

Configuring Ethernet link aggregation

Ethernet link aggregation bundles multiple physical Ethernet links into one logical link, called an aggregate link. Link aggregation has the following benefits:

- Increased bandwidth beyond the limits of any single link. In an aggregate link, traffic is distributed across the member ports.
- Improved link reliability. The member ports dynamically back up one another. When a member port fails, its traffic is automatically switched to other member ports.

As shown in [Figure 8](#), Device A and Device B are connected by three physical Ethernet links. These physical Ethernet links are combined into an aggregate link called link aggregation 1. The bandwidth of this aggregate link can reach up to the total bandwidth of the three physical Ethernet links. At the same time, the three Ethernet links back up one another. When a physical Ethernet link fails, the traffic previously carried on the failed link is switched to the other two links.

Figure 8 Ethernet link aggregation diagram



Basic concepts

Aggregation group, member port, and aggregate interface

Link bundling is implemented through interface bundling. An aggregation group is a group of Ethernet interfaces bundled together, which are called member ports of the aggregation group. For each aggregation group, a logical interface (called an aggregate interface), is created. When an upper layer entity uses the link aggregation service, the following events occur:

- A link aggregation group appears the same as a single logical link.
- An aggregate interface transmits data traffic.

When you create an aggregate interface, the device automatically creates an aggregation group of the same type and number as the aggregate interface. For example, when you create aggregate interface 1, aggregation group 1 is created.

Aggregate interfaces include Layer 2 aggregate interfaces and Layer 3 aggregate interfaces.

You can assign Layer 2 Ethernet interfaces only to a Layer 2 aggregation group, and Layer 3 Ethernet interfaces only to a Layer 3 aggregation group.

NOTE:

You can use the **port link-mode** command to configure an Ethernet port as a Layer 2 or Layer 3 interface (see "[Configuring Ethernet interfaces](#)").

The port rate of an aggregate interface equals the total rate of its Selected member ports. Its duplex mode is the same as that of the Selected member ports. For more information about the states of member ports in an aggregation group, see "[Aggregation states of member ports in an aggregation group](#)."

Aggregation states of member ports in an aggregation group

A member port in an aggregation group can be in any of the following aggregation states:

- **Selected**—A Selected port can forward traffic.
- **Unselected**—An Unselected port cannot forward traffic.
- **Individual**—An Individual port can forward traffic as a normal physical port. A port is placed in the Individual state when the following conditions are met:
 - The corresponding aggregate interface is configured as an edge aggregate interface.
 - The port does not receive LACPDU from its peer port.

Operational key

When aggregating ports, the system automatically assigns each port an operational key based on port information, such as port rate and duplex mode. Any change to this information triggers a recalculation of the operational key.

In an aggregation group, all Selected ports are assigned the same operational key.

Configuration types

Every configuration setting on a port might affect its aggregation state. Port configurations include the following types:

- **Attribute configurations**—To become a Selected port, a member port must have the same attribute configurations as the aggregate interface. [Table 2](#) describes the attribute configurations.

Attribute configurations made on an aggregate interface are automatically synchronized to all member ports. These configurations are retained on the member ports even after the aggregate interface is removed.

Any attribute configuration change on a member port might affect the aggregation states and running services of the member ports. The system displays a warning message every time you try to change an attribute configuration setting on a member port.

Table 2 Attribute configurations

Feature	Considerations
Port isolation	Indicates whether the port has joined an isolation group, and which isolation group the port belongs to.
QinQ	QinQ enable state (enabled/disabled), TPID for VLAN tags, and VLAN transparent transmission. For information about QinQ, see " Configuring QinQ ."
VLAN mapping	Different types of VLAN mapping configured on the port. For more information about VLAN mapping, see " Configuring VLAN mapping ."
VLAN	<p>VLAN attribute configurations include the following:</p> <ul style="list-style-type: none"> • Permitted VLAN IDs. • PVID. • Link type (trunk, hybrid, or access). • Operating mode (promiscuous, trunk promiscuous, host). • VLAN tagging mode. <p>For information about VLAN, see "Configuring VLANs."</p>

- **Protocol configurations**—Protocol configurations do not affect the aggregation state of the member ports. MAC address learning and spanning tree settings are examples of protocol configurations.

NOTE:

The protocol configuration for a member port is effective only when the member port leaves the aggregation group.

Link aggregation modes

Link aggregation has dynamic and static modes:

- **Static aggregation mode**—Aggregation is stable. The aggregation state of the member ports are not affected by the peer ports.
- **Dynamic aggregation mode**—The peering system automatically maintains the aggregation state of the member ports, which reduces the administrators' workload.

An aggregation group in static mode is called a "static aggregation group" and an aggregation group in dynamic mode is called a "dynamic aggregation group."

Aggregating links in static mode

Choosing a reference port

When setting the aggregation state of the ports in an aggregation group, the system automatically chooses a member port as the reference port. A Selected port must have the same operational key and attribute configurations as the reference port.

The system chooses a reference port from the member ports that are in up state with the same attribute configurations as the aggregate interface.

The candidate ports are sorted in the following order:

1. Highest port priority
2. Full duplex/high speed
3. Full duplex/low speed
4. Half duplex/high speed
5. Half duplex/low speed

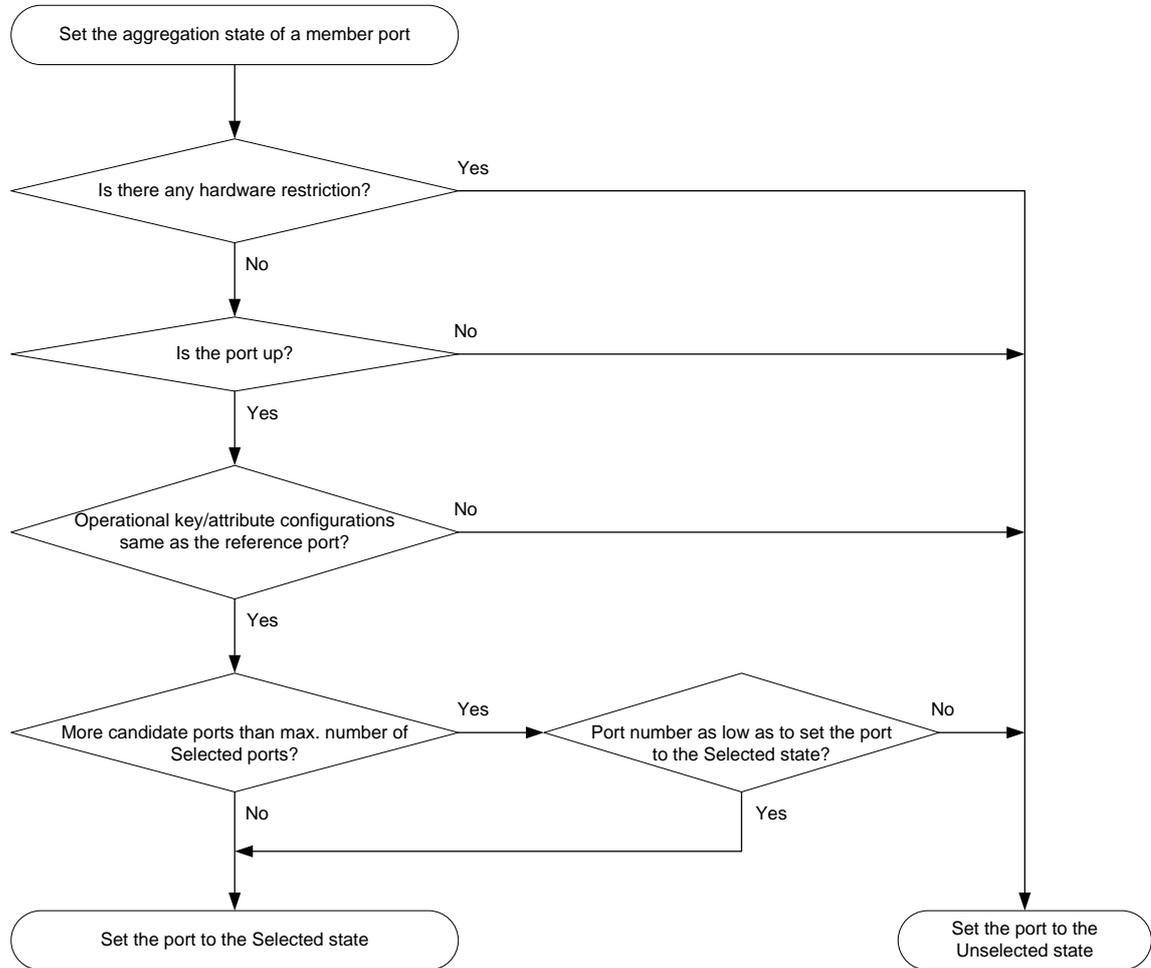
The candidate port at the top is chosen as the reference port.

