval*	The default IRF link down report delay is 4 seconds.

Configuring MAD

When you configure MAD, follow these restrictions and guidelines:

- You can configure ARP MAD and ND MAD together in an IRF fabric for prompt IRF split detection. However, do not configure any of these mechanisms together with LACP MAD, because they handle collisions differently.

- If LACP MAD, ARP MAD, or ND MAD runs between two IRF fabrics, assign each fabric a unique IRF domain ID.

- An IRF fabric has only one IRF domain ID. You can change the IRF domain ID by using the following commands: **irf domain**, **mad enable**, **mad arp enable**, or **mad nd enable**. The IRF domain IDs configured by using these commands overwrite each other.

- To prevent a port from being shut down when the IRF fabric transits to the Recovery state, use the **mad exclude interface** command. To bring up ports in a Recovery-state IRF fabric, use the **mad restore** command instead of the **undo shutdown** command. The **mad restore** command activates the Recovery-state IRF fabric.

Configuring LACP MAD

When you use LACP MAD, follow these guidelines:

- The intermediate device must be a device that supports extended LACP for MAD.
- If the intermediate device is also an IRF fabric, assign the two IRF fabrics different domain IDs for correct split detection.
- 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 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** *domain-id*	The default IRF domain ID is 0.
3. Create a Layer 2 aggregate interface and enter aggregate interface view.	**interface bridge-aggregation** *interface-number*	Perform this step also on the intermediate device.
4. Configure the aggregation group to operate in dynamic aggregation mode.	**link-aggregation mode dynamic**	By default, an aggregation group operates in static aggregation mode. Perform this step also on the intermediate device.
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 or interface range view.	• Enter interface range view: **interface range** { *interface-type interface-number* [**to** *interface-type interface-number*] } &<1-24> • Enter Ethernet interface view: **interface** *interface-type interface-number*	To assign a range of ports to the aggregation group, enter interface range view. To assign one port to the aggregation group, enter Ethernet interface view.
8. Assign the Ethernet port or the range of Ethernet ports to the specified aggregation group.	**port link-aggregation group** *number*	Multichassis link aggregation is allowed. Also perform this step on the intermediate device.

Configuring ARP MAD

When you use ARP MAD, follow these guidelines:

- Do not configure ARP MAD on VLAN-interface 1.
- Do not use the VLAN configured for ARP MAD for any other purpose.

- 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. Make sure there is only one ARP MAD link in forwarding state. For more information about the spanning tree feature and its configuration, see *Layer 2—LAN Switching Configuration Guide*.
  - Enable the IRF fabric to change its bridge MAC address as soon as the address owner leaves.
  - Create an ARP MAD VLAN and assign the ports on the ARP MAD links to the VLAN.
  - If the intermediate device is also an IRF fabric, assign the two IRF fabrics different domain IDs for correct split detection.

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** *domain-id*	The default IRF domain ID is 0.
3. Configure the IRF bridge MAC address to change as soon as the address owner leaves.	**undo irf mac-address persistent**	By default, the IRF fabric's bridge MAC address remains unchanged for 6 minutes after the address owner leaves.
4. Create a VLAN dedicated to ARP MAD.	**vlan** *vlan-id*	The default VLAN on the device is VLAN 1.
5. Return to system view.	**quit**	N/A
6. Enter Ethernet interface view.	**interface** *interface-type interface-number*	N/A
7. Assign the port to the ARP MAD VLAN.	• Assign the port to the VLAN as an access port: **port access vlan** *vlan-id* • Assign the port to the VLAN as a trunk port: **port trunk permit vlan** *vlan-id* • Assign the port to the VLAN as a hybrid port: **port hybrid vlan** *vlan-id* { **tagged** | **untagged** }	The link type of ARP MAD ports can be access, trunk, or hybrid. The default link type of a port is access.
8. Return to system view.	**quit**	N/A
9. Enter VLAN interface view.	**interface vlan-interface** *vlan-interface-id*	N/A
10. Assign the interface an IP address.	**ip address** *ip-address* { *mask* | *mask-length* }	By default, no IP address is assigned to any VLAN interface.
11. Enable ARP MAD.	**mad arp enable**	By default, ARP MAD is disabled.

Configuring ND MAD

When you use ND MAD, follow these guidelines:

- Do not configure ND MAD on VLAN-interface 1.
- Do not use the VLAN configured for ND MAD for any other purpose.
- If an intermediate device is used, you can use common data links as ND MAD links. If no intermediate device is used, set up dedicated ND 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. Make sure there is only one ND MAD link in forwarding state. For more information about the spanning tree feature and its configuration, see *Layer 2—LAN Switching Configuration Guide*.
  - Enable the IRF fabric to change its bridge MAC address as soon as the address owner leaves.
  - Create an ND MAD VLAN and assign the ports on the ND MAD links to the VLAN.
  - If the intermediate device is also an IRF fabric, assign the two IRF fabrics different domain IDs for correct split detection.

To configure ND MAD:

Step		Command	Remarks
1.	Enter system view.	**system-view**	N/A
2.	Assign a domain ID to the IRF fabric.	**irf domain** *domain-id*	The default IRF domain ID is 0.
3.	Configure the IRF bridge MAC address to change as soon as the address owner leaves.	**undo irf mac-address persistent**	By default, the IRF fabric's bridge MAC address remains unchanged for 6 minutes after the address owner leaves.
4.	Create a VLAN dedicated to ND MAD.	**vlan** *vlan-id*	The default VLAN on the device is VLAN 1.
5.	Return to system view.	**quit**	N/A
6.	Enter Ethernet interface view.	**interface** *interface-type interface-number*	N/A
7.	Assign the port to the ND MAD VLAN.	- Assign the port to the VLAN as an access port: **port access vlan** *vlan-id* - Assign the port to the VLAN as a trunk port: **port trunk permit vlan** *vlan-id* - Assign the port to the VLAN as a hybrid port: **port hybrid vlan** *vlan-id* { **tagged** | **untagged** }	The link type of ND MAD ports can be access, trunk, or hybrid. The default link type of a port is access.
8.	Return to system view.	**quit**	N/A
9.	Enter VLAN interface view.	**interface vlan-interface** *vlan-interface-id*	N/A

Step	Command	Remarks
10. Assign the interface an IP address.	**ipv6 address** { *ipv6-address/pre-length* | *ipv6 address pre-length* }	By default, no IPv6 address is assigned to any VLAN interface.
11. Enable ND MAD.	**mad nd enable**	By default, ND MAD is disabled.

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

> △ CAUTION:
>
> Do not exclude a VLAN interface and its Layer 2 ports from the shutdown action if the Layer 2 ports are distributed on multiple member devices. The exclusion introduces IP collision risks because the VLAN interface might be up on both active and inactive IRF fabrics.

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

You can exclude a network 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 device.
- 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 network port to not shut down when the IRF fabric transits to the Recovery state.	**mad exclude interface** *interface-type interface-number*	By default, all network ports on a Recovery-state IRF fabric are shut down, except for the IRF physical interfaces and console port.

Recovering an IRF fabric

When the failed IRF link between two split IRF fabrics is recovered, all member devices in the inactive IRF fabric automatically join the active IRF fabric as subordinate members. The network ports that have been shut down by MAD automatically restore their original physical state, as shown in Figure 11.

