

Contents

Configuring Ethernet link aggregation	1
About Ethernet link aggregation	1
Ethernet link aggregation application scenario	1
Aggregate interface, aggregation group, and member port	1
Operational key	2
Configuration types	2
Link aggregation modes	2
How static link aggregation works	3
Dynamic link aggregation	4
How dynamic link aggregation works	6
Edge aggregate interface	9
Load sharing modes for link aggregation groups	9
Ethernet link aggregation tasks at a glance	9
Configuring an aggregation group	10
Restrictions and guidelines for aggregation group configuration	10
Configuring a Layer 2 aggregation group	11
Configuring an aggregate interface	12
Setting the minimum and maximum numbers of Selected ports for an aggregation group	12
Configuring the description of an aggregate interface	13
Configuring jumbo frame support	13
Setting the expected bandwidth for an aggregate interface	14
Configuring an edge aggregate interface	14
Shutting down an aggregate interface	15
Restoring the default settings for an aggregate interface	15
Configuring a dynamic aggregation group to use port speed as the prioritized criterion for reference port selection	16
Configuring load sharing for link aggregation groups	16
Setting load sharing modes for link aggregation groups	16
Enabling local-first load sharing for link aggregation	17
Enabling a Layer 2 aggregate interface to reflect incoming packets back	18
Enabling BFD for an aggregation group	19
Enabling link-aggregation traffic redirection	20
About link-aggregation traffic redirection	20
Restrictions and guidelines for link-aggregation traffic redirection	20
Enabling link-aggregation traffic redirection globally	20
Enabling link-aggregation traffic redirection for an aggregation group	20
Display and maintenance commands for Ethernet link aggregation	21
Ethernet link aggregation configuration examples	21
Example: Configuring a Layer 2 static aggregation group	21
Example: Configuring a Layer 2 dynamic aggregation group	23
Example: Configuring a Layer 2 edge aggregate interface	25

Configuring Ethernet link aggregation

About Ethernet link aggregation

Ethernet link aggregation bundles multiple physical Ethernet links into one logical link (called an aggregate link). Link aggregation provides the following benefits:

- Increased bandwidth beyond the limits of a single individual 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.

Ethernet link aggregation application scenario

As shown in [Figure 1](#), 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 transmitted on the failed link is switched to the other two links.

Figure 1 Ethernet link aggregation diagram



Aggregate interface, aggregation group, and member port

Each link aggregation is represented by a logical aggregate interface. Each aggregate interface has an automatically created aggregation group, which contains member ports to be used for aggregation. The type and number of an aggregation group are the same as its aggregate interface.

Supported aggregate interface types

A Layer 2 aggregate interface is created manually. The member ports in a Layer 2 aggregation group can only be Layer 2 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 Selected member ports, 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. This state is peculiar to the member ports of edge aggregate interfaces. A Selected or Unselected member port of an edge aggregate interface is placed in Individual state if the following events occur in sequence:

- a. The member port goes down and then comes up.
- b. The LACP timeout timer expires because it has not received LACPDUs.

For more information about edge aggregate interfaces, see "[Edge aggregate interface.](#)"

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 have the same operational key.

Configuration types

Port configuration includes the attribute configuration and protocol configuration. Attribute configuration affects the aggregation state of the port but the protocol configuration does not.

Attribute configuration

To become a Selected port, a member port must have the same attribute configuration as the aggregate interface. [Table 1](#) describes the attribute configuration.

Table 1 Attribute configuration

Feature	Attribute configuration
Port isolation	Membership of the port in an isolation group. Isolation group number.
QinQ	QinQ status (enabled/disabled), TPID for VLAN tags, and VLAN transparent transmission. For information about QinQ, see "Configuring QinQ."
VLAN mapping	VLAN mapping configured on the port. For more information about VLAN mapping, see "Configuring VLAN mapping."
VLAN	VLAN attribute settings: <ul style="list-style-type: none">• Permitted VLAN IDs.• PVID.• Link type (trunk, hybrid, or access).• PVLAN port type (promiscuous, trunk promiscuous, host, or trunk secondary).• IP subnet-based VLAN configuration.• Protocol-based VLAN configuration.• VLAN tagging mode. For information about VLANs, see "Configuring VLANs."

Protocol configuration

Protocol configuration of a member port does not affect the aggregation state of the member port. MAC address learning and spanning tree settings are examples of the protocol configuration.

Link aggregation modes

An aggregation group operates in one of the following modes:

- **Static**—Static aggregation is stable. An aggregation group in static mode is called a static aggregation group. The aggregation states of the member ports in a static aggregation group are not affected by the peer ports.
- **Dynamic**—An aggregation group in dynamic mode is called a dynamic aggregation group. Dynamic aggregation is implemented through IEEE 802.3ad Link Aggregation Control Protocol (LACP). The local system and the peer system automatically maintain the aggregation states of the member ports. Dynamic link aggregation reduces the administrators' workload.

How static link aggregation works

Reference port selection process

When setting the aggregation states 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 in up state.

The candidate reference ports are organized into different priority levels following these rules:

1. In descending order of port priority.
2. Full duplex.
3. In descending order of speed.
4. Half duplex.
5. In descending order of speed.