- If multiple ports have the same port priority, duplex mode, and speed, the port that has been a Selected port (if any) is chosen. If multiple ports have been Selected ports, the one with the smallest port number is chosen.
- If multiple ports have the same port priority, duplex mode, and speed and none of them has been a Selected port, the port with the smallest port number is chosen.

Setting the aggregation state of each member port

After a static aggregation group reaches the limit on Selected ports, ports attempting to join the group are put in Unselected state. This prevents traffic interruption on the existing Selected ports.

Figure 9 Setting the aggregation state of a member port in a static aggregation group



For more information about configuring the maximum number of Selected ports in a static aggregation group, see "[Setting the minimum and maximum numbers of Selected ports for an aggregation group.](#)"

Any operational key or attribute configuration change might affect the aggregation states of link aggregation member ports and ongoing traffic.

Aggregating links in dynamic mode

Dynamic aggregation mode is implemented through IEEE 802.3ad Link Aggregation Control Protocol (LACP).

LACP

LACP uses LACPDU to exchange aggregation information between LACP-enabled devices.

Each member port in an LACP-enabled aggregation group exchanges information with its peer. When a member port receives an LACPDU, it compares the received information with information received on the other member ports. In this way, the two systems reach an agreement on which ports are placed in the Selected state.

LACP functions

LACP offers basic LACP functions and extended LACP functions, as described in [Table 3](#).

Table 3 Basic and extended LACP functions

Category	Description
Basic LACP functions	Implemented through the basic LACPDU fields, including the system LACP priority, system MAC address, port priority, port number, and operational key.
Extended LACP functions	Implemented by extending the LACPDU with new TLV fields. This is how the LACP MAD mechanism of the IRF feature is implemented. The switch series can participate in LACP MAD as either an IRF member device or an intermediate device. For more information about IRF and the LACP MAD mechanism, see <i>IRF Configuration Guide</i> .

LACP operating modes

LACP can operate in active mode or passive mode.

When LACP is operating in passive mode on a local member port and its peer port, both ports cannot send LACPDUs. When LACP is operating in active mode on the port on either end of a link, both ports can send LACPDUs.

LACP priorities

LACP priorities include system LACP priority and port priority, as described in [Table 4](#). The smaller the priority value, the higher the priority.

Table 4 LACP priorities

Type	Description
System LACP priority	Used by two peer devices (or systems) to determine which one is superior in link aggregation. In dynamic link aggregation, the system that has higher system LACP priority sets the Selected state of member ports on its side. The system that has lower priority sets the aggregation state of local member ports the same as their respective peer ports.
Port priority	Determines the likelihood of a member port to be selected on a system. A port with a higher port priority is more likely to become Selected.

LACP timeout interval

The LACP timeout interval specifies how long a member port waits to receive LACPDUs from the peer port. If a local member port does not receive LACPDUs from the peer within the LACP timeout interval, the member port considers the peer as failed.

The LACP timeout interval also determines the LACPDU sending rate of the peer. LACP timeout intervals include the following types:

- **Short timeout interval**—3 seconds. If you use the short timeout interval, the peer sends one LACPDU per second.
- **Long timeout interval**—90 seconds. If you use the long timeout interval, the peer sends one LACPDU every 30 seconds.

How dynamic link aggregation works

Choosing a reference port

The system chooses a reference port from the member ports that are in up state and have the same attribute configurations as the aggregate interface. A Selected port must have the same operational key and attribute configurations as the reference port.

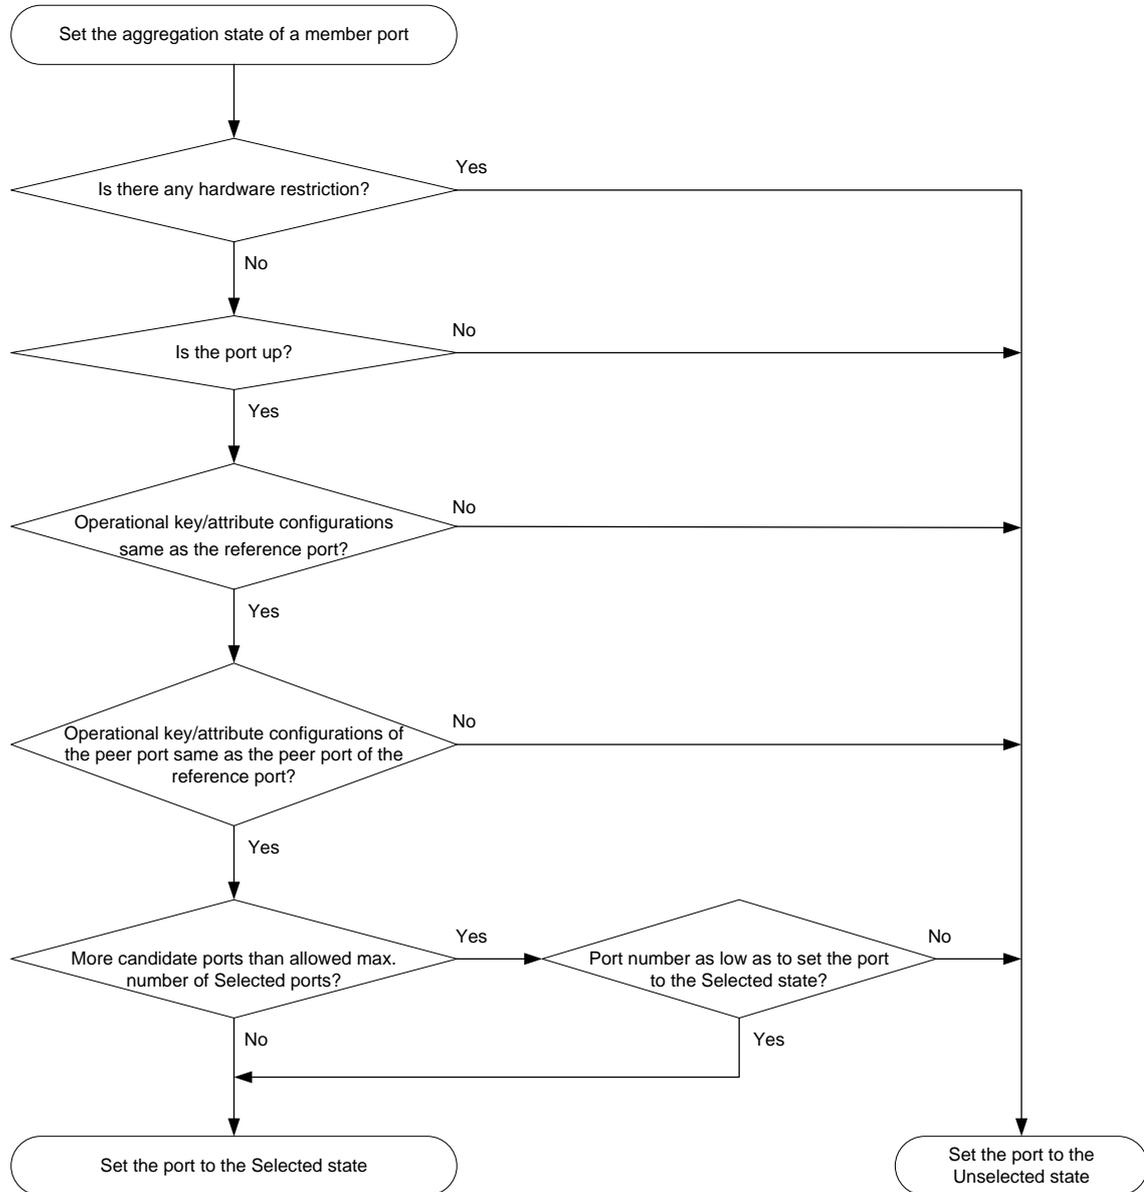
The local system (the actor) and the remote system (the partner) negotiate a reference port by using the following workflow:

1. The two systems determine the system with the smaller system ID.
A system ID contains the system LACP priority and the system MAC address.
 - a. The two systems compare their LACP priority values.
The lower the LACP priority, the smaller the system ID. If the LACP priority values are the same, the two systems proceed to step b.
 - b. The two systems compare their MAC addresses.
The lower the MAC address, the smaller the system ID.
2. The system with the smaller system ID chooses the port with the smallest port ID as the reference port.
A port ID contains a port priority and a port number. The lower the port priority, the smaller the port ID.
 - a. The system chooses the port with the lowest priority value as the reference port.
If the ports have the same priority, the system proceeds to step b.
 - b. The system compares their port numbers.
The smaller the port number, the smaller the port ID.
The port with the smallest port number and the same attribute configurations as the aggregate interface is chosen as the reference port.