Figure 11 Recovering the IRF fabric
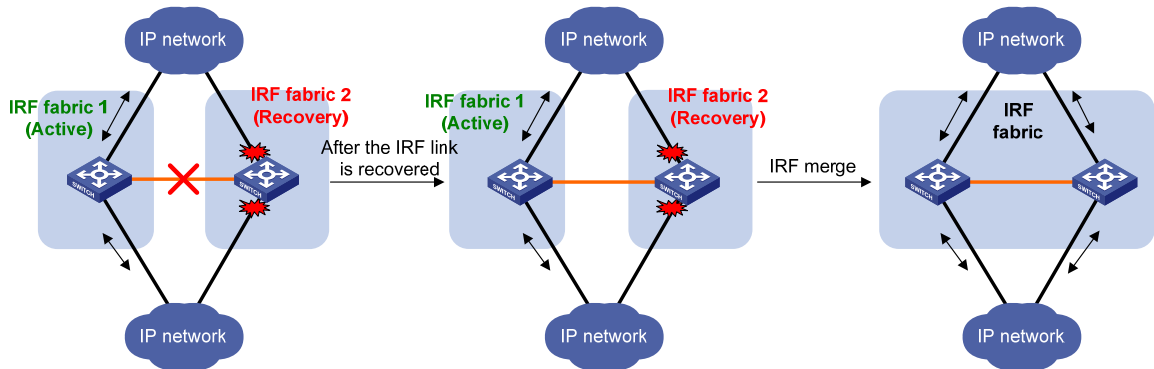
If the active IRF fabric fails before the IRF link is recovered (see Figure 12), use the **mad restore** command on the inactive IRF fabric to recover the inactive IRF fabric. The command also brings up all network ports that were shut down by MAD. After you repair the IRF link, the two parts merge into a unified IRF fabric.

Figure 12 Active IRF fabric fails before the IRF link is recovered
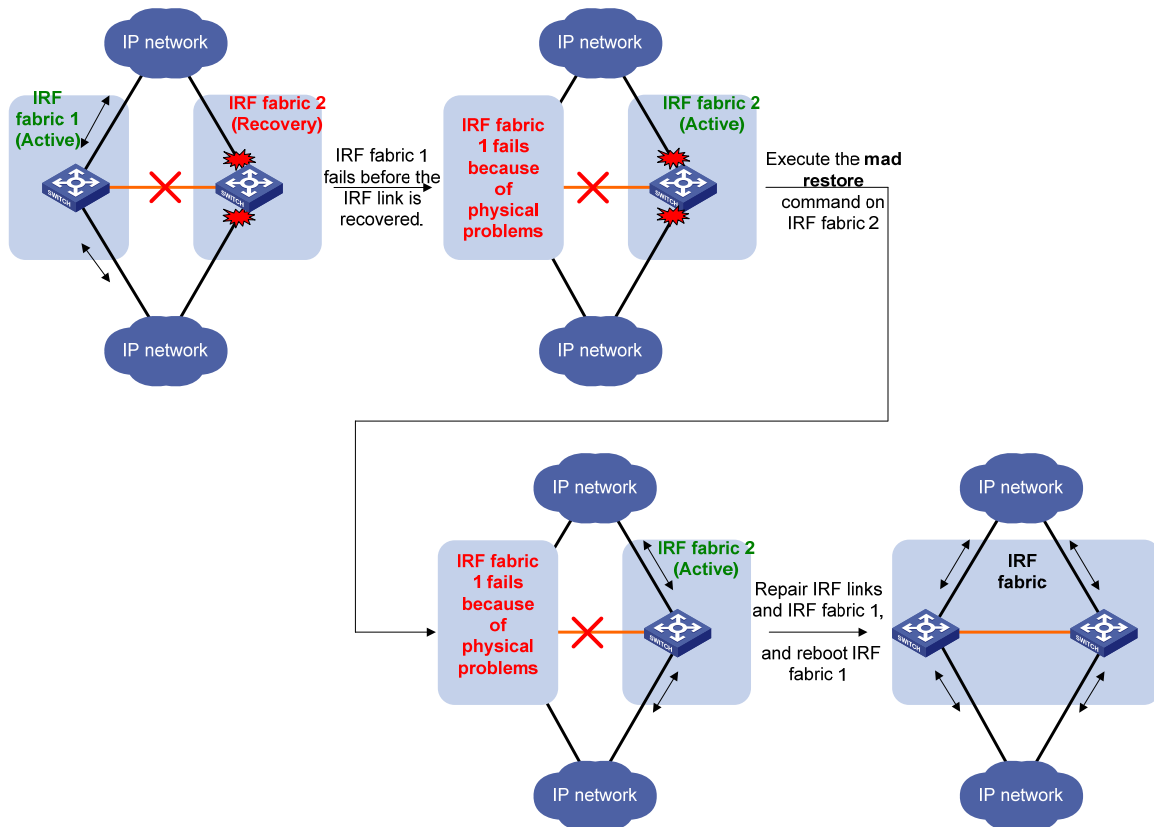
To manually recover an inactive IRF fabric:

Step	Command
1. Enter system view.	**system-view**
2. Recover the inactive IRF fabric.	**mad restore**

28

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

Displaying and maintaining an IRF fabric

Execute **display** commands in any view.

Task	Command
Display information about all IRF members.	**display irf**
Display the IRF fabric topology.	**display irf topology**
Display IRF link information.	**display irf link**
Display IRF configuration.	**display irf configuration**
Display the load sharing mode for IRF links.	**display irf-port load-sharing mode** [**irf-port** [*member-id/port-number*]]
Display MAD configuration.	**display mad** [**verbose**]

Configuration examples

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

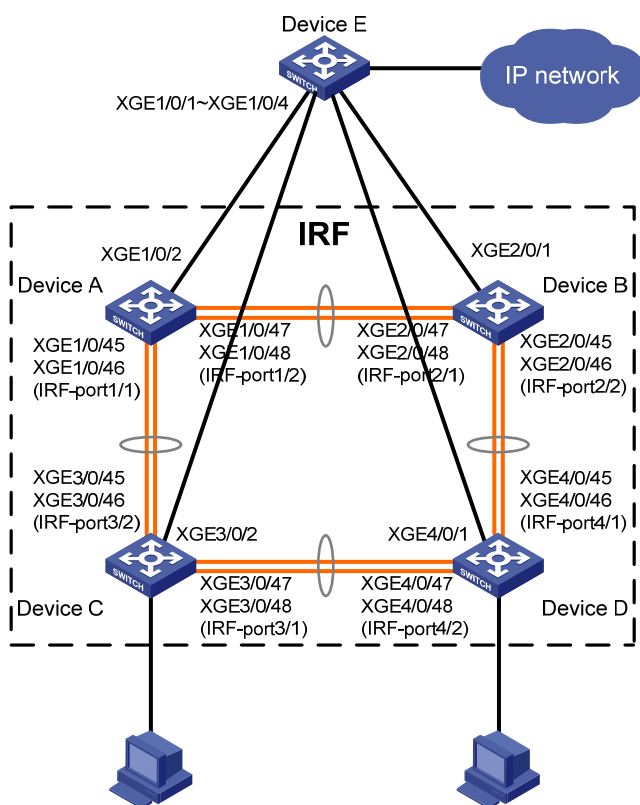## LACP MAD-enabled IRF configuration example

Network requirements

As shown in Figure 13, set up a four-chassis IRF fabric at the access layer of the enterprise network.

Configure LACP MAD on the multichassis aggregation to Device E, an HP device that supports extended LACP.

29

Figure 13 Network diagram

Configuration procedure

1. Configure Device A:

   # Shut down the physical interfaces used for IRF links. This example uses the SFP+ port group that contains Ten-GigabitEthernet 1/0/45 to Ten-GigabitEthernet 1/0/48 for IRF links.