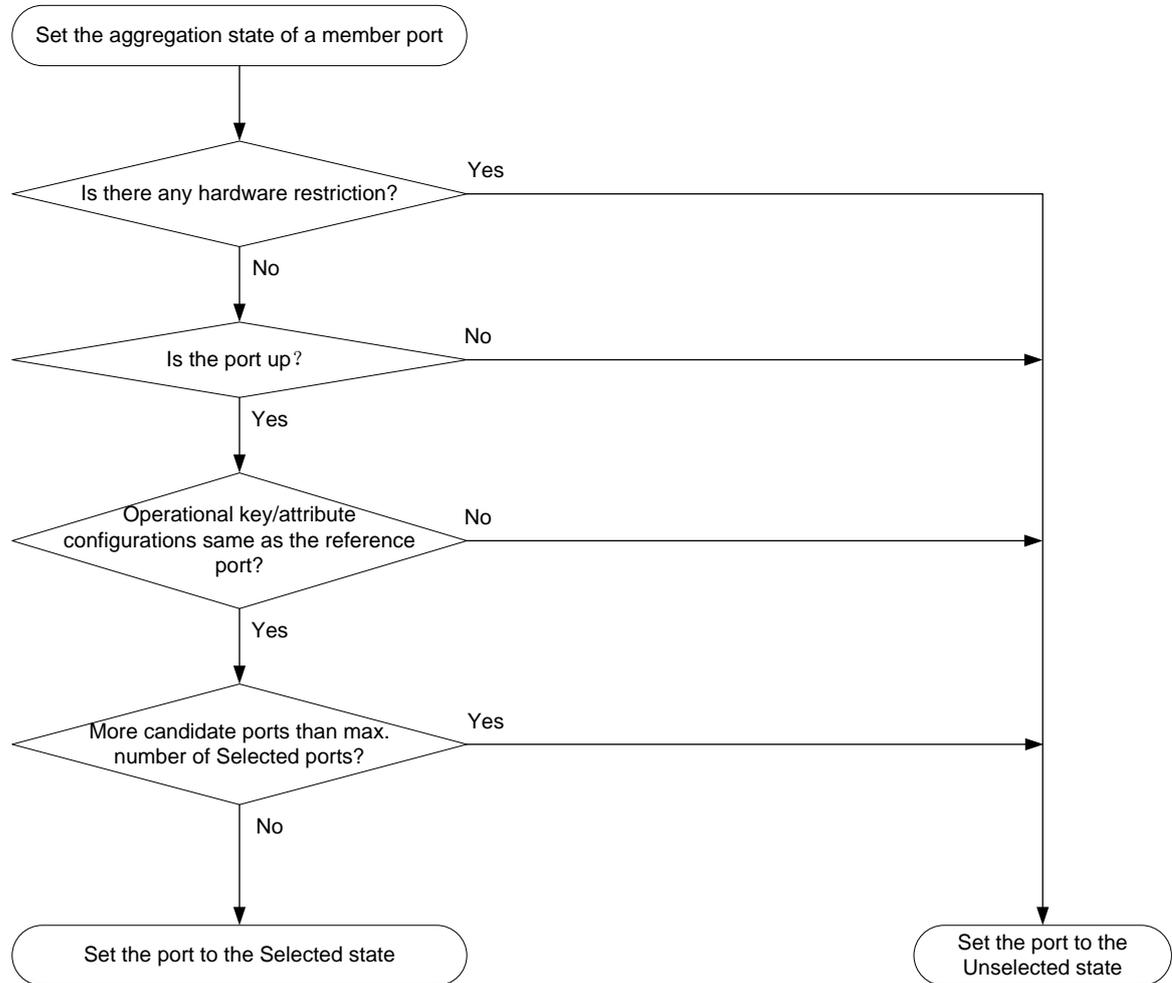
From the candidate ports with the same attribute configurations as the aggregate interface, the one with the highest priority level is chosen as the reference port.

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

Setting the aggregation state of each member port

After the reference port is chosen, the system sets the aggregation state of each member port in the static aggregation group.

Figure 2 Setting the aggregation state of a member port in a static aggregation group



After the limit on Selected ports is reached, the aggregation state of a new member port varies by following conditions:

- The port is placed in Unselected state if the port and the Selected ports have the same port priority. This mechanism prevents traffic interruption on the existing Selected ports. A device reboot can cause the device to recalculate the aggregation states of member ports.
- The port is placed in Selected state when the following conditions are met:
 - The port and the Selected ports have different port priorities, and the port has a higher port priority than a minimum of one Selected port.
 - The port has the same attribute configurations as the aggregate interface.

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

Dynamic link aggregation

About LACP

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

LACP uses LACPDUs to exchange aggregation information between LACP-enabled devices. Each member port in a dynamic aggregation group can exchange 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 Selected state.

LACP functions

LACP offers basic LACP functions and extended LACP functions, as described in [Table 2](#).

Table 2 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. Extended LACP can implement LACP MAD for the IRF feature. The device 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>Virtual Technologies Configuration Guide</i> .

LACP operating modes

LACP can operate in active 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 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 3](#). The smaller the priority value, the higher the priority.

Table 3 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 a Selected port 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 has not received 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.

Methods to assign interfaces to a dynamic link aggregation group

You can use one of the following methods to assign interfaces to a dynamic link aggregation group:

- **Manual assignment**—Manually assign interfaces to the dynamic link aggregation group.
- **Automatic assignment**—Enable automatic assignment on interfaces to have them automatically join a dynamic link aggregation group depending on the peer information in the received LACPDU.