Setting the aggregation state of each member port

After the reference port is chosen, the system with the smaller system ID sets the state of each member port on its side.

Figure 10 Setting the state of a member port in a dynamic aggregation group



The system with the higher system ID is aware of the aggregation state changes on the remote system. The system sets the aggregation state of local member ports the same as their peer ports.

When you aggregate interfaces in dynamic mode, follow these guidelines:

- A dynamic link aggregation group preferably chooses full-duplex ports as the Selected ports. The group chooses only one half-duplex port as a Selected port when either of the following conditions exist:
 - None of the full-duplex ports can become Selected ports.
 - Only half-duplex ports exist in the group.
- To ensure stable aggregation and service continuity, do not change the operational key or attribute configurations on any member port.
- When the aggregation state of a local port changes in a dynamic aggregation group, the aggregation state of the peer port also changes.
- After the Selected port limit has been reached, a port joining the aggregation group is placed in the Selected state if it is more eligible than a current member port.

For more information about configuring the maximum number of Selected ports in a dynamic aggregation group, see "[Setting the minimum and maximum numbers of Selected ports for an aggregation group.](#)"

Edge aggregate interface

Dynamic link aggregation fails on a server-facing aggregate interface if dynamic link aggregation is configured only on the device. The device forwards traffic by using only one of the physical ports that are connected to the server.

To improve link reliability, configure the aggregate interface as an edge aggregate interface. This feature enables all member ports of the aggregation group to forward traffic. When a member port fails, its traffic is automatically switched to other member ports.

After dynamic link aggregation is configured on the server, the device can receive LACPDU from the server. Then, link aggregation between the device and the server operates correctly.

An edge aggregate interface takes effect only when it is configured on an aggregate interface corresponding to a dynamic aggregation group.

Load sharing modes for link aggregation groups

In a link aggregation group, traffic can be load shared across the Selected ports based on any of the following modes:

- **Per-flow load sharing**—Load shares traffic on a per-flow basis. The load sharing mode classifies packets into flows and forwards packets of the same flow on the same link. This mode can be one or any combination of the following criteria that classify traffic:
 - Source or destination MAC address.
 - Source or destination port number.
 - Ingress port.
 - Source or destination IP address.
- **Packet type-based load sharing**—Load shares traffic automatically based on packet types (Layer 2, IPv4, or IPv6 for example).

Ethernet link aggregation configuration task list

Tasks at a glance
(Required.) Configuring an aggregation group: <ul style="list-style-type: none">• Configuring a static aggregation group• Configuring a dynamic aggregation group
(Optional.) Configuring an aggregate interface: <ul style="list-style-type: none">• Configuring the description of an aggregate interface• Specifying ignored VLANs for a Layer 2 aggregate interface• Setting the MTU for a Layer 3 aggregate interface• Setting the minimum and maximum numbers of Selected ports for an aggregation group• Setting the expected bandwidth for an aggregate interface• Configuring an edge aggregate interface• Enabling BFD for an aggregation group• Shutting down an aggregate interface• Restoring the default settings for an aggregate interface

Tasks at a glance

(Optional.) Configuring load balancing for link aggregation group:

- [Setting load sharing modes for link aggregation groups](#)
- [Enabling local-first load sharing for link aggregation](#)
- [Configuring per-flow load sharing algorithm settings for Ethernet link aggregation](#)

[Enabling link-aggregation traffic redirection](#)

[Specifying link aggregation management VLANs and management port](#)

Configuring an aggregation group

This section explains how to configure an aggregation group.

Configuration restrictions and guidelines

When you configure an aggregation group, follow these restrictions and guidelines:

- You cannot assign a port to a Layer 2 aggregation group if any of the following features are configured on the port:
 - MAC authentication. (See *Security Configuration Guide*.)
 - Port security. (See *Security Configuration Guide*.)
 - 802.1X. (See *Security Configuration Guide*.)
 - Association between AC and cross connection. (See *MPLS Configuration Guide*.)
 - AC-VSI association. (See *MPLS Configuration Guide*.)
- You cannot assign a port to a Layer 3 aggregation group if any of the following features are configured on the port:
 - Association between AC and cross connection. (See *MPLS Configuration Guide*.)
 - AC-VSI association. (See *MPLS Configuration Guide*.)
- If a port is used as a reflector port for port mirroring, do not assign it to an aggregation group. For more information about reflector ports, see *Network Management and Monitoring Configuration Guide*.
- Removing an aggregate interface also removes its aggregation group and causes all member ports to leave the aggregation group.
- You must configure the same aggregation mode on the two ends of an aggregate link.
- This switch series supports a maximum of 128 aggregation groups. To ensure the operation of the service loopback groups, configure no more than 126 aggregation groups.

Configuring a static aggregation group

To guarantee a successful static aggregation, make sure that the ports at both ends of each link are in the same aggregation state.

Avoid assigning ports to a static aggregation group that has reached the limit on Selected ports. These ports will be placed in the Unselected state to avoid traffic interruption on the current Selected ports. However, a device reboot can cause the aggregation state of member ports to change.

Configuring a Layer 2 static aggregation group

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

Step	Command	Remarks
2. Create a Layer 2 aggregate interface and enter Layer 2 aggregate interface view.	interface bridge-aggregation <i>interface-number</i>	When you create a Layer 2 aggregate interface, the system automatically creates a Layer 2 static aggregation group numbered the same.
3. Exit to system view.	quit	N/A
4. Assign an interface to the specified Layer 2 aggregation group.	<ul style="list-style-type: none"> a. Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> b. Assign the interface to the specified Layer 2 aggregation group: port link-aggregation group <i>number</i> 	Repeat these two sub-steps to assign more Layer 2 Ethernet interfaces to the aggregation group.

Configuring a Layer 3 static aggregation group

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Create a Layer 3 aggregate interface and enter Layer 3 aggregate interface view.	interface route-aggregation <i>interface-number</i>	When you create a Layer 3 aggregate interface, the system automatically creates a Layer 3 static aggregation group numbered the same.
3. Exit to system view.	quit	N/A
4. Assign an interface to the specified Layer 3 aggregation group.	<ul style="list-style-type: none"> a. Enter Layer 3 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> b. Assign the interface to the specified Layer 3 aggregation group: port link-aggregation group <i>number</i> 	Repeat these two sub-steps to assign more Layer 3 Ethernet interfaces to the aggregation group.

Configuring a dynamic aggregation group

To guarantee a successful dynamic aggregation, make sure that the peer ports of the ports aggregated at one end are also aggregated. The two ends can automatically negotiate the aggregation state of each member port.

Configuring a Layer 2 dynamic aggregation group

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Set the system LACP priority.	lacp system-priority <i>system-priority</i>	By default, the system LACP priority is 32768. Changing the system LACP priority might affect the aggregation state of the ports in a

Step	Command	Remarks
		dynamic aggregation group.
3. Create a Layer 2 aggregate interface and enter Layer 2 aggregate interface view.	interface bridge-aggregation <i>interface-number</i>	When you create a Layer 2 aggregate interface, the system automatically creates a Layer 2 static aggregation group numbered the same.
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.
5. Exit to system view.	quit	N/A
6. Assign an interface to the specified Layer 2 aggregation group.	<p>a. Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i></p> <p>b. Assign the interface to the specified Layer 2 aggregation group: port link-aggregation group <i>number</i></p>	Repeat these two sub-steps to assign more Layer 2 Ethernet interfaces to the aggregation group.
7. Set the LACP operating mode for the interface.	<ul style="list-style-type: none"> Set LACP to operate in passive mode: lacp mode passive Set LACP to operate in active mode: undo lacp mode 	By default, LACP is operating in active mode.
8. Configure the port priority for the interface.	link-aggregation port-priority <i>port-priority</i>	The default setting is 32768.
9. Set the short LACP timeout interval (3 seconds) for the interface.	lacp period short	By default, the long LACP timeout interval (90 seconds) is adopted by the interface. The peer sends LACPDUs slowly.