   ```
   <Sysname> system-view
   [Sysname] interface range ten-gigabitethernet 1/0/45 to ten-gigabitethernet 1/0/48
   [Sysname-if-range] shutdown
   [Sysname-if-range] quit
   ```

   # Bind Ten-GigabitEthernet 1/0/45 and Ten-GigabitEthernet 1/0/46 to IRF-port 1/1.

   ```
   [Sysname] irf-port 1/1
   [Sysname-irf-port1/1] port group interface ten-gigabitethernet 1/0/45
   [Sysname-irf-port1/1] port group interface ten-gigabitethernet 1/0/46
   [Sysname-irf-port1/1] quit
   ```

   # Bind Ten-GigabitEthernet 1/0/47 and Ten-GigabitEthernet 1/0/48 to IRF-port 1/2.

   ```
   [Sysname] irf-port 1/2
   [Sysname-irf-port1/2] port group interface ten-gigabitethernet 1/0/47
   [Sysname-irf-port1/2] port group interface ten-gigabitethernet 1/0/48
   [Sysname-irf-port1/2] quit
   ```

   # Bring up the SFP+ ports and save the configuration.

   ```
   [Sysname] interface range ten-gigabitethernet 1/0/45 to ten-gigabitethernet 1/0/48
   [Sysname-if-range] undo shutdown
   [Sysname-if-range] quit
   ```

```
[Sysname] save
```
Activate the IRF port configuration.
```
[Sysname] irf-port-configuration active
```

2. Configure Device B:

Change the member ID of Device B to 2 and reboot the device to validate the change.
```
<Sysname> system-view
[Sysname] irf member 1 renumber 2
Renumbering the member ID may result in configuration change or loss. Continue? [Y/N]:y
[Sysname] quit
<Sysname> reboot
```

Connect Device B to Device A as shown in , and log in to Device B. This example uses the SFP+ group that contains Ten-GigabitEthernet 2/0/45 to Ten-GigabitEthernet 2/0/48 for IRF links.

Shut down the physical interfaces.
```
<Sysname> system-view
[Sysname] interface range ten-gigabitethernet 2/0/45 to ten-gigabitethernet 2/0/48
[Sysname-if-range] shutdown
[Sysname-if-range] quit
```

Bind Ten-GigabitEthernet 2/0/47 and Ten-GigabitEthernet 2/0/48 to IRF-port 2/1.
```
[Sysname] irf-port 2/1
[Sysname-irf-port2/1] port group interface ten-gigabitethernet 2/0/47
[Sysname-irf-port2/1] port group interface ten-gigabitethernet 2/0/48
[Sysname-irf-port2/1] quit
```

Bind Ten-GigabitEthernet 2/0/45 and Ten-GigabitEthernet 2/0/46 to IRF-port 2/2.
```
[Sysname] irf-port 2/2
[Sysname-irf-port2/2] port group interface ten-gigabitethernet 2/0/45
[Sysname-irf-port2/2] port group interface ten-gigabitethernet 2/0/46
[Sysname-irf-port2/2] quit
```

Bring up the SFP+ ports and save the configuration.
```
[Sysname] interface range ten-gigabitethernet 2/0/45 to ten-gigabitethernet 2/0/48
[Sysname-if-range] undo shutdown
[Sysname-if-range] quit
[Sysname] save
```

Activate the IRF port configuration.
```
[Sysname] irf-port-configuration active
```

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

3. Configure Device C:

Change the member ID of Device C to 3 and reboot the device to validate the change.
```
<Sysname> system-view
[Sysname] irf member 1 renumber 3
Renumbering the member ID may result in configuration change or loss. Continue? [Y/N]:y
[Sysname] quit
<Sysname> reboot
```

Connect Device C to Device A as shown in Figure 13, and log in to Device C. This example uses the SFP+ port group that contains Ten-GigabitEthernet 3/0/45 to Ten-GigabitEthernet 3/0/48 for IRF links.

Shut down the physical interfaces.

```
<Sysname> system-view
[Sysname] interface range ten-gigabitethernet 3/0/45 to ten-gigabitethernet 3/0/48
[Sysname-if-range] shutdown
[Sysname-if-range] quit
```

Bind Ten-GigabitEthernet 3/0/47 and Ten-GigabitEthernet 3/0/48 to IRF-port 3/1.

```
[Sysname] irf-port 3/1
[Sysname-irf-port3/1] port group interface ten-gigabitethernet 3/0/47
[Sysname-irf-port3/1] port group interface ten-gigabitethernet 3/0/48
[Sysname-irf-port3/1] quit
```

Bind Ten-GigabitEthernet 3/0/45 and Ten-GigabitEthernet 3/0/46 to IRF-port 3/2.

```
[Sysname] irf-port 3/2
[Sysname-irf-port3/2] port group interface ten-gigabitethernet 3/0/45
[Sysname-irf-port3/2] port group interface ten-gigabitethernet 3/0/46
[Sysname-irf-port3/2] quit
```

Bring up the SFP+ ports and save the configuration.

```
[Sysname] interface range ten-gigabitethernet 3/0/45 to ten-gigabitethernet 3/0/48
[Sysname-if-range] undo shutdown
[Sysname-if-range] quit
[Sysname] save
```

Activate the IRF port configuration.

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

Device C reboots to join the IRF fabric.

4. Configure Device D:

Change the member ID of Device D to 4 and reboot the device to validate the change.

```
<Sysname> system-view
[Sysname] irf member 1 renumber 4
Renumbering the member ID may result in configuration change or loss. Continue? [Y/N]:y
[Sysname] quit
<Sysname> reboot
```

Connect Device D to Device B and Device C as shown in Figure 13, and log in to Device D. This example uses the SFP+ port group that contains Ten-GigabitEthernet 4/0/45 to Ten-GigabitEthernet 4/0/48 for IRF links.

Shut down the physical interfaces.

```
<Sysname> system-view
[Sysname] interface range ten-gigabitethernet 4/0/45 to ten-gigabitethernet 4/0/48
[Sysname-if-range] shutdown
[Sysname-if-range] quit
```

Bind Ten-GigabitEthernet 4/0/45 and Ten-GigabitEthernet 4/0/46 to IRF-port 4/1.

```
[Sysname] irf-port 4/1
[Sysname-irf-port4/1] port group interface ten-gigabitethernet 4/0/45
[Sysname-irf-port4/1] port group interface ten-gigabitethernet 4/0/46
[Sysname-irf-port4/1] quit
```

Bind Ten-GigabitEthernet 4/0/47 and Ten-GigabitEthernet 4/0/48 to IRF-port 4/2.

```
[Sysname] irf-port 4/2
[Sysname-irf-port4/2] port group interface ten-gigabitethernet 4/0/47
[Sysname-irf-port4/2] port group interface ten-gigabitethernet 4/0/48
[Sysname-irf-port4/2] quit
```

Bring up the SFP+ ports and save the configuration.

```
[Sysname] interface range ten-gigabitethernet 4/0/45 to ten-gigabitethernet 4/0/48
[Sysname-if-range] undo shutdown
[Sysname-if-range] quit
[Sysname] save
```

Activate the IRF port configuration.

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

Device D reboots to join the IRF fabric. A four-chassis IRF fabric is formed.

5. Configure LACP MAD on the IRF fabric:

Set the domain ID of the IRF fabric to 1.

```
<Sysname> system-view
[Sysname] irf domain 1
```

Create a dynamic aggregate interface and enable LACP MAD.

```
[Sysname] interface bridge-aggregation 2
[Sysname-Bridge-Aggregation2] link-aggregation mode dynamic
[Sysname-Bridge-Aggregation2] mad enable
You need to assign a domain ID (range: 0-4294967295)
[Current domain is: 1]:
 The assigned  domain ID is: 1
 Info: MAD LACP only enable on dynamic aggregation interface.
[Sysname-Bridge-Aggregation2] quit
```

Assign Ten-GigabitEthernet 1/0/2, Ten-GigabitEthernet 2/0/1, Ten-GigabitEthernet 3/0/2, and Ten-GigabitEthernet 4/0/1 to the aggregate interface.

```
[Sysname] interface range ten-gigabitethernet 1/0/2 ten-gigabitethernet 2/0/1
ten-gigabitethernet 3/0/2 ten-gigabitethernet 4/0/1
[Sysname-if-range] port link-aggregation group 2
[Sysname-if-range] quit
```

6. Configure Device E as the intermediate device:

⚠ **CAUTION:**

If the intermediate device is also an IRF fabric, assign the two IRF fabrics different domain IDs for correct split detection. False detection causes IRF split.

Create a dynamic aggregate interface.