NOTE:

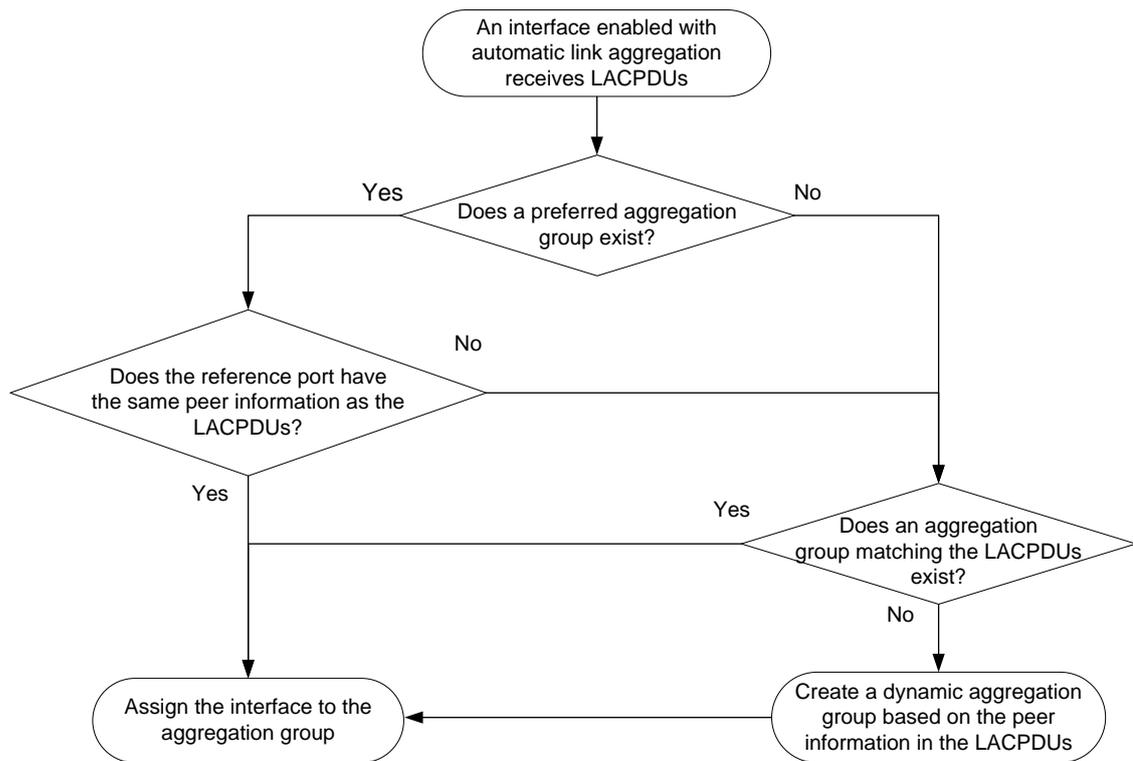
When you use automatic assignment on one end, you must use manual assignment on the other end.

Automatic member port assignment

As shown in [Figure 3](#), an interface enabled with automatic assignment joins a dynamic aggregation group based on the peer information in the LACPDU received from the aggregation peer. If none of the existing dynamic aggregation groups is qualified, the device automatically creates a new dynamic aggregation group. Then, the device assigns the interface to that group and synchronizes the interface's attribute configurations to the aggregate interface.

A dynamic aggregation group that contains automatically assigned member ports selects a reference port and Selected ports as described in "[How dynamic link aggregation works.](#)" The assignment methods of member ports do not change the processes of reference port and Selected port selection.

Figure 3 Automatic member port assignment process



How dynamic link aggregation works

Choosing a reference port

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

The local system (the actor) and the peer 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.