Configuring a Layer 3 dynamic aggregation group

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Set the system LACP priority.	lacp system-priority <i>system-priority</i>	By default, the system LACP priority is 32768. Changing the system LACP priority might affect the aggregation state of the ports in the dynamic aggregation group.
3. Create a Layer 3 aggregate interface and enter Layer 3 aggregate interface view.	interface route-aggregation <i>interface-number</i>	When you create a Layer 3 aggregate interface, the system automatically creates a Layer 3 static aggregation group numbered the same.
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.
5. Exit to system view.	quit	N/A

Step	Command	Remarks
6. Assign an interface to the specified Layer 3 aggregation group.	<ul style="list-style-type: none"> a. Enter Layer 3 Ethernet interface view: interface <i>interface-type interface-number</i> b. Assign the interface to the specified Layer 3 aggregation group: port link-aggregation group <i>number</i> 	Repeat these two sub-steps to assign more Layer 3 Ethernet interfaces to the aggregation group.
7. Set the LACP operating mode for the interface.	<ul style="list-style-type: none"> • Set LACP to operate in passive mode: lacp mode passive • Set LACP to operate in active mode: undo lacp mode 	By default, LACP is operating in active mode.
8. Configure the port priority for the interface.	link-aggregation port-priority <i>port-priority</i>	The default setting is 32768.
9. Set the short LACP timeout interval (3 seconds) for the interface.	lacp period short	By default, the long LACP timeout interval (90 seconds) is adopted by the interface. The peer sends LACPDUs slowly.

Configuring an aggregate interface

Most configurations that can be made on Layer 2 or Layer 3 Ethernet interfaces can also be made on Layer 2 or Layer 3 aggregate interfaces.

Configuring the description of an aggregate interface

You can configure the description of an aggregate interface for administration purposes, for example, describing the purpose of the interface.

To configure the description of an aggregate interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	<ul style="list-style-type: none"> • Enter Layer 2 aggregate interface view: interface bridge-aggregation <i>interface-number</i> • Enter Layer 3 aggregate interface view: interface route-aggregation <i>interface-number</i> 	N/A
3. Configure the description of the aggregate interface.	description <i>text</i>	By default, the description of an interface is <i>interface-name</i> Interface .

Specifying ignored VLANs for a Layer 2 aggregate interface

By default, to become Selected ports, the member ports must have the same VLAN permit state and VLAN tagging mode as the corresponding Layer 2 aggregate interface.

The system ignores the permit state and tagging mode of an ignored VLAN when choosing Selected ports.

To specify ignored VLANs for a Layer 2 aggregate interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 2 aggregate interface view.	interface bridge-aggregation <i>interface-number</i>	N/A
3. Specify ignored VLANs.	link-aggregation ignore vlan <i>vlan-id-list</i>	By default, a Layer 2 aggregate interface does not ignore any VLANs.

Setting the MTU for a Layer 3 aggregate interface

The MTU of an interface affects IP packets fragmentation and reassembly on the interface.

To set the MTU for a Layer 3 aggregate interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 3 aggregate interface view.	interface route-aggregation <i>interface-number</i>	N/A
3. Set the MTU for the Layer 3 aggregate interface.	mtu <i>size</i>	The default setting is 1500 bytes.

Setting the minimum and maximum numbers of Selected ports for an aggregation group

ⓘ IMPORTANT:

The minimum and maximum numbers of Selected ports must be the same for the local and peer aggregation groups.

The bandwidth of an aggregate link increases as the number of Selected member ports increases. To avoid congestion, you can set the minimum number of Selected ports required for bringing up an aggregate interface.

This minimum threshold setting affects the aggregation state of both aggregation member ports and the aggregate interface.

- When the number of member ports eligible to be Selected ports is smaller than the minimum threshold:
 - All member ports are placed in Unselected state.
 - The link of the aggregate interface goes down.
- When the minimum threshold is reached, the eligible member ports are placed in the Selected state, and the link of the aggregate interface goes up.

The maximum number of Selected ports allowed in an aggregation group is limited by either manual configuration or hardware limitation, whichever value is smaller.

You can implement backup between two ports by performing the following tasks:

- Assigning two ports to an aggregation group.
- Setting the maximum number of Selected ports to 1 for the aggregation group.

Then, only one Selected port is allowed in the aggregation group, and the Unselected port acts as a backup port.

To set the minimum and maximum numbers of Selected ports for an aggregation group:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	<ul style="list-style-type: none"> • Enter Layer 2 aggregate interface view: interface bridge-aggregation <i>interface-number</i> • Enter Layer 3 aggregate interface view: interface route-aggregation <i>interface-number</i> 	N/A
3. Set the minimum number of Selected ports for the aggregation group.	link-aggregation selected-port minimum <i>number</i>	By default, the minimum number of Selected ports is not specified for an aggregation group.
4. Set the maximum number of Selected ports for the aggregation group.	link-aggregation selected-port maximum <i>number</i>	By default, the maximum number of Selected ports for an aggregation group is 32.

Setting the expected bandwidth for an aggregate interface

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	<ul style="list-style-type: none"> • Enter Layer 2 aggregate interface view: interface bridge-aggregation <i>interface-number</i> • Enter Layer 3 aggregate interface view: interface route-aggregation <i>interface-number</i> 	N/A
3. Set the expected bandwidth for the interface.	bandwidth <i>bandwidth-value</i>	By default, the expected bandwidth (in kbps) is the interface baud rate divided by 1000.

Configuring an edge aggregate interface

When you configure an edge aggregate interface, follow these restrictions and guidelines:

- This configuration takes effect only on the aggregate interface corresponding to a dynamic aggregation group.
- Link-aggregation traffic redirection does not operate correctly on an edge aggregate interface. For more information about link-aggregation traffic redirection, see "[Enabling link-aggregation traffic redirection.](#)"

To configure an edge aggregate interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	<ul style="list-style-type: none"> • Enter Layer 2 aggregate interface view: interface bridge-aggregation interface-number • Enter Layer 3 aggregate interface view: interface route-aggregation interface-number 	N/A
3. Configure the aggregate interface as an edge aggregate interface.	lACP edge-port	By default, an aggregate interface does not operate as an edge aggregate interface.

Enabling BFD for an aggregation group

BFD for Ethernet link aggregation can monitor member link status in an aggregation group. After you enable BFD on an aggregate interface, each Selected port in the aggregation group establishes a BFD session with its peer port. BFD operates differently depending on the aggregation mode.

- **BFD for static aggregation**—When BFD detects a link failure, BFD notifies the Ethernet link aggregation module that the peer port is unreachable. The local port is placed in Unselected state. The BFD session between the local and peer ports remains, and the local port keeps sending BFD packets. When the link is recovered, the local port receives BFD packets from the peer port, and BFD notifies the Ethernet link aggregation module that the peer port is reachable. The local port is placed in Selected state again. This mechanism ensures that the local and peer ports of a static aggregate link have the same aggregation state.
- **BFD for dynamic aggregation**—When BFD detects a link failure, BFD notifies the Ethernet link aggregation module that the peer port is unreachable. BFD clears the session and stops sending BFD packets. When the link is recovered and the local port is placed in Selected state again, the local port establishes a new session with the peer port. BFD notifies the Ethernet link aggregation module that the peer port is reachable. Because BFD provides fast failure detection, the local and peer systems of a dynamic aggregate link can negotiate the aggregation state of their member ports faster.

For more information about BFD, see *High Availability Configuration Guide*.

Configuration restrictions and guidelines

When you enable BFD for an aggregation group, follow these restrictions and guidelines:

- Make sure the source and destination IP addresses are consistent at the two ends of an aggregate link. For example, if you execute **link-aggregation bfd ipv4 source 1.1.1.1 destination 2.2.2.2** on the local end, execute **link-aggregation bfd ipv4 source 2.2.2.2 destination 1.1.1.1** on the peer end.
- The BFD parameters configured on an aggregate interface take effect on all BFD sessions in the aggregation group. BFD sessions for link aggregation do not support the echo packet mode and the Demand mode.

- As a best practice, do not configure other protocols to collaborate with BFD on a BFD-enabled aggregate interface.

Configuration procedure

To enable BFD for an aggregation group:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	<ul style="list-style-type: none"> • Enter Layer 2 aggregate interface view: interface bridge-aggregation <i>interface-number</i> • Enter Layer 3 aggregate interface view: interface route-aggregation <i>interface-number</i> 	N/A
3. Enable BFD for the aggregation group.	link-aggregation bfd ipv4 source <i>ip-address</i> destination <i>ip-address</i>	By default, BFD is disabled for an aggregation group.

Shutting down an aggregate interface

Make sure no member port in an aggregation group is configured with the **loopback** command when you shut down the aggregate interface. Similarly, a port configured with the **loopback** command cannot be assigned to an aggregate interface already shut down. For more information about the **loopback** command, see *Layer 2—LAN Switching Command Reference*.

Shutting down or bringing up an aggregate interface affects the aggregation state and link state of ports in the corresponding aggregation group in the following ways:

- When an aggregate interface is shut down, all Selected ports in the corresponding aggregation group become unselected and their link state becomes down.
- When an aggregate interface is brought up, the aggregation state of ports in the corresponding aggregation group is recalculated.