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

Assign Ten-GigabitEthernet 1/0/1, Ten-GigabitEthernet 1/0/2, Ten-GigabitEthernet 1/0/3, and Ten-GigabitEthernet 1/0/4 to the aggregate interface.

```
[Sysname] interface range ten-gigabitethernet 1/0/1 to ten-gigabitethernet 1/0/4
[Sysname-if-range] port link-aggregation group 2
```
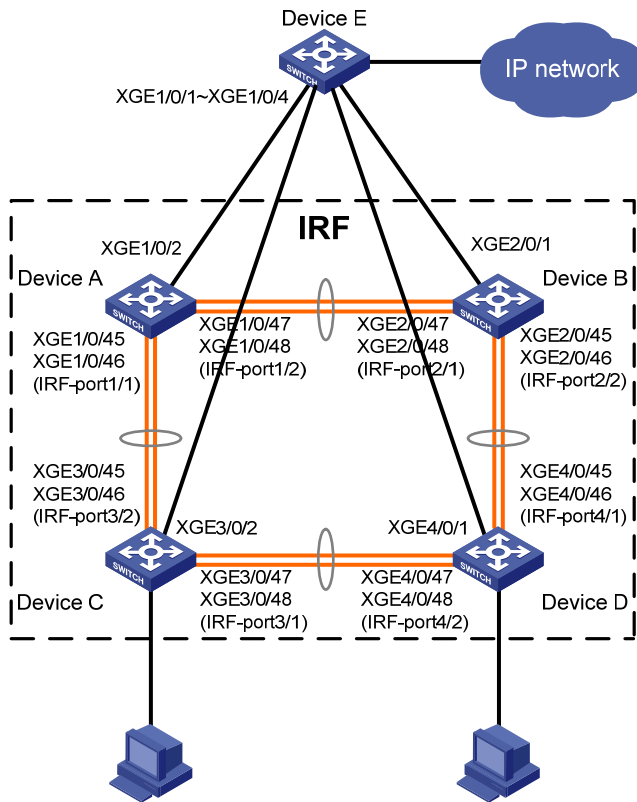
```
[Sysname-if-range] quit
```

ARP MAD-enabled IRF configuration example

Network requirements

As shown in Figure 14, set up a four-chassis IRF fabric in the enterprise network.

- Configure ARP MAD in the IRF fabric and use the links connected to Device E for transmitting ARP MAD packets.
- To prevent loops, run the spanning tree feature between Device E and the IRF fabric.

Figure 14 Network diagram

Configuration procedure

1. Configure Device A:

   # Shut down the physical interfaces used for IRF links. This example uses the SFP+ port group that contains Ten-GigabitEthernet 1/0/45 to Ten-GigabitEthernet 1/0/48 for IRF links.

   ```
   <Sysname> system-view
   [Sysname] interface range ten-gigabitethernet 1/0/45 to ten-gigabitethernet 1/0/48
   [Sysname-if-range] shutdown
   [Sysname-if-range] quit
   ```

   # Bind Ten-GigabitEthernet 1/0/45 and Ten-GigabitEthernet 1/0/46 to IRF-port 1/1.

   ```
   [Sysname] irf-port 1/1
   [Sysname-irf-port1/1] port group interface ten-gigabitethernet 1/0/45
   [Sysname-irf-port1/1] port group interface ten-gigabitethernet 1/0/46
   [Sysname-irf-port1/1] quit
   ```

34

Bind Ten-GigabitEthernet 1/0/47 and Ten-GigabitEthernet 1/0/48 to IRF-port 1/2.

```
[Sysname] irf-port 1/2
[Sysname-irf-port1/2] port group interface ten-gigabitethernet 1/0/47
[Sysname-irf-port1/2] port group interface ten-gigabitethernet 1/0/48
[Sysname-irf-port1/2] quit
```

Bring up the SFP+ ports and save the configuration.

```
[Sysname] interface range ten-gigabitethernet 1/0/45 to ten-gigabitethernet 1/0/48
[Sysname-if-range] undo shutdown
[Sysname-if-range] quit
[Sysname] save
```

Activate the IRF port configuration.

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

2. Configure Device B:

Change the member ID of Device B to 2 and reboot the device to validate the change.

```
<Sysname> system-view
[Sysname] irf member 1 renumber 2
Renumbering the member ID may result in configuration change or loss. Continue? [Y/N]:y
[Sysname] quit
<Sysname> reboot
```

Connect Device B to Device A as shown in Figure 14, and log in to Device B. This example uses the SFP+ group that contains Ten-GigabitEthernet 2/0/45 to Ten-GigabitEthernet 2/0/48 for IRF links.

Shut down the physical interfaces.

```
<Sysname> system-view
[Sysname] interface range ten-gigabitethernet 2/0/45 to ten-gigabitethernet 2/0/48
[Sysname-if-range] shutdown
[Sysname-if-range] quit
```

Bind Ten-GigabitEthernet 2/0/47 and Ten-GigabitEthernet 2/0/48 to IRF-port 2/1.

```
[Sysname] irf-port 2/1
[Sysname-irf-port2/1] port group interface ten-gigabitethernet 2/0/47
[Sysname-irf-port2/1] port group interface ten-gigabitethernet 2/0/48
[Sysname-irf-port2/1] quit
```

Bind Ten-GigabitEthernet 2/0/45 and Ten-GigabitEthernet 2/0/46 to IRF-port 2/2.

```
[Sysname] irf-port 2/2
[Sysname-irf-port2/2] port group interface ten-gigabitethernet 2/0/45
[Sysname-irf-port2/2] port group interface ten-gigabitethernet 2/0/46
[Sysname-irf-port2/2] quit
```

Bring up the SFP+ ports and save the configuration.

```
[Sysname] interface range ten-gigabitethernet 2/0/45 to ten-gigabitethernet 2/0/48
[Sysname-if-range] undo shutdown
[Sysname-if-range] quit
[Sysname] save
```

Activate the IRF port configuration.

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

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

3. Configure Device C:

\# Change the member ID of Device C to 3 and reboot the device to validate the change.

```
<Sysname> system-view
[Sysname] irf member 1 renumber 3
Renumbering the member ID may result in configuration change or loss. Continue? [Y/N]:y
[Sysname] quit
<Sysname> reboot
```

\# Connect Device C to Device A as shown in Figure 14, and log in to Device C. This example uses the SFP+ port group that contains Ten-GigabitEthernet 3/0/45 to Ten-GigabitEthernet 3/0/48 for IRF links.

\# Shut down the physical interfaces.

```
<Sysname> system-view
[Sysname] interface range ten-gigabitethernet 3/0/45 to ten-gigabitethernet 3/0/48
[Sysname-if-range] shutdown
[Sysname-if-range] quit
```

\# Bind Ten-GigabitEthernet 3/0/47 and Ten-GigabitEthernet 3/0/48 to IRF-port 3/1.

```
[Sysname] irf-port 3/1
[Sysname-irf-port3/1] port group interface ten-gigabitethernet 3/0/47
[Sysname-irf-port3/1] port group interface ten-gigabitethernet 3/0/48
[Sysname-irf-port3/1] quit
```

\# Bind Ten-GigabitEthernet 3/0/45 and Ten-GigabitEthernet 3/0/46 to IRF-port 3/2.

```
[Sysname] irf-port 3/2
[Sysname-irf-port3/2] port group interface ten-gigabitethernet 3/0/45
[Sysname-irf-port3/2] port group interface ten-gigabitethernet 3/0/46
[Sysname-irf-port3/2] quit
```

\# Bring up the SFP+ ports and save the configuration.

```
[Sysname] interface range ten-gigabitethernet 3/0/45 to ten-gigabitethernet 3/0/48
[Sysname-if-range] undo shutdown
[Sysname-if-range] quit
[Sysname] save
```

\# Activate the IRF port configuration.

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

Device C reboots to join the IRF fabric.