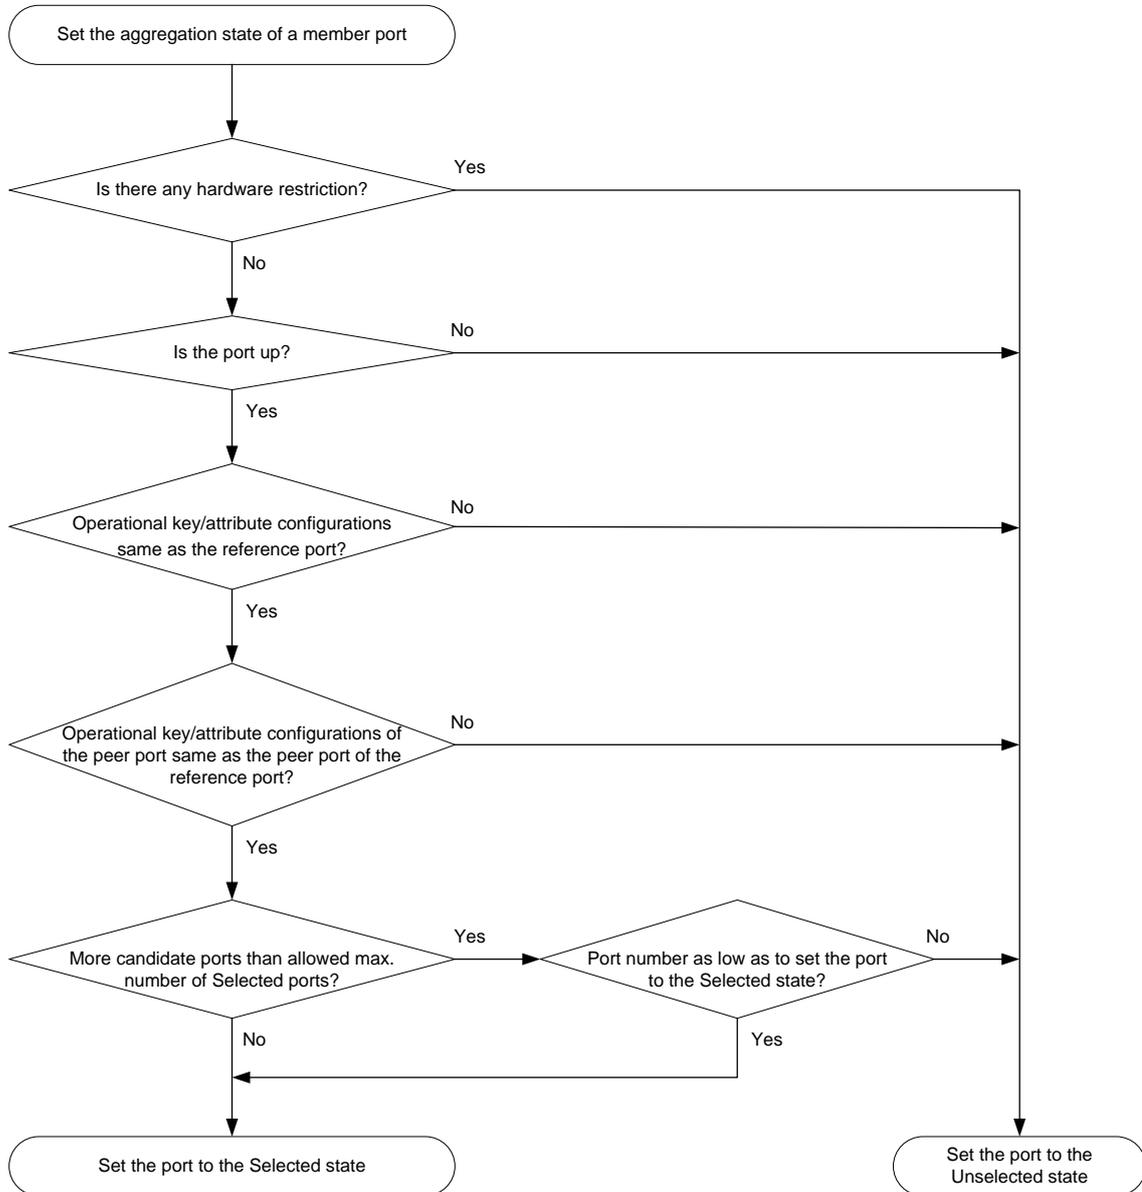
NOTE:

To identify the port numbers of aggregation member ports, execute the **display link-aggregation verbose** command and examine the **Index** field in the command output.

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 4 Setting the state of a member port in a dynamic aggregation group



The system with the greater system ID can detect the aggregation state changes on the peer system. The system with the greater system ID 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 chooses only full-duplex ports as the Selected ports.
- For stable aggregation and service continuity, do not change the operational key or attribute configurations on any member port.
- When a member port changes to the Selected or Unselected state, its peer port changes to the same aggregation state.
- After the Selected port limit is reached, a newly joining port becomes a Selected port if it is more eligible than a current Selected port.

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 LACPDUs 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**—Distributes 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 of or a combination of the following traffic classification criteria:
 - Ingress port.
 - Source or destination IP.
 - Source or destination MAC.
 - Source or destination port number.
- **Packet type-based load sharing**—Distributes traffic automatically based on packet types.

Ethernet link aggregation tasks at a glance

To configure Ethernet link aggregation, perform the following tasks:

1. [Configuring an aggregation group](#)
2. (Optional.) [Configuring an aggregate interface](#)
 - [Setting the minimum and maximum numbers of Selected ports for an aggregation group](#)
 - [Configuring the description of an aggregate interface](#)
 - [Configuring jumbo frame support](#)
 - [Setting the expected bandwidth for an aggregate interface](#)
 - [Configuring an edge aggregate interface](#)

An edge aggregate interface uses all member ports to forward traffic when the aggregation peer is not enabled with dynamic link aggregation.
 - [Shutting down an aggregate interface](#)
 - [Restoring the default settings for an aggregate interface](#)
3. (Optional.) [Configuring a dynamic aggregation group to use port speed as the prioritized criterion for reference port selection](#)
4. (Optional.) [Configuring load sharing for link aggregation groups](#)
 - [Setting load sharing modes for link aggregation groups](#)
 - [Enabling local-first load sharing for link aggregation](#)
5. (Optional.) [Enabling a Layer 2 aggregate interface to reflect incoming packets back](#)

Perform this task to have a Layer 2 aggregate interface reflect a packet back when it is both the incoming and outgoing interfaces of that packet.

6. (Optional.) [Enabling BFD for an aggregation group](#)
7. (Optional.) [Enabling link-aggregation traffic redirection](#)

This feature redirects traffic on an unavailable Selected port to the remaining available Selected ports of an aggregation group to avoid traffic interruption.

Configuring an aggregation group

Restrictions and guidelines for aggregation group configuration

Layer 2 aggregation group restrictions

You cannot assign an interface to a Layer 2 aggregation group if any features in [Table 4](#) are configured on that interface.