To shut down an aggregate interface:

Step	Command
1. Enter system view.	system-view
2. Enter aggregate interface view.	<ul style="list-style-type: none"> • Enter Layer 2 aggregate interface view: interface bridge-aggregation <i>interface-number</i> • Enter Layer 3 aggregate interface view: interface route-aggregation <i>interface-number</i>
3. Shut down the aggregate interface.	shutdown

Restoring the default settings for an aggregate interface

You can return all configurations on an aggregate interface to default settings.

To restore the default settings for an aggregate interface:

Step	Command
1. Enter system view.	system-view
2. Enter aggregate interface view.	<ul style="list-style-type: none"> Enter Layer 2 aggregate interface view: interface bridge-aggregation <i>interface-number</i> Enter Layer 3 aggregate interface view: interface route-aggregation <i>interface-number</i>
3. Restore the default settings for the aggregate interface.	default

Configuring load sharing for link aggregation groups

Setting load sharing modes for link aggregation groups

You can set the global or group-specific load sharing mode. A link aggregation group preferentially uses the group-specific load sharing mode. If the group-specific load sharing mode is not available, the group uses the global load sharing mode.

Setting the global link-aggregation load sharing mode

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the global link-aggregation load sharing mode.	link-aggregation global load-sharing mode { destination-ip destination-mac destination-port ingress-port source-ip source-mac source-port } *	By default, the system automatically chooses the global link-aggregation load sharing mode according to the packet type.

Setting the group-specific load sharing mode

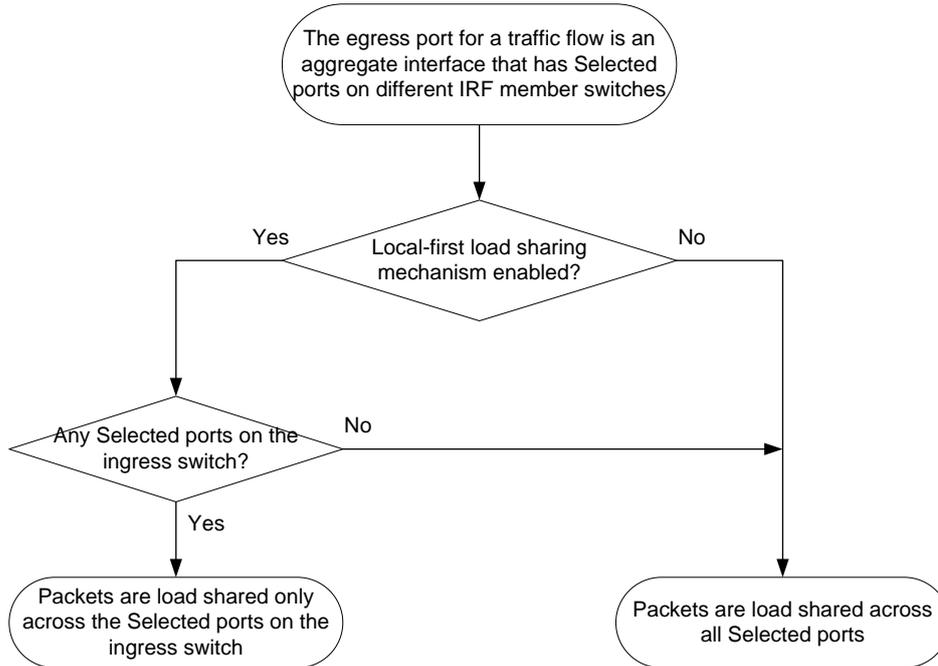
Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	<ul style="list-style-type: none"> Enter Layer 2 aggregate interface view interface bridge-aggregation <i>interface-number</i> Enter Layer 3 aggregate interface view: interface route-aggregation <i>interface-number</i> 	N/A
3. Configure the load sharing mode for the aggregation group.	link-aggregation load-sharing mode { destination-ip destination-mac source-ip source-mac } *	The default settings are the same as the global load sharing mode.

Enabling local-first load sharing for link aggregation

Use local-first load sharing in a multidevice link aggregation scenario to distribute traffic preferentially across member ports on the ingress card or device.

When you aggregate ports on different member devices in an IRF fabric, you can use local-first load sharing to reduce traffic on IRF links, as shown in [Figure 11](#). For more information about IRF, see *IRF Configuration Guide*.

Figure 11 Load sharing for multiswitch link aggregation in an IRF fabric



To enable local-first load sharing for link aggregation:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable local-first load sharing for link aggregation.	link-aggregation load-sharing mode local-first	By default, local-first load sharing for link aggregation is enabled.

Configuring per-flow load sharing algorithm settings for Ethernet link aggregation

Configure the per-flow load sharing algorithm and algorithm seed to optimize traffic distribution on aggregate links based on existing per-flow load sharing settings. You can set only the algorithm or the algorithm seed, or both. You can combine an algorithm with different algorithm seeds to obtain different effects.

This feature does not take effect on per-flow load sharing that uses the following traffic classification criteria:

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

- Destination MAC address.
- Source and destination IP addresses.
- Source and destination MAC addresses.

To configure per-flow load sharing algorithm settings for Ethernet link aggregation:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the load sharing algorithm.	link-aggregation global load-sharing algorithm <i>algorithm-number</i>	By default, algorithm 0 is used.
3. Configure the load sharing algorithm seed.	link-aggregation global load-sharing seed <i>seed-number</i>	By default, algorithm seed 0 is used.

Enabling link-aggregation traffic redirection

ⓘ IMPORTANT:

In Release 1111, this feature is supported only in system view. In Release 1121 and later, this feature is supported in system view, Layer 2 aggregate interface view, and Layer 3 aggregate interface view.

Link-aggregation traffic redirection prevents traffic interruption.

When you shut down a Selected port in an aggregation group, this feature redirects traffic to other Selected ports.

When you restart an IRF member device that contains Selected ports, this feature redirects traffic to other IRF member devices.

You can enable link-aggregation traffic redirection globally or for an aggregation group. Global link-aggregation traffic redirection settings take effect on all aggregation groups. A link aggregation group preferentially uses the group-specific link-aggregation traffic redirection settings. If group-specific link-aggregation traffic redirection is not configured, the group uses the global link-aggregation traffic redirection settings.

Configuration restrictions and guidelines

When you enable link-aggregation traffic redirection, follow these restrictions and guidelines:

- Link-aggregation traffic redirection applies only to dynamic link aggregation groups and takes effect on only known unicast packets.
- To prevent traffic interruption, enable link-aggregation traffic redirection on devices at both ends of the aggregate link.
- To prevent packet loss that might occur at a reboot, do not enable spanning tree together with link-aggregation traffic redirection.
- Link-aggregation traffic redirection does not operate correctly on an edge aggregate interface.
- As a best practice, enable link-aggregation traffic redirection on aggregate interfaces. If you enable this feature globally, communication with a third-party peer device might be affected if the peer is not compatible with this feature.

Configuration procedure

To enable link-aggregation traffic redirection globally:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable link-aggregation traffic redirection globally.	link-aggregation lacp traffic-redirect-notification enable	By default, link-aggregation traffic redirection is disabled globally.

To enable link-aggregation traffic redirection for an aggregation group:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	<ul style="list-style-type: none">Enter Layer 2 aggregate interface view: interface bridge-aggregation interface-numberEnter Layer 3 aggregate interface view: interface route-aggregation interface-number	N/A
3. Enable link-aggregation traffic redirection for the aggregation group.	link-aggregation lacp traffic-redirect-notification enable	By default, link-aggregation traffic redirection is disabled for an aggregation group.

Specifying link aggregation management VLANs and management port

ⓘ IMPORTANT:

This feature is available in Release 1122P02 and later.

For an aggregation group to forward Layer 3 data traffic of some VLANs through a specific port, specify the VLANs as management VLANs and the port as a management port.

To specify link aggregation management VLANs and management port for an aggregation group:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Specify link aggregation management VLANs.	link-aggregation management-vlan <i>vlan-id-list</i>	By default, no link aggregation management VLANs are specified.
3. Enter Layer 2 Ethernet interface view.	interface <i>interface-type</i> interface-number	N/A
4. Configure the port as a management port for its aggregation group.	link-aggregation management-port	By default, a port does not act as a management port in its aggregation group.

Displaying and maintaining Ethernet link aggregation

Execute **display** commands in any view and **reset** commands in user view.