4. Configure Device D:

\# Change the member ID of Device D to 4 and reboot the device to validate the change.

```
<Sysname> system-view
[Sysname] irf member 1 renumber 4
Renumbering the member ID may result in configuration change or loss. Continue? [Y/N]:y
[Sysname] quit
<Sysname> reboot
```

\# Connect Device D to Device B and Device C as shown in Figure 14, and log in to Device D. This example uses the SFP+ port group that contains Ten-GigabitEthernet 4/0/45 to Ten-GigabitEthernet 4/0/48 for IRF links.

\# Shut down the physical interfaces.

```
<Sysname> system-view
[Sysname] interface range ten-gigabitethernet 4/0/45 to ten-gigabitethernet 4/0/48
```

```
[Sysname-if-range] shutdown
[Sysname-if-range] quit
```

Bind Ten-GigabitEthernet 4/0/45 and Ten-GigabitEthernet 4/0/46 to IRF-port 4/1.

```
[Sysname] irf-port 4/1
[Sysname-irf-port4/1] port group interface ten-gigabitethernet 4/0/45
[Sysname-irf-port4/1] port group interface ten-gigabitethernet 4/0/46
[Sysname-irf-port4/1] quit
```

Bind Ten-GigabitEthernet 4/0/47 and Ten-GigabitEthernet 4/0/48 to IRF-port 4/2.

```
[Sysname] irf-port 4/2
[Sysname-irf-port4/2] port group interface ten-gigabitethernet 4/0/47
[Sysname-irf-port4/2] port group interface ten-gigabitethernet 4/0/48
[Sysname-irf-port4/2] quit
```

Bring up the SFP+ ports and save the configuration.

```
[Sysname] interface range ten-gigabitethernet 4/0/45 to ten-gigabitethernet 4/0/48
[Sysname-if-range] undo shutdown
[Sysname-if-range] quit
[Sysname] save
```

Activate the IRF port configuration.

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

Device D reboots to join the IRF fabric. A four-chassis IRF fabric is formed.

5. Configure ARP MAD on the IRF fabric:

Enable the spanning tree feature globally. Map the ARP MAD VLAN to MSTI 1 in the MST region.

```
<Sysname> system-view
[Sysname] stp global enable
[Sysname] stp region-configuration
[Sysname-mst-region] region-name arpmad
[Sysname-mst-region] instance 1 vlan 3
[Sysname-mst-region] active region-configuration
[Sysname-mst-region] quit
```

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

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

Set the domain ID of the IRF fabric to 1.

```
[Sysname] irf domain 1
```

Create VLAN 3, and assign Ten-GigabitEthernet 1/0/2, Ten-GigabitEthernet 2/0/1, Ten-GigabitEthernet 3/0/2, and Ten-GigabitEthernet 4/0/1 to VLAN 3.

```
[Sysname] vlan 3
[Sysname-vlan3] port ten-gigabitethernet 1/0/2 ten-gigabitethernet 2/0/1
ten-gigabitethernet 3/0/2 ten-gigabitethernet 4/0/1
[Sysname-vlan3] quit
```

Create VLAN-interface 3, assign it an IP address, and enable ARP MAD on the interface.

```
[Sysname] interface vlan-interface 3
[Sysname-Vlan-interface3] ip address 192.168.2.1 24
[Sysname-Vlan-interface3] mad arp enable
You need to assign a domain ID (range: 0-4294967295)
 [Current domain is: 1]:
```

```
    The assigned  domain ID is: 1
```

6. Configure Device E as the intermediate device:

> △ CAUTION:
>
> If the intermediate device is also in an IRF fabric, assign the two IRF fabrics different domain IDs for correct split detection. False detection causes IRF split.

\# Enable the spanning tree feature globally. Map the ARP MAD VLAN to MSTI 1 in the MST region.

```
<DeviceE> system-view
[DeviceE] stp global enable
[DeviceC] stp region-configuration
[DeviceC-mst-region] region-name arpmad
[DeviceC-mst-region] instance 1 vlan 3
[DeviceC-mst-region] active region-configuration
[DeviceC-mst-region] quit
```

\# Create VLAN 3, and assign ports Ten-GigabitEthernet 1/0/1, Ten-GigabitEthernet 1/0/2, Ten-GigabitEthernet 1/0/3, and Ten-GigabitEthernet 1/0/4 to VLAN 3 for forwarding ARP MAD packets.

```
[DeviceE] vlan 3
[DeviceE-vlan3] port ten-gigabitethernet 1/0/1 to ten-gigabitethernet 1/0/4
[DeviceE-vlan3] quit
```

# ND MAD-enabled IRF configuration example

Network requirements

As shown in Figure 15, set up a four-chassis IRF fabric in the IPv6 enterprise network.

- Configure ND MAD in the IRF fabric and use the links connected to Device E for transmitting ND MAD packets.
- To prevent loops, run the spanning tree feature between Device E and the IRF fabric.

Figure 15 Network diagram

Configuration procedure

1. Configure Device A:

   # Shut down the physical interfaces used for IRF links. This example uses the SFP+ port group that contains Ten-GigabitEthernet 1/0/45 to Ten-GigabitEthernet 1/0/48 for IRF links.

   ```
   <Sysname> system-view
   [Sysname] interface range ten-gigabitethernet 1/0/45 to ten-gigabitethernet 1/0/48
   [Sysname-if-range] shutdown
   [Sysname-if-range] quit
   ```

   # Bind Ten-GigabitEthernet 1/0/45 and Ten-GigabitEthernet 1/0/46 to IRF-port 1/1.

   ```
   [Sysname] irf-port 1/1
   [Sysname-irf-port1/1] port group interface ten-gigabitethernet 1/0/45
   [Sysname-irf-port1/1] port group interface ten-gigabitethernet 1/0/46
   [Sysname-irf-port1/1] quit
   ```

   # Bind Ten-GigabitEthernet 1/0/47 and Ten-GigabitEthernet 1/0/48 to IRF-port 1/2.

   ```
   [Sysname] irf-port 1/2
   [Sysname-irf-port1/2] port group interface ten-gigabitethernet 1/0/47
   [Sysname-irf-port1/2] port group interface ten-gigabitethernet 1/0/48
   [Sysname-irf-port1/2] quit
   ```

   # Bring up the SFP+ ports and save the configuration.

   ```
   [Sysname] interface range ten-gigabitethernet 1/0/45 to ten-gigabitethernet 1/0/48
   [Sysname-if-range] undo shutdown
   [Sysname-if-range] quit
   [Sysname] save
   ```

Activate the IRF port configuration.

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

2. Configure Device B:

Change the member ID of Device B to 2 and reboot the device to validate the change.

```
<Sysname> system-view
[Sysname] irf member 1 renumber 2
Renumbering the member ID may result in configuration change or loss. Continue? [Y/N]:y
[Sysname] quit
<Sysname> reboot
```

Connect Device B to Device A as shown in Figure 15, and log in to Device B. This example uses the SFP+ group that contains Ten-GigabitEthernet 2/0/45 to Ten-GigabitEthernet 2/0/48 for IRF links.

Shut down the physical interfaces.

```
<Sysname> system-view
[Sysname] interface range ten-gigabitethernet 2/0/45 to ten-gigabitethernet 2/0/48
[Sysname-if-range] shutdown
[Sysname-if-range] quit
```

Bind Ten-GigabitEthernet 2/0/47 and Ten-GigabitEthernet 2/0/48 to IRF-port 2/1.

```
[Sysname] irf-port 2/1
[Sysname-irf-port2/1] port group interface ten-gigabitethernet 2/0/47
[Sysname-irf-port2/1] port group interface ten-gigabitethernet 2/0/48
[Sysname-irf-port2/1] quit
```

Bind Ten-GigabitEthernet 2/0/45 and Ten-GigabitEthernet 2/0/46 to IRF-port 2/2.

```
[Sysname] irf-port 2/2
[Sysname-irf-port2/2] port group interface ten-gigabitethernet 2/0/45
[Sysname-irf-port2/2] port group interface ten-gigabitethernet 2/0/46
[Sysname-irf-port2/2] quit
```

Bring up the SFP+ ports and save the configuration.

```
[Sysname] interface range ten-gigabitethernet 2/0/45 to ten-gigabitethernet 2/0/48
[Sysname-if-range] undo shutdown
[Sysname-if-range] quit
[Sysname] save
```

Activate the IRF port configuration.