Table 4 Features incompatible with Layer 2 aggregation member ports

Feature on the interface	Reference
MAC authentication	MAC authentication in <i>Security Configuration Guide</i>
Port security	Port security in <i>Security Configuration Guide</i>
802.1X	802.1X in <i>Security Configuration Guide</i>

Aggregation member port restrictions

Deleting an aggregate interface also deletes its aggregation group and causes all member ports to leave the aggregation group.

Do not assign a reflector port for port mirroring to an aggregation group. For more information about reflector ports, see *Network Management and Monitoring Configuration Guide*.

Attribute and protocol configuration restrictions

Attribute configuration changes made on an aggregate interface are automatically synchronized to all member ports. If the changes fail to be synchronized to a Selected port, the port might change to the Unselected state. To make the port become Selected again, you can change the attribute configurations on the aggregate interface or on the port. The synchronization failure does not affect the attribute configuration changes made on the aggregate interface. The configurations that have been synchronized from the aggregate interface are retained on the member ports even after the aggregate interface is deleted.

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.

The protocol configurations for an aggregate interface take effect only on the current aggregate interface. The protocol configurations for a member port take effect only when the port leaves its aggregation group.

Configuration consistency requirements

You must configure the same aggregation mode at the two ends of an aggregate link.

- For a successful static aggregation, make sure the ports at both ends of each link are in the same aggregation state.
- For a successful dynamic aggregation:

- Make sure the ports at both ends of a link are assigned to the correct aggregation group. The two ends can automatically negotiate the aggregation state of each member port.
- If you use automatic interface assignment on one end, you must use manual assignment on the other end.

Configuring a Layer 2 aggregation group

Configuring a Layer 2 static aggregation group

1. Enter system view.
system-view
2. Create a Layer 2 aggregate interface and enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
When you create a Layer 2 aggregate interface, the system automatically creates a Layer 2 static aggregation group numbered the same as that interface.
3. Return to system view.
quit
4. Assign an interface to the Layer 2 aggregation group:
 - a. Enter Layer 2 Ethernet interface view.
interface *interface-type interface-number*
 - b. Assign the interface to the Layer 2 aggregation group.
port link-aggregation group *group-id* [**force**]
 Repeat the substeps to assign more interfaces to the aggregation group.
To synchronize the attribute configurations from the aggregate interface when the current interface joins the aggregation group, specify the **force** keyword.
5. (Optional.) Set the port priority of the interface.
link-aggregation port-priority *priority*
The default port priority of an interface is 32768.

Configuring a Layer 2 dynamic aggregation group

1. Enter system view.
system-view
2. Set the system LACP priority.
lacp system-priority *priority*
By default, the system LACP priority is 32768.
Changing the system LACP priority might affect the aggregation states of the ports in a dynamic aggregation group.
3. Create a Layer 2 aggregate interface and enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
When you create a Layer 2 aggregate interface, the system automatically creates a Layer 2 static aggregation group numbered the same as that interface.
4. Configure the aggregation group to operate in dynamic mode.
link-aggregation mode dynamic
By default, an aggregation group operates in static mode.
5. Return to system view.
quit
6. Assign an interface to the Layer 2 aggregation group:

- a. Enter Layer 2 Ethernet interface view.

```
interface interface-type interface-number
```

- b. Assign the interface to the Layer 2 aggregation group or enable automatic assignment on that interface.

```
port link-aggregation group { group-id [ force ] | auto [ group-id ] }
```

Repeat these two substeps to assign more Layer 2 Ethernet interfaces to the aggregation group.

To synchronize the attribute configurations from the aggregate interface when the current interface joins the aggregation group, specify the **force** keyword.

To enable automatic assignment, specify the **auto** keyword.

7. Set the LACP operating mode for the interface.

- o Set the LACP operating mode to passive.

```
lacp mode passive
```

- o Set the LACP operating mode to active.

```
undo lacp mode
```

By default, LACP is operating in active mode.

8. (Optional.) Set the port priority for the interface.

```
link-aggregation port-priority priority
```

The default setting is 32768.

9. (Optional.) Set the short LACP timeout interval (3 seconds) for the interface.

```
lacp period short
```

By default, the long LACP timeout interval (90 seconds) is used by the interface.

Configuring an aggregate interface

Most settings that can be made on Layer 2 Ethernet interfaces can also be made on Layer 2 aggregate interfaces.

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

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

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 states of aggregation member ports and the state of the aggregate interface.

- When the number of member ports eligible to be Selected ports is smaller than the minimum threshold, the following events occur:
 - o The eligible member ports are placed in Unselected state.
 - o The link layer state of the aggregate interface becomes down.
- When the number of member ports eligible to be Selected ports reaches or exceeds the minimum threshold, the following events occur:
 - o The eligible member ports are placed in Selected state.
 - o The link layer state of the aggregate interface becomes 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.

Restrictions and guidelines

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

Procedure

1. Enter system view.
system-view
2. Enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
3. Set the minimum number of Selected ports for the aggregation group.
link-aggregation selected-port minimum *min-number*

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 *max-number*

By default, an aggregation group can have a maximum of 8 Selected ports.

Configuring the description of an aggregate interface

About the aggregate interface description

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

Procedure

1. Enter system view.
system-view
2. Enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
3. Configure the interface description.
description *text*

By default, the description of an interface is *interface-name* **Interface**.