Task	Command
Display information for an aggregate interface or multiple aggregate interfaces.	display interface [bridge-aggregation route-aggregation] [brief [down]] display interface [{ bridge-aggregation route-aggregation } [<i>interface-number</i>]] [brief [description]]
Display the local system ID.	display lacp system-id
Display the global or group-specific link-aggregation load sharing modes.	display link-aggregation load-sharing mode [interface [{ bridge-aggregation route-aggregation } <i>interface-number</i>]]
Display detailed link aggregation information for link aggregation member ports.	display link-aggregation member-port [<i>interface-list</i>]
Display summary information about all aggregation groups.	display link-aggregation summary
Display detailed information about the specified aggregation groups.	display link-aggregation verbose [{ bridge-aggregation route-aggregation } [<i>interface-number</i>]]
Clear LACP statistics for the specified link aggregation member ports.	reset lacp statistics [interface <i>interface-list</i>]
Clear statistics for the specified aggregate interfaces.	reset counters interface [{ bridge-aggregation route-aggregation } [<i>interface-number</i>]]

Ethernet link aggregation configuration examples

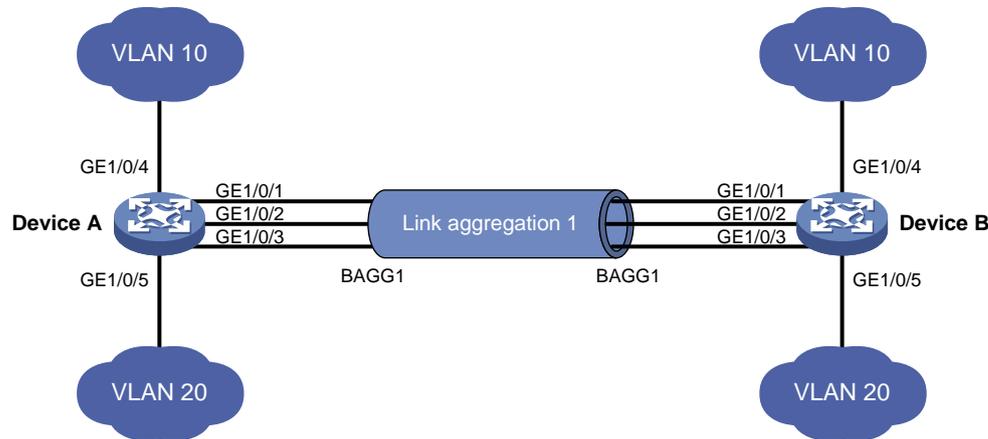
Layer 2 static aggregation configuration example

Network requirements

On the network shown in [Figure 12](#), perform the following tasks:

- Configure a Layer 2 static aggregation group on both Device A and Device B.
- Enable VLAN 10 at one end of the aggregate link to communicate with VLAN 10 at the other end.
- Enable VLAN 20 at one end of the aggregate link to communicate with VLAN 20 at the other end.

Figure 12 Network diagram



Configuration procedure

1. Configure Device A:

Create VLAN 10, and assign port GigabitEthernet 1/0/4 to VLAN 10.

```
<DeviceA> system-view
[DeviceA] vlan 10
[DeviceA-vlan10] port gigabitethernet 1/0/4
[DeviceA-vlan10] quit
```

Create VLAN 20, and assign port GigabitEthernet 1/0/5 to VLAN 20.

```
[DeviceA] vlan 20
[DeviceA-vlan20] port gigabitethernet 1/0/5
[DeviceA-vlan20] quit
```

Create Layer 2 aggregate interface Bridge-Aggregation 1.

```
[DeviceA] interface bridge-aggregation 1
[DeviceA-Bridge-Aggregation1] quit
```

Assign ports GigabitEthernet 1/0/1 through GigabitEthernet 1/0/3 to link aggregation group 1.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/1] quit
[DeviceA] interface gigabitethernet 1/0/2
[DeviceA-GigabitEthernet1/0/2] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/2] quit
[DeviceA] interface gigabitethernet 1/0/3
[DeviceA-GigabitEthernet1/0/3] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/3] quit
```

Configure Layer 2 aggregate interface Bridge-Aggregation 1 as a trunk port and assign it to VLANs 10 and 20.

```
[DeviceA] interface bridge-aggregation 1
[DeviceA-Bridge-Aggregation1] port link-type trunk
[DeviceA-Bridge-Aggregation1] port trunk permit vlan 10 20
[DeviceA-Bridge-Aggregation1] quit
```

2. Configure Device B in the same way Device A is configured. (Details not shown.)

Verifying the configuration

Display detailed information about all aggregation groups on Device A.

```
[DeviceA] display link-aggregation verbose
Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Port Status: S -- Selected, U -- Unselected, I -- Individual
Flags: A -- LACP_Activity, B -- LACP_Timeout, C -- Aggregation,
      D -- Synchronization, E -- Collecting, F -- Distributing,
      G -- Defaulted, H -- Expired
```

```
Aggregate Interface: Bridge-Aggregation1
```

```
Aggregation Mode: Static
```

```
Loadsharing Type: Shar
```

```
Port          Status  Priority Oper-Key
```

Port	Status	Priority	Oper-Key
GE1/0/1	S	32768	1
GE1/0/2	S	32768	1
GE1/0/3	S	32768	1

The output shows that:

- Link aggregation group 1 is a Layer 2 static aggregation group.
- The aggregation group contains three Selected ports.

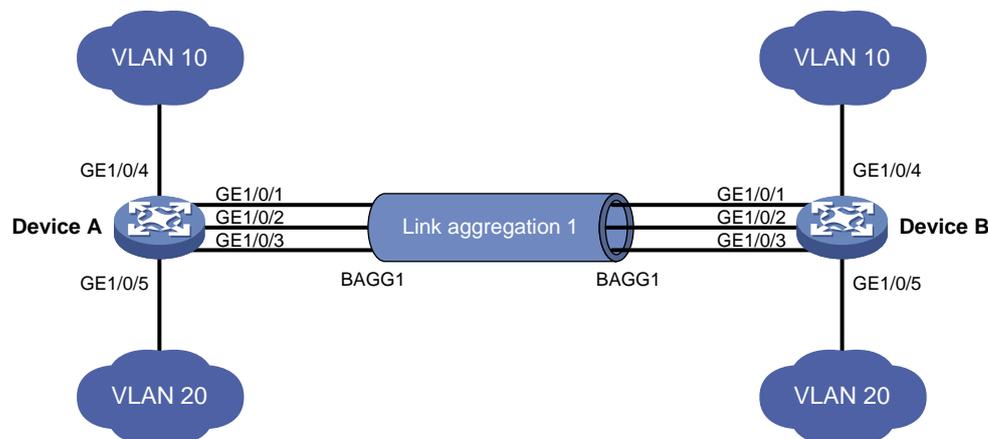
Layer 2 dynamic aggregation configuration example

Network requirements

On the network shown in [Figure 13](#), perform the following tasks:

- Configure a Layer 2 dynamic aggregation group on both Device A and Device B.
- Enable VLAN 10 at one end of the aggregate link to communicate with VLAN 10 at the other end.
- Enable VLAN 20 at one end of the aggregate link to communicate with VLAN 20 at the other end.

Figure 13 Network diagram



Configuration procedure

1. Configure Device A:


```
# Create VLAN 10, and assign the port GigabitEthernet 1/0/4 to VLAN 10.
<DeviceA> system-view
[DeviceA] vlan 10
```

```

[DeviceA-vlan10] port gigabitethernet 1/0/4
[DeviceA-vlan10] quit
# Create VLAN 20, and assign the port GigabitEthernet 1/0/5 to VLAN 20.
[DeviceA] vlan 20
[DeviceA-vlan20] port gigabitethernet 1/0/5
[DeviceA-vlan20] quit
# Create Layer 2 aggregate interface Bridge-Aggregation 1, and set the link aggregation mode
to dynamic.
[DeviceA] interface bridge-aggregation 1
[DeviceA-Bridge-Aggregation1] link-aggregation mode dynamic
[DeviceA-Bridge-Aggregation1] quit
# Assign ports GigabitEthernet 1/0/1 through GigabitEthernet 1/0/3 to link aggregation group 1.
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/1] quit
[DeviceA] interface gigabitethernet 1/0/2
[DeviceA-GigabitEthernet1/0/2] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/2] quit
[DeviceA] interface gigabitethernet 1/0/3
[DeviceA-GigabitEthernet1/0/3] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/3] quit
# Configure Layer 2 aggregate interface Bridge-Aggregation 1 as a trunk port and assign it to
VLANs 10 and 20.
[DeviceA] interface bridge-aggregation 1
[DeviceA-Bridge-Aggregation1] port link-type trunk
[DeviceA-Bridge-Aggregation1] port trunk permit vlan 10 20
[DeviceA-Bridge-Aggregation1] quit

```

2. Configure Device B in the same way Device A is configured. (Details not shown.)

Verifying the configuration

Display detailed information about all aggregation groups on Device A.