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

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

3. Configure Device C:

Change the member ID of Device C to 3 and reboot the device to validate the change.

```
<Sysname> system-view
[Sysname] irf member 1 renumber 3
Renumbering the member ID may result in configuration change or loss. Continue? [Y/N]:y
[Sysname] quit
<Sysname> reboot
```

Connect Device C to Device A as shown in Figure 15, and log in to Device C. This example uses the SFP+ port group that contains Ten-GigabitEthernet 3/0/45 to Ten-GigabitEthernet 3/0/48 for IRF links.

Shut down the physical interfaces.
```
<Sysname> system-view
[Sysname] interface range ten-gigabitethernet 3/0/45 to ten-gigabitethernet 3/0/48
[Sysname-if-range] shutdown
[Sysname-if-range] quit
```
Bind Ten-GigabitEthernet 3/0/47 and Ten-GigabitEthernet 3/0/48 to IRF-port 3/1.
```
[Sysname] irf-port 3/1
[Sysname-irf-port3/1] port group interface ten-gigabitethernet 3/0/47
[Sysname-irf-port3/1] port group interface ten-gigabitethernet 3/0/48
[Sysname-irf-port3/1] quit
```
Bind Ten-GigabitEthernet 3/0/45 and Ten-GigabitEthernet 3/0/46 to IRF-port 3/2.
```
[Sysname] irf-port 3/2
[Sysname-irf-port3/2] port group interface ten-gigabitethernet 3/0/45
[Sysname-irf-port3/2] port group interface ten-gigabitethernet 3/0/46
[Sysname-irf-port3/2] quit
```
Bring up the SFP+ ports and save the configuration.
```
[Sysname] interface range ten-gigabitethernet 3/0/45 to ten-gigabitethernet 3/0/48
[Sysname-if-range] undo shutdown
[Sysname-if-range] quit
[Sysname] save
```
Activate the IRF port configuration.
```
[Sysname] irf-port-configuration active
```
Device C reboots to join the IRF fabric.

4. Configure Device D:

Change the member ID of Device D to 4 and reboot the device to validate the change.
```
<Sysname> system-view
[Sysname] irf member 1 renumber 4
Renumbering the member ID may result in configuration change or loss. Continue?[Y/N]:y
[Sysname] quit
<Sysname> reboot
```
Connect Device D to Device B and Device C as shown in Figure 15, and log in to Device D. This example uses the SFP+ port group that contains Ten-GigabitEthernet 4/0/45 to Ten-GigabitEthernet 4/0/48 for IRF links.

Shut down the physical interfaces.
```
<Sysname> system-view
[Sysname] interface range ten-gigabitethernet 4/0/45 to ten-gigabitethernet 4/0/48
[Sysname-if-range] shutdown
[Sysname-if-range] quit
```
Bind Ten-GigabitEthernet 4/0/45 and Ten-GigabitEthernet 4/0/46 to IRF-port 4/1.
```
[Sysname] irf-port 4/1
[Sysname-irf-port4/1] port group interface ten-gigabitethernet 4/0/45
[Sysname-irf-port4/1] port group interface ten-gigabitethernet 4/0/46
[Sysname-irf-port4/1] quit
```
Bind Ten-GigabitEthernet 4/0/47 and Ten-GigabitEthernet 4/0/48 to IRF-port 4/2.
```
[Sysname] irf-port 4/2
[Sysname-irf-port4/2] port group interface ten-gigabitethernet 4/0/47
```

```
[Sysname-irf-port4/2] port group interface ten-gigabitethernet 4/0/48
[Sysname-irf-port4/2] quit
```
Bring up the SFP+ ports and save the configuration.
```
[Sysname] interface range ten-gigabitethernet 4/0/45 to ten-gigabitethernet 4/0/48
[Sysname-if-range] undo shutdown
[Sysname-if-range] quit
[Sysname] save
```
Activate the IRF port configuration.
```
[Sysname] irf-port-configuration active
```
Device D reboots to join the IRF fabric. A four-chassis IRF fabric is formed.

5. Configure ND MAD on the IRF fabric:

   # Enable the spanning tree feature globally. Map the ND MAD VLAN to MSTI 1 in the MST region.
   ```
   <Sysname> system-view
   [Sysname] stp global enable
   [Sysname] stp region-configuration
   [Sysname-mst-region] region-name ndmad
   [Sysname-mst-region] instance 1 vlan 3
   [Sysname-mst-region] active region-configuration
   [Sysname-mst-region] quit
   ```
   # Configure the IRF fabric to change its bridge MAC address as soon as the address owner leaves.
   ```
   [Sysname] undo irf mac-address persistent
   ```
   # Set the domain ID of the IRF fabric to 1.
   ```
   [Sysname] irf domain 1
   ```
   # Create VLAN 3, and add Ten-GigabitEthernet 1/0/2, Ten-GigabitEthernet 2/0/1, Ten-GigabitEthernet 3/0/2, and Ten-GigabitEthernet 4/0/1 to VLAN 3.
   ```
   [Sysname] vlan 3
   [Sysname-vlan3] port ten-gigabitethernet 1/0/2 ten-gigabitethernet 2/0/1
   ten-gigabitethernet 3/0/2 ten-gigabitethernet 4/0/1
   [Sysname-vlan3] quit
   ```
   # Create VLAN-interface 3, assign it an IPv6 address, and enable ND MAD on the interface.
   ```
   [Sysname] interface vlan-interface 3
   [Sysname-Vlan-interface3] ipv6 address 2001::1 64
   [Sysname-Vlan-interface3] mad nd enable
   You need to assign a domain ID (range: 0-4294967295)
    [Current domain is: 1]:
    The assigned  domain ID is: 1
   ```

6. Configure Device E as the intermediate device:

⚠ CAUTION:

If the intermediate device is also in an IRF fabric, assign the two IRF fabrics different domain IDs for correct split detection. False detection causes IRF split.

Enable the spanning tree feature globally. Map the ND MAD VLAN to MSTI 1 in the MST region.
```
<DeviceE> system-view
[DeviceE] stp global enable
[DeviceC] stp region-configuration
[DeviceC-mst-region] region-name ndmad
```

```
[DeviceC-mst-region] instance 1 vlan 3
[DeviceC-mst-region] active region-configuration
[DeviceC-mst-region] quit
```

Create VLAN 3, and add ports Ten-GigabitEthernet 1/0/1, Ten-GigabitEthernet 1/0/2, Ten-GigabitEthernet 1/0/3, and Ten-GigabitEthernet 1/0/4 to VLAN 3 for forwarding ND MAD packets.

```
[DeviceE] vlan 3
[DeviceE-vlan3] port ten-gigabitethernet 1/0/1 to ten-gigabitethernet 1/0/4
[DeviceE-vlan3] quit
```

Setting up an eIRF system

Overview

eIRF integrates multiple lower-layer devices with a higher-layer IRF fabric to provide high-density, low-cost connectivity at the access layer.

In an eIRF system, the higher-layer IRF fabric is called the parent fabric and the lower-layer devices are called port extenders (PEXs). You can manage and configure the PEXs from the parent fabric as if they were interface cards on the parent fabric.

Typically, eIRF works at the access layer of data centers. As shown in Figure 16, the access layer of a network is virtualized into an eIRF system. The system contains one parent fabric (a two-chassis IRF fabric) and multiple PEXs to provide connectivity for servers and hosts.