Configuring jumbo frame support

About jumbo frames

An aggregate interface might receive frames larger than 1522 bytes during high-throughput data exchanges, such as file transfers. These frames are called jumbo frames.

How an aggregate interface processes jumbo frames depends on whether jumbo frame support is enabled on the interface.

- If configured to deny jumbo frames, the aggregate interface discards jumbo frames.

- If enabled with jumbo frame support, the aggregate interface performs the following operations:
 - Processes jumbo frames within the allowed length.
 - Discards jumbo frames that exceed the allowed length.

Procedure

1. Enter system view.
system-view
2. Enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
3. Allow jumbo frames.
jumboframe enable [*size*]
By default, an aggregate interface allows jumbo frames with a maximum length of 10240 bytes to pass through.
If you execute this command multiple times, the most recent configuration takes effect.

Setting the expected bandwidth for an aggregate interface

About expected bandwidth

Expected bandwidth is an informational parameter used only by higher-layer protocols for calculation. You cannot adjust the actual bandwidth of an interface by performing this task.

Procedure

1. Enter system view.
system-view
2. Enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
3. Set the expected bandwidth for the interface.
bandwidth *bandwidth-value*
By default, the expected bandwidth (in kbps) is the interface baud rate divided by 1000.

Configuring an edge aggregate interface

Restrictions and guidelines

This configuration takes effect only on aggregate interfaces in dynamic mode.

Configure only terminal-facing aggregate interfaces as edge aggregate interfaces. Do not configure aggregate interfaces that are connected to network devices as edge aggregate interfaces.

Link-aggregation traffic redirection cannot operate correctly on an edge aggregate interface. For more information about link-aggregation traffic redirection, see "[Enabling link-aggregation traffic redirection.](#)"

Procedure

1. Enter system view.
system-view
2. Enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
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.

Shutting down an aggregate interface

Restrictions and guidelines

Shutting down or bringing up an aggregate interface affects the aggregation states and link states of member ports in the corresponding aggregation group as follows:

- When an aggregate interface is shut down, all its Selected ports become Unselected and all member ports go down.
- When an aggregate interface is brought up, the aggregation states of all its member ports are recalculated.

Procedure

1. Enter system view.
system-view
2. Enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
3. Shut down the interface.
shutdown

Restoring the default settings for an aggregate interface

Restrictions and guidelines

△ CAUTION:

The **default** command might interrupt ongoing network services. Make sure you are fully aware of the impacts of this command when you execute it on a live network.

The **default** command might fail to restore the default settings for some commands for reasons such as command dependencies and system restrictions.

To resolve this issue:

1. Use the **display this** command in interface view to identify these commands.
2. Use their **undo** forms or follow the command reference to restore their default settings.
3. If the restoration attempt still fails, follow the error message instructions to resolve the issue.

Procedure

1. Enter system view.
system-view
2. Enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
3. Restore the default settings for the aggregate interface.
default

Configuring a dynamic aggregation group to use port speed as the prioritized criterion for reference port selection

About prioritizing port speed in reference port selection

Perform this task to ensure that a dynamic aggregation group selects a high-speed member port as the reference port. After you perform this task, the reference port will be selected based on the criteria in order of device ID, port speed, and port ID.

Restrictions and guidelines

Changing reference port selection criteria might cause transient traffic interruption. Make sure you understand the impact of this task on your network.

You must perform this task at both ends of the aggregate link so the peer aggregation systems use the same criteria for reference port selection.

As a best practice, shut down the peer aggregate interfaces before you execute this command and bring up the interfaces after this command is executed on both of them.

Procedure

1. Enter system view.
system-view
2. Enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
3. Specify port speed as the prioritized criterion for reference port selection.
lacp select speed

By default, port ID is the prioritized criterion for reference port selection of a dynamic aggregation group.

Configuring load sharing for link aggregation groups

Setting load sharing modes for link aggregation groups

About load sharing modes

ⓘ IMPORTANT:

Configuring link aggregation load sharing might cause equal-cost routes to load share traffic in an unbalanced manner.

On the device, you can only set a global load sharing mode, which takes effect on all aggregation groups.

Restrictions and guidelines

The following are global load sharing modes supported on the device:

- Load sharing mode automatically determined based on the packet type.
- Source IP.

- Destination IP.
- Source MAC.
- Destination MAC.
- Ingress port.
- Source IP and destination IP.
- Source IP and source port.
- Destination IP and destination port.
- Source IP, source port, destination IP, and destination port.
- Any combinations of ingress port, source MAC, and destination MAC.

Setting the global link-aggregation load sharing mode

1. Enter system view.
`system-view`
2. Set 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, packets are load shared based on the following information:

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

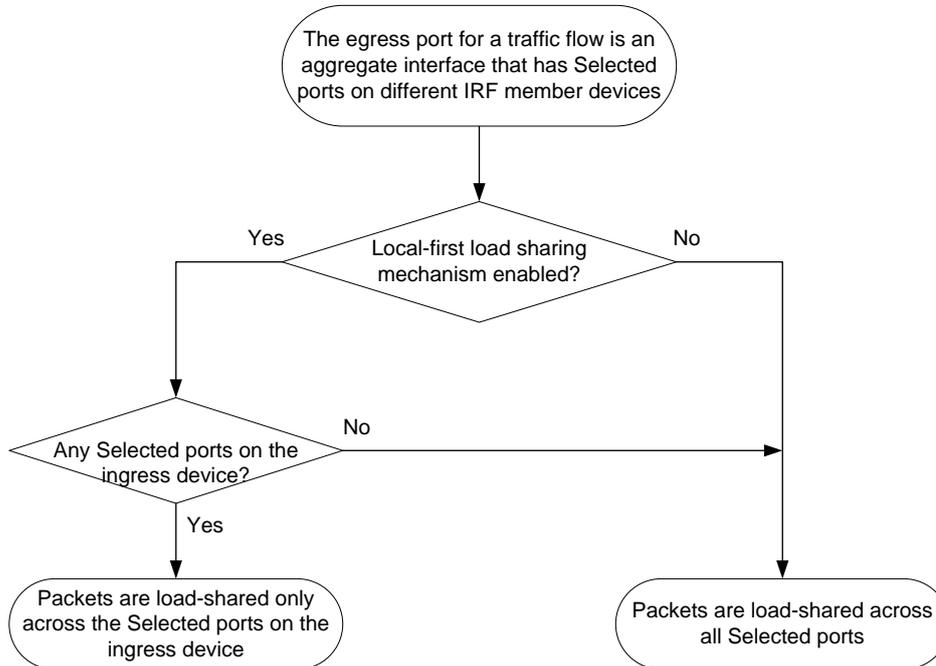
Enabling local-first load sharing for link aggregation

About 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 slot.

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 5](#). For more information about IRF, see *Virtual Technologies Configuration Guide*.

Figure 5 Load sharing for multidevice link aggregation in an IRF fabric



Enabling local-first load sharing for link aggregation globally

1. Enter system view.
system-view
2. Enable local-first load sharing for link aggregation globally.
link-aggregation load-sharing mode local-first
By default, local-first load sharing is enabled globally.

Enabling a Layer 2 aggregate interface to reflect incoming packets back

About reflecting incoming packets on a Layer 2 aggregate interface

By default, the device drops a packet if its outgoing interface is the incoming interface where the packet arrived. To have a Layer 2 aggregate interface reflect a packet back when it is both the incoming and outgoing interfaces of that packet, perform this task.

Procedure

1. Enter system view.
system-view
2. Enter Layer 2 aggregate interface view.
interface bridge-aggregation *interface-number*
3. Enable port bridging.
port bridge enable

By default, port bridging is disabled. A Layer 2 aggregate interface cannot reflect incoming packets back.

Enabling BFD for an aggregation group

About BFD for Ethernet link aggregation

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 BFD on the local port receives packets from the peer port upon link recovery, BFD notifies the Ethernet link aggregation module that the peer port is reachable. Then, 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 local port is placed in Selected state again upon link recovery, the local port establishes a new session with the peer port and 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*.

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 reversed between 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` at the local end, execute `link-aggregation bfd ipv4 source 2.2.2.2 destination 1.1.1.1` at the peer end. The source and destination IP addresses cannot be the same.
- 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 or the Demand mode.
- As a best practice, do not configure a protocol to collaborate with BFD on a BFD-enabled aggregate interface.
- Make sure the number of member ports in a BFD-enabled aggregation group is less than or identical to the number of BFD sessions supported by the device. If the aggregation group contains more member ports than the supported sessions, some Selected ports might change to the Unselected state.
- If the number of BFD sessions differs between the two ends of an aggregate link, check their settings for inconsistency in the maximum number of Selected ports. You must make sure the two ends have the same setting for the maximum number of Selected ports.

Procedure

1. Enter system view.
`system-view`
2. Enter Layer 2 aggregate interface view.
`interface bridge-aggregation interface-number`
3. Enable BFD for the aggregation group.
`link-aggregation bfd ipv4 source ip-address destination ip-address`
By default, BFD is disabled for an aggregation group.

Enabling link-aggregation traffic redirection

About link-aggregation traffic redirection

This feature operates on dynamic link aggregation groups. It redirects traffic on a Selected port to the remaining available Selected ports of an aggregation group if one of the following events occurs:

- The port is shut down by using the **shutdown** command.
- The slot that hosts the port reboots, and the aggregation group spans multiple slots.

NOTE:

The device does not redirect traffic to member ports that become Selected during the traffic redirection process.

This feature ensures zero packet loss for known unicast traffic, but it does not protect unknown unicast traffic.

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.

Restrictions and guidelines for link-aggregation traffic redirection

Link-aggregation traffic redirection applies only to dynamic link aggregation groups.

As a best practice, enable link-aggregation traffic redirection on a per-interface basis. 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.

To prevent traffic interruption, enable link-aggregation traffic redirection at both ends of the aggregate link.

To prevent packet loss that might occur at a reboot, do not enable the spanning tree feature together with link-aggregation traffic redirection.

Link-aggregation traffic redirection does not operate correctly on an edge aggregate interface.

Enabling link-aggregation traffic redirection globally

1. Enter system view.
system-view
2. Enable link-aggregation traffic redirection globally.
link-aggregation lacp traffic-redirect-notification enable
By default, link-aggregation traffic redirection is disabled globally.

Enabling link-aggregation traffic redirection for an aggregation group

1. Enter system view.

system-view

2. Enter Layer 2 aggregate interface view.

interface bridge-aggregation *interface-number*

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.