```

[DeviceA] display link-aggregation verbose
Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Port Status: S -- Selected, U -- Unselected, I -- Individual
Flags: A -- LACP_Activity, B -- LACP_Timeout, C -- Aggregation,
       D -- Synchronization, E -- Collecting, F -- Distributing,
       G -- Defaulted, H -- Expired

```

Aggregate Interface: Bridge-Aggregation1

Aggregation Mode: Dynamic

Loadsharing Type: Shar

System ID: 0x8000, 000f-e267-6c6a

Local:

Port	Status	Priority	Oper-Key	Flag
GE1/0/1	S	32768	1	{ACDEF}
GE1/0/2	S	32768	1	{ACDEF}
GE1/0/3	S	32768	1	{ACDEF}

Remote:

Actor	Partner	Priority	Oper-Key	SystemID	Flag
GE1/0/1	1	32768	1	0x8000, 000f-e267-57ad	{ACDEF}
GE1/0/2	2	32768	1	0x8000, 000f-e267-57ad	{ACDEF}
GE1/0/3	3	32768	1	0x8000, 000f-e267-57ad	{ACDEF}

The output shows that:

- Link aggregation group 1 is a Layer 2 dynamic aggregation group.
- The aggregation group contains three Selected ports.

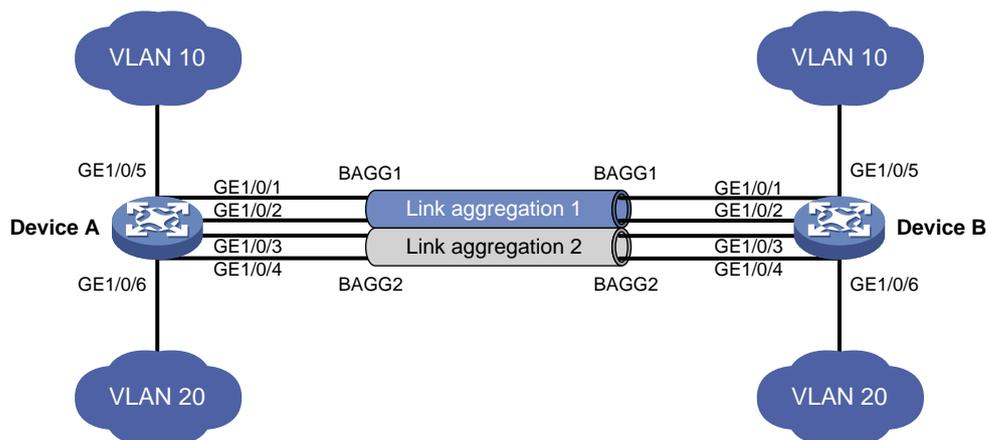
Layer 2 aggregation load sharing configuration example

Network requirements

On the network shown in [Figure 14](#), perform the following tasks:

- Configure two Layer 2 static aggregation groups 1 and 2 on Device A and Device B, respectively.
- Enable VLAN 10 at one end of the aggregate link to communicate with VLAN 10 at the other end.
- Enable VLAN 20 at one end of the aggregate link to communicate with VLAN 20 at the other end.
- Configure link aggregation groups 1 and 2 to load share traffic across aggregation group member ports.
 - Configure link aggregation group 1 to load share packets based on source MAC addresses.
 - Configure link aggregation group 2 to load share packets based on destination MAC addresses.

Figure 14 Network diagram



Configuration procedure

1. Configure Device A:

Create VLAN 10, and assign the port GigabitEthernet 1/0/5 to VLAN 10.

```
<DeviceA> system-view
[DeviceA] vlan 10
[DeviceA-vlan10] port gigabitethernet 1/0/5
[DeviceA-vlan10] quit
```

Create VLAN 20, and assign the port GigabitEthernet 1/0/6 to VLAN 20.

```
[DeviceA] vlan 20
```

```

[DeviceA-vlan20] port gigabitethernet 1/0/6
[DeviceA-vlan20] quit
# Create Layer 2 aggregate interface Bridge-Aggregation 1.
[DeviceA] interface bridge-aggregation 1
# Configure Layer 2 aggregation group 1 to load share packets based on source MAC
addresses.
[DeviceA-Bridge-Aggregation1] link-aggregation load-sharing mode source-mac
[DeviceA-Bridge-Aggregation1] quit
# Assign ports GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to link aggregation group 1.
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/1] quit
[DeviceA] interface gigabitethernet 1/0/2
[DeviceA-GigabitEthernet1/0/2] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/2] quit
# Configure Layer 2 aggregate interface Bridge-Aggregation 1 as a trunk port and assign it to
VLAN 10.
[DeviceA] interface bridge-aggregation 1
[DeviceA-Bridge-Aggregation1] port link-type trunk
[DeviceA-Bridge-Aggregation1] port trunk permit vlan 10
[DeviceA-Bridge-Aggregation1] quit
# Create Layer 2 aggregate interface Bridge-Aggregation 2.
[DeviceA] interface bridge-aggregation 2
# Configure Layer 2 aggregation group 2 to load share packets based on destination MAC
addresses.
[DeviceA-Bridge-Aggregation2] link-aggregation load-sharing mode destination-mac
[DeviceA-Bridge-Aggregation2] quit
# Assign ports GigabitEthernet 1/0/3 and GigabitEthernet 1/0/4 to link aggregation group 2.
[DeviceA] interface gigabitethernet 1/0/3
[DeviceA-GigabitEthernet1/0/3] port link-aggregation group 2
[DeviceA-GigabitEthernet1/0/3] quit
[DeviceA] interface gigabitethernet 1/0/4
[DeviceA-GigabitEthernet1/0/4] port link-aggregation group 2
[DeviceA-GigabitEthernet1/0/4] quit
# Configure Layer 2 aggregate interface Bridge-Aggregation 2 as a trunk port and assign it to
VLAN 20.
[DeviceA] interface bridge-aggregation 2
[DeviceA-Bridge-Aggregation2] port link-type trunk
[DeviceA-Bridge-Aggregation2] port trunk permit vlan 20
[DeviceA-Bridge-Aggregation2] quit

```

2. Configure Device B in the same way Device A is configured. (Details not shown.)

Verifying the configuration

```

# Display detailed information about all aggregation groups on Device A.
[DeviceA] display link-aggregation verbose
Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Port Status: S -- Selected, U -- Unselected, I -- Individual
Flags: A -- LACP_Activity, B -- LACP_Timeout, C -- Aggregation,
      D -- Synchronization, E -- Collecting, F -- Distributing,

```

G -- Defaulted, H -- Expired

Aggregate Interface: Bridge-Aggregation1

Aggregation Mode: Static

Loadsharing Type: Shar

Port	Status	Priority	Oper-Key
GE1/0/1	S	32768	1
GE1/0/2	S	32768	1

Aggregate Interface: Bridge-Aggregation2

Aggregation Mode: Static

Loadsharing Type: Shar

Port	Status	Priority	Oper-Key
GE1/0/3	S	32768	2
GE1/0/4	S	32768	2

The output shows that:

- Link aggregation groups 1 and 2 are both load-shared Layer 2 static aggregation groups.
- Each aggregation group contains two Selected ports.

Display all the group-specific load sharing modes on Device A.

[DeviceA] display link-aggregation load-sharing mode interface

Bridge-Aggregation1 Load-Sharing Mode:

source-mac address

Bridge-Aggregation2 Load-Sharing Mode:

destination-mac address

The output shows that:

- Link aggregation group 1 load shares packets based on source MAC addresses.
- Link aggregation group 2 load shares packets based on destination MAC addresses.

Layer 2 edge aggregate interface configuration example

Network requirements

As shown in [Figure 15](#), a Layer 2 dynamic aggregation group is configured on the device. The server is not configured with dynamic link aggregation.

Configure an edge aggregate interface so that both GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 can forward traffic to improve link reliability.

Figure 15 Network diagram



Configuration procedure

Create Layer 2 aggregate interface Bridge-Aggregation 1, and set the link aggregation mode to dynamic.

```
<Device> system-view
[Device] interface bridge-aggregation 1
[Device-Bridge-Aggregation1] link-aggregation mode dynamic
```

Configure Layer 2 aggregate interface Bridge-Aggregation 1 as an edge aggregate interface.

```
[Device-Bridge-Aggregation1] lacp edge-port
[Device-Bridge-Aggregation1] quit
```

Assign ports GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to link aggregation group 1.

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

Verifying the configuration

Display detailed information about all aggregation groups on the device when the server is not configured with dynamic link aggregation.