Figure 16 eIRF application scenario
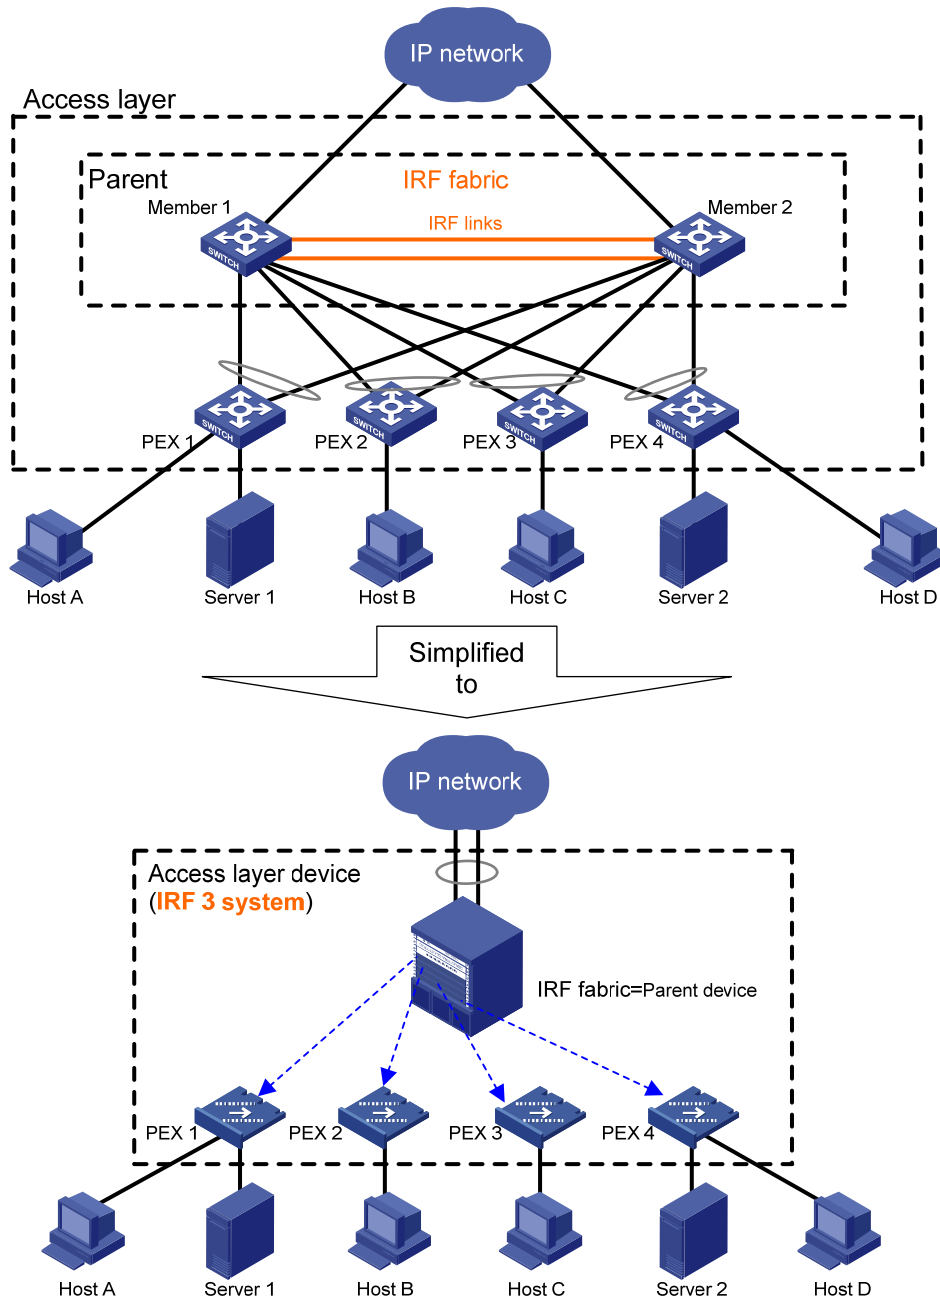
eIRF provides the following benefits:

- **Simplified topology**—Devices in an eIRF system appear as one node. For redundancy and load balancing, a downstream or upstream device can connect to the eIRF system through multichassis link aggregation. Together with link aggregation, eIRF creates a loop-free Layer 2 network. The spanning tree feature is not needed among devices in the eIRF system or on the link aggregations. eIRF also simplifies the Layer 3 network topology because it reduces the number of routing peers. The network topology does not change when a device is added to or removed from the eIRF system.

- **Single point of management**—An eIRF system is accessible at a single IP address on the network. You can use this IP address to log in through any network port to manage all the devices in the system. For an SNMP NMS, an eIRF system is one managed network node.

- **Network scalability and resiliency**—You can increase the number of ports in an eIRF system by adding PEXs without changing network topology.

- **High availability**—Each PEX has multiple high-speed physical ports for uplink connectivity to the parent fabric. The links on these ports are aggregated and load balanced automatically.

- **Decreased TCO**—eIRF decreases hardware investments and management costs. In an eIRF system, the parent fabric performs all the management and routing functions, and the PEXs only forwards traffic. You can add low-performance devices as PEXs to an eIRF system for network scalability. In addition, PEXs can load software and synchronize configuration from the parent fabric without administrative intervention.

Basic concepts

eIRF includes IRF concepts and adds the concepts in this section.

eIRF roles

The devices in an eIRF system have the following roles:

- **Parent fabric**—Higher-layer single-chassis or multichassis IRF fabric that controls the entire eIRF system, including PEXs. Each eIRF system has one parent fabric.

- **Parent device**—Member devices in the parent fabric.

- **Master device**—Controls and manages the entire eIRF system, including all parent devices and PEXs. The master device in the IRF fabric is also the master device for the eIRF system. You configure all devices (including PEXs and parent devices) from the master device.

- **PEX**—Operates as I/O modules of the parent fabric to receive and transmit traffic. All forwarding decisions are made on the parent fabric. Table 3 shows the operating states of PEXs.

Table 3 PEX operating states

State	Description
Loading	The PEX is starting up. To avoid problems, do not reboot a PEX while it is in Loading state.
Online	The PEX has started up and registered with the parent fabric. You can configure and manage the PEX from the parent fabric.
Offline	Any of the following conditions exist: • The PEX has not started up. • The PEX has lost its uplink connectivity to the parent fabric. • The PEX configuration on the parent fabric is incomplete.

PEX port

A PEX port is a logical port created on the parent fabric for managing a PEX. For each PEX, you must create a unique PEX port, and assign all physical interfaces connected to the PEX to the PEX port.

PEX physical interface

PEX physical interfaces connect PEXs and the parent fabric.

You can set up multiple PEX links between the parent fabric and a PEX. These links aggregate automatically for backup and load balancing.

On the parent fabric, you must assign the PEX physical interfaces for one PEX to the same PEX port. On a PEX, you must connect all its PEX physical interfaces to the physical interfaces in the same PEX port.

Table 4 describes the states of PEX physical interfaces.

Table 4 PEX physical interface states

State	Description
Forwarding	The PEX physical interface is operating correctly and can forward data traffic.
Down	The physical link is disconnected. The PEX physical interface cannot forward any packets.
Blocked	The PEX physical interface cannot forward any packets except for eIRF packets. The Blocked state is a transitional state between Forwarding and Down. A PEX physical interface changes to the Blocked state in the following situations: • Incorrect physical connection: ○ The PEX has PEX links to more than one PEX port on the parent fabric. ○ The PEX port on the parent fabric contains physical links to more than one PEX. • The data link is forced to the Blocked state. In the startup phase, the PEX blocks a PEX physical interface if the interface's physical link is up, but it is not used for loading startup software. • The physical state of the interface is up, but the PEX link to the parent fabric has been disconnected. The PEX and the parent fabric cannot receive eIRF heartbeat packets from each other.

Virtual slot number

Each PEX is identified by a unique virtual slot number in an eIRF system.