Display and maintenance commands for Ethernet link aggregation

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

Task	Command
Display information about aggregate interfaces.	display interface [bridge-aggregation [<i>interface-number</i>]] [brief [description down]]
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 <i>interface-number</i>]]
Display detailed link aggregation information about link aggregation member ports.	display link-aggregation member-port [<i>interface-list</i> auto]
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 [<i>interface-number</i>]]
Clear statistics for the specified aggregate interfaces.	reset counters interface [bridge-aggregation [<i>interface-number</i>]]
Clear LACP statistics for the specified link aggregation member ports.	reset lacp statistics [interface <i>interface-list</i>]

Ethernet link aggregation configuration examples

Example: Configuring a Layer 2 static aggregation group

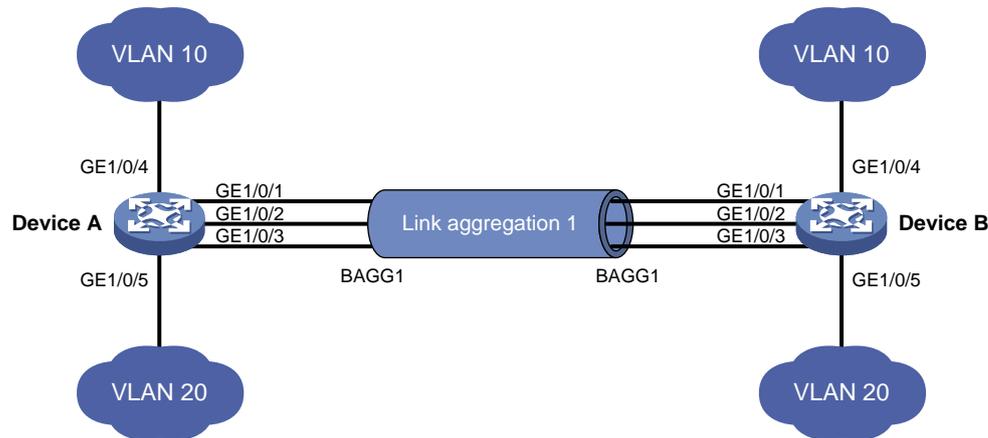
Network configuration

On the network shown in [Figure 6](#), 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 6 Network diagram



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
Port: A -- Auto port, M -- Management port, R -- Reference port
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
Management VLANs: None

  Port                Status  Priority Oper-Key
  ---                ---
GE1/0/1(R)           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 that contains three Selected ports.

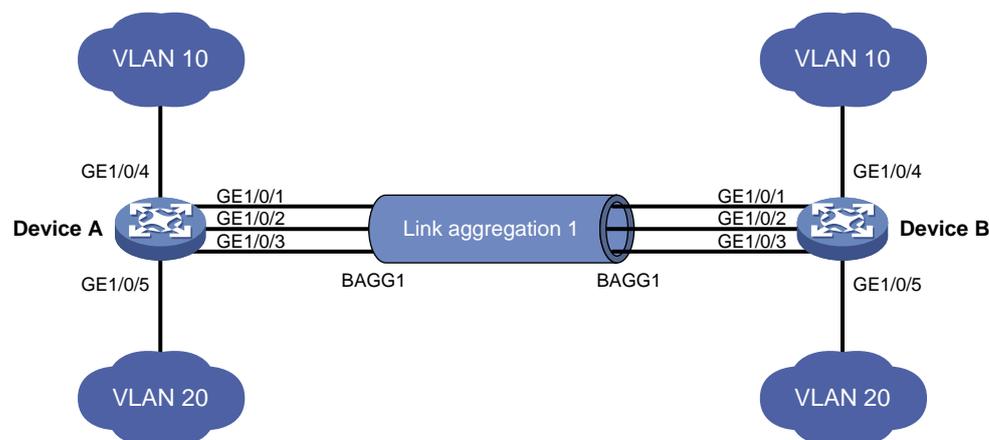
Example: Configuring a Layer 2 dynamic aggregation group

Network configuration

On the network shown in [Figure 7](#), 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 7 Network diagram



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
Port: A -- Auto port, M -- Management port, R -- Reference port
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
```

```
Management VLANs: None
```

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

```
Local:
```

Port	Status	Priority	Index	Oper-Key	Flag
GE1/0/1(R)	S	32768	11	1	{ACDEF}

```

GE1/0/2          S          32768   12      1              {ACDEF}
GE1/0/3          S          32768   13      1              {ACDEF}
Remote:
Actor           Priority Index   Oper-Key SystemID      Flag
GE1/0/1         32768   81     1         0x8000, 000f-e267-57ad {ACDEF}
GE1/0/2         32768   82     1         0x8000, 000f-e267-57ad {ACDEF}
GE1/0/3         32768   83     1         0x8000, 000f-e267-57ad {ACDEF}

```

The output shows that link aggregation group 1 is a Layer 2 dynamic aggregation group that contains three Selected ports.

Example: Configuring a Layer 2 edge aggregate interface

Network configuration

As shown in [Figure 8](#), 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 8 Network diagram



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
Port: A -- Auto port, M -- Management port, R -- Reference port
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

Management VLANs: None

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

Local:

Port	Status	Priority	Index	Oper-Key	Flag
GE1/0/1	I	32768	11	1	{AG}
GE1/0/2	I	32768	12	1	{AG}

Remote:

Actor	Priority	Index	Oper-Key	SystemID	Flag
GE1/0/1	32768	81	0	0x8000, 0000-0000-0000	{DEF}
GE1/0/2	32768	82	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.