```
[Device] display link-aggregation verbose
Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Port Status: S -- Selected, U -- Unselected, I -- Individual
Flags:  A -- LACP_Activity, B -- LACP_Timeout, C -- Aggregation,
        D -- Synchronization, E -- Collecting, F -- Distributing,
        G -- Defaulted, H -- Expired
```

Aggregate Interface: Bridge-Aggregation1

Aggregation Mode: Dynamic

Loadsharing Type: NonS

System ID: 0x8000, 000f-e267-6c6a

Local:

Port	Status	Priority	Oper-Key	Flag
GE1/0/1	I	32768	1	{AG}
GE1/0/2	I	32768	1	{AG}

Remote:

Actor	Partner	Priority	Oper-Key	SystemID	Flag
GE1/0/1	0	32768	0	0x8000, 0000-0000-0000	{DEF}
GE1/0/2	0	32768	0	0x8000, 0000-0000-0000	{DEF}

The output shows that GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 are in Individual state when they do not receive LACPDUs from the server. Both GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 can forward traffic. When one port fails, its traffic is automatically switched to the other port.

Layer 3 static aggregation configuration example

Network requirements

On the network shown in [Figure 16](#), perform the following tasks:

- Configure a Layer 3 static aggregation group on both Device A and Device B.
- Configure IP addresses and subnet masks for the corresponding Layer 3 aggregate interfaces.

Figure 16 Network diagram



Configuration procedure

1. Configure Device A:

Create Layer 3 aggregate interface Route-Aggregation 1, and configure an IP address and subnet mask for the aggregate interface.

```
<DeviceA> system-view
[DeviceA] interface route-aggregation 1
[DeviceA-Route-Aggregation1] ip address 192.168.1.1 24
[DeviceA-Route-Aggregation1] quit
```

Assign Layer 3 Ethernet interfaces GigabitEthernet 1/0/1 through GigabitEthernet 1/0/3 to aggregation group 1.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/1] quit
[DeviceA] interface gigabitethernet 1/0/2
[DeviceA-GigabitEthernet1/0/2] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/2] quit
[DeviceA] interface gigabitethernet 1/0/3
[DeviceA-GigabitEthernet1/0/3] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/3] quit
```

2. Configure Device B in the same way Device A is configured. (Details not shown.)

Verifying the configuration

Display detailed information about all aggregation groups on Device A.

```
[DeviceA] display link-aggregation verbose
Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Port Status: S -- Selected, U -- Unselected, I -- Individual
Flags: A -- LACP_Activity, B -- LACP_Timeout, C -- Aggregation,
       D -- Synchronization, E -- Collecting, F -- Distributing,
       G -- Defaulted, H -- Expired
```

```
Aggregate Interface: Route-Aggregation1
Aggregation Mode: Static
Loadsharing Type: Shar
  Port          Status  Priority Oper-Key
```

```

GE1/0/1          S          32768    1
GE1/0/2          S          32768    1
GE1/0/3          S          32768    1

```

The output shows that link aggregation group 1 is a non-load-shared Layer 3 static aggregation group that contains three Selected ports.

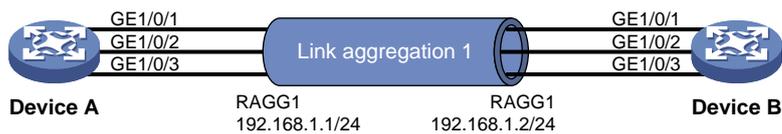
Layer 3 dynamic aggregation configuration example

Network requirements

On the network shown in [Figure 17](#), perform the following tasks:

- Configure a Layer 3 dynamic aggregation group on both Device A and Device B.
- Configure IP addresses and subnet masks for the corresponding Layer 3 aggregate interfaces.

Figure 17 Network diagram



Configuration procedure

1. Configure Device A:

Create Layer 3 aggregate interface Route-Aggregation 1.

```

<DeviceA> system-view
[DeviceA] interface route-aggregation 1

```

Set the link aggregation mode to dynamic.

```

[DeviceA-Route-Aggregation1] link-aggregation mode dynamic

```

Configure an IP address and subnet mask for Route-Aggregation 1.

```

[DeviceA-Route-Aggregation1] ip address 192.168.1.1 24
[DeviceA-Route-Aggregation1] quit

```

Assign Layer 3 Ethernet interfaces GigabitEthernet 1/0/1 through GigabitEthernet 1/0/3 to aggregation group 1.

```

[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/1] quit
[DeviceA] interface gigabitethernet 1/0/2
[DeviceA-GigabitEthernet1/0/2] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/2] quit
[DeviceA] interface gigabitethernet 1/0/3
[DeviceA-GigabitEthernet1/0/3] port link-aggregation group 1
[DeviceA-GigabitEthernet1/0/3] quit

```

2. Configure Device B in the same way Device A is configured. (Details not shown.)

Verifying the configuration

Display detailed information about all aggregation groups on Device A.

```

[DeviceA] display link-aggregation verbose
Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Port Status: S -- Selected, U -- Unselected, I -- Individual
Flags: A -- LACP_Activity, B -- LACP_Timeout, C -- Aggregation,

```

D -- Synchronization, E -- Collecting, F -- Distributing,
 G -- Defaulted, H -- Expired

Aggregate Interface: Route-Aggregation1

Aggregation Mode: Dynamic

Loadsharing Type: Shar

System ID: 0x8000, 000f-e267-6c6a

Local:

Port	Status	Priority	Oper-Key	Flag
GE1/0/1	S	32768	1	{ACDEF}
GE1/0/2	S	32768	1	{ACDEF}
GE1/0/3	S	32768	1	{ACDEF}

Remote:

Actor	Partner	Priority	Oper-Key	SystemID	Flag
GE1/0/1	1	32768	1	0x8000, 000f-e267-57ad	{ACDEF}
GE1/0/2	2	32768	1	0x8000, 000f-e267-57ad	{ACDEF}
GE1/0/3	3	32768	1	0x8000, 000f-e267-57ad	{ACDEF}

The output shows that:

- Link aggregation group 1 is a non-load-shared Layer 3 dynamic aggregation group.
- The aggregation group contains three Selected ports.

Layer 3 edge aggregate interface configuration example

Network requirements

As shown in [Figure 18](#), a Layer 3 dynamic aggregation group is configured on the device. The server is not configured with dynamic link aggregation.

Configure an edge aggregate interface so that both GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 can forward traffic to improve link reliability.

Figure 18 Network diagram



Configuration procedure

Create Layer 3 aggregate interface Route-Aggregation 1, and set the link aggregation mode to dynamic.

```
<Device> system-view
[Device] interface route-aggregation 1
[Device-Route-Aggregation1] link-aggregation mode dynamic
```

Configure an IP address and subnet mask for Layer 3 aggregate interface Route-Aggregation 1.

```
[Device-Route-Aggregation1] ip address 192.168.1.1 24
```

Configure Layer 3 aggregate interface Route-Aggregation 1 as an edge aggregate interface.

```
[Device-Route-Aggregation1] lacp edge-port
```

```
[Device-Route-Aggregation1] quit
```

Assign Layer 3 Ethernet interfaces GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to aggregation group 1.

```
[Device] interface gigabitethernet 1/0/1
```

```
[Device-GigabitEthernet1/0/1] port link-aggregation group 1
```

```
[Device-GigabitEthernet1/0/1] quit
```

```
[Device] interface gigabitethernet 1/0/2
```

```
[Device-GigabitEthernet1/0/2] port link-aggregation group 1
```

```
[Device-GigabitEthernet1/0/2] quit
```

Verifying the configuration

Display detailed information about all aggregation groups on the device when the server is not configured with dynamic link aggregation.

```
[Device] display link-aggregation verbose
```

```
Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
```

```
Port Status: S -- Selected, U -- Unselected, I -- Individual
```

```
Flags: A -- LACP_Activity, B -- LACP_Timeout, C -- Aggregation,  
D -- Synchronization, E -- Collecting, F -- Distributing,  
G -- Defaulted, H -- Expired
```

```
Aggregate Interface: Route-Aggregation1
```

```
Aggregation Mode: Dynamic
```

```
Loadsharing Type: NonS
```

```
System ID: 0x8000, 000f-e267-6c6a
```

```
Local:
```

Port	Status	Priority	Oper-Key	Flag
GE1/0/1	I	32768	1	{AG}
GE1/0/2	I	32768	1	{AG}

```
Remote:
```

Actor	Partner	Priority	Oper-Key	SystemID	Flag
GE1/0/1	0	32768	0	0x8000, 0000-0000-0000	{DEF}
GE1/0/2	0	32768	0	0x8000, 0000-0000-0000	{DEF}

The output shows that GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 are in Individual state when they do not receive LACPDUs from the server. Both GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 can forward traffic. When one port fails, its traffic is automatically switched to the other port.