After a PEX joins an eIRF system, the first segment in its interface numbers changes to the virtual slot number assigned to the PEX. For example, a PEX has an interface numbered 1/0/1 before it is added to an eIRF system. After it is added to an eIRF system as slot 100, the interface number changes to 100/0/1.

eIRF operating mechanisms

eIRF membership establishment

After you complete PEX configuration for a PEX, the master device monitors the PEX physical interfaces for the slot number request from the PEX. The PEX uses the following process to join the eIRF system:

1. Virtual slot assignment:
   a. At startup, the PEX sends a slot number request to the master device.
   b. The master device assigns the user-configured slot number to the PEX.
2. Software loading and PEX registration:
   a. The PEX checks the local storage medium for valid startup images.
   b. If the local startup images are valid, the PEX verifies their compatibility with the running software images on the master device:
      – If all the images are compatible, the PEX starts up with the local startup images.
      – If any incompatibility exists, the PEX goes to step c.
   c. The PEX sends a startup software request to the master device.

d. The master device provides the correct Boot ROM and startup software images to the PEX.

e. The PEX downloads the software images to the local storage medium, and then loads the software images.

f. The PEX automatically reboots to register with the master device.

3. Configuration synchronization:

a. The parent device issues its running configuration to the PEX.

b. The PEX runs with the configuration received from the master device instead of reading the configuration from its local startup configuration file.

4. PEX link maintenance:

The master device and the PEX send heartbeat packets on the PEX links to detect link failure. You can determine whether a PEX link is available by checking the state of its physical interfaces.

Configuration management

An eIRF system manages all its settings (including settings for PEXs) on the master device. You can configure and manage PEXs only from the master device. The running configuration on the master device has all settings in the eIRF system, including settings for PEXs. When a PEX reboots or is added, the master device issues its running configuration to the PEX.

Data forwarding

The PEXs do not have local forwarding capability.

The PEXs send any incoming traffic to the parent fabric. The parent fabric makes the forwarding decisions and sends the traffic to the outgoing interfaces. Figure 17 shows the data forwarding model.

Figure 17 Data forwarding model
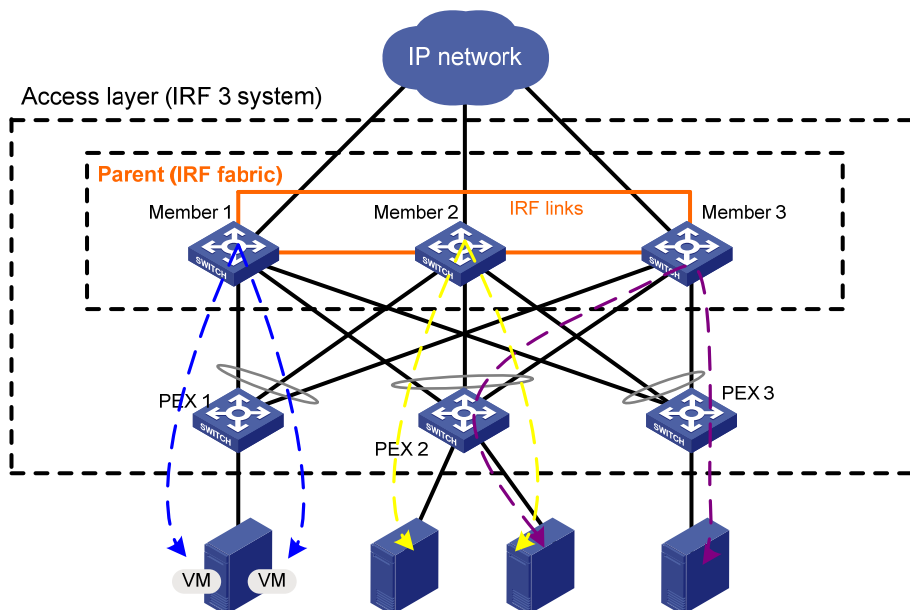
Hardware compatibility

HP FlexFabric 5700 switches can only operate as PEXs. The parent devices are from HP 5900 switch series. For more information, see the IRF configuration guide for the HP 5900 switch.

Configuration restrictions and guidelines

For a successful eIRF system setup, read the configuration restrictions and guidelines carefully before you connect and set up a PEX.

PEX physical interface requirements

Use Table 5 to identify physical interfaces for PEX links.

Table 5 Candidate PEX physical interfaces

Switch model	Candidate PEX physical interfaces
HP FlexFabric 5700 switches	• The highest-numbered four SFP+ ports on the front panel. • QSFP+ ports on the front panel. NOTE: Do not use both SFP+ and QSFP+ ports for PEX links.

PEX cabling requirements

When you connect the parent fabric and PEXs, follow these cabling restrictions and guidelines:

- On the parent fabric, connect a PEX port's all physical interfaces to the same PEX.
- On a PEX, connect all its PEX physical interfaces to the physical interfaces in the same PEX port.
- Do not connect PEXs to each other.
- eIRF only supports one layer of PEXs. You cannot attach a lower-layer PEX to a higher-layer PEX.

Configuring the operating mode of the device

An eIRF-capable device can operate in the following modes:

- **Auto**—The device automatically changes to the PEX mode to join an eIRF system when it detects any PEX links in up state. Before that, the device operates as an independently node and is configurable from its CLI.

  NOTE:
  - The auto mode does not take effect if IRF port settings are configured on the device. For the auto mode to take effect, you must remove these settings.
  - The auto mode runs only for 24 hours. The auto mode stops monitoring PEX links if the device fails to join an eIRF system before the timer expires. To add the device to an eIRF system, you must manually change its mode to PEX mode.

- **PEX**—The device can only be managed from the parent fabric in an eIRF system. You cannot manage the device from the CLI of the device. If no PEX links are available to the parent fabric, you can only access the device from the Boot menu.

- **Switch**—The device operates independently as a node. It is not part of an eIRF system, whether or not it has connections to the parent fabric.

Configuring the device as a PEX

To add the device to an eIRF system as a PEX, you must change its operating mode to auto or PEX.

To configure the device as a PEX:

Step		Command	Remarks
1.	Enter system view.	**system-view**	N/A
2.	Change the operating mode to auto or PEX.	**pex working-mode** { **auto** | **pex** } { **all** | **slot** *slot-number1* [**to** *slot-number2*] }	By default, the device operates in auto mode. To cancel the change, configure the **pex working-mode** { **auto** | **switch** } { **all** | **slot** *slot-number1* [**to** *slot-number2*] } command before the reboot.
3.	Return to user view.	**quit**	N/A
4.	Save the running configuration.	**save**	N/A
5.	Reboot the device.	**reboot**	The change takes effect at reboot. If PEX mode is used, you will be unable to access the CLI of the device. To change the operating mode of the PEX, you must perform the task from its Boot menu or the parent fabric.

Configuring the device as an independent switch

You can set the device to switch mode from its CLI when the device is operating in auto mode.

To configure the device as an independent switch from its CLI:

Step		Command	Remarks
1.	Enter system view.	**system-view**	N/A
2.	Change the operating mode to switch.	**pex working-mode switch** { **all** | **slot** *slot-number1* [**to** *slot-number2*] }	By default, the device operates in auto mode. To cancel the change, configure the **pex working-mode** { **auto** | **pex** } { **all** | **slot** *slot-number1* [**to** *slot-number2*] } command before the reboot.
3.	Return to user view.	**quit**	N/A
4.	Save the running configuration.	**save**	N/A
5.	Reboot the device.	**reboot**	The change takes effect at reboot.

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.
Italic	*Italic* 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.
① IMPORTANT	An alert that calls attention to essential information.
NOTE	An alert that contains additional or supplementary information.
·ϙ· 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 mesh access point.
	Represents omnidirectional signals.
	Represents directional signals.
	Represents a security product, such as a firewall, UTM, multiservice security gateway, or load-balancing device.
	Represents a security card, such as a firewall, load-balancing, NetStream, SSL VPN, IPS, or ACG card.

Port numbering in examples

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

Index