HP A5120 EI Switch Series
Layer 2 - LAN Switching

Configuration Guide

Abstract

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

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

Part number: 5998-1791
Software version: Release 2208
Document version: 5W100-20110530

Legal and notice information

© Copyright 2011 Hewlett-Packard Development Company, L.P.

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

The information contained herein is subject to change without notice.

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

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Contents

Ethernet interface configuration

Ethernet interface naming conventions

The GE and 10-GE interfaces on the A5120 EI Switch Series are named in the format of *interface-type* A/B/C, where the following definitions apply:

- If the switch does not support Intelligent Resilient Framework (IRF), A takes 1. If the switch supports IRF, A represents the ID of the switch in an IRF fabric. If the switch is not assigned to any IRF fabric, A uses 1.
- B represents a slot number on the switch. It uses 0 for fixed interfaces, 1 for interfaces on interface expansion card 1, and 2 for interfaces on interface expansion card 2.
- C represents the number of an interface on a slot.

NOTE:

- For more information about the expansion cards, see the *HP A5120 EI Switch Series Installation Guide.*
- The HP A5120-24G EI Switch(JE066A) and the HP A5120-48G EI Switch(JE067A) do not support IRF.

Configuring basic settings of an Ethernet interface

Configuring a combo interface

Introduction to combo interfaces

A combo interface is a logical interface that comprises one optical (fiber) port and one electrical (copper) port. The two ports share one forwarding interface, so they cannot work simultaneously. When you enable one port, the other is automatically disabled.

The fiber combo port and the copper combo port are Layer 2 Ethernet interfaces. They have their own separate interface views, in which you can activate the fiber or copper combo port and configure other port attributes such as the interface rate and duplex mode.

Configuration prerequisites

Before you configure a combo interface, complete the following tasks:

- Use the **display port combo** command to identify the combo interfaces on your device and identify the two physical ports that compose each combo interface.
- Use the **display interface** command to determine, of the two physical ports that compose a combo interface, which is the fiber combo port and which is the cooper combo port. If the current port is the copper port, the output will include "Media type is twisted pair, Port hardware type is 1000_BASE_T". If the current port is the fiber port, the output will include "Media type is not sure, Port hardware type is No connector".

Changing the active port of a combo interface

Follow these steps to change the active port of a double combo interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Activate the current interface	**undo shutdown**	Optional By default, of the two ports that compose a combo interface, the one with a smaller port ID is active.

Configuring basic settings of an Ethernet interface

You can set an Ethernet interface to operate in one of the following duplex modes:

- Full-duplex mode (full): Interfaces that operates in this mode can send and receive packets simultaneously.
- Half-duplex mode (half): Interfaces that operates in this mode cannot send and receive packets simultaneously.
- Auto-negotiation mode (auto): Interfaces that operates in this mode negotiate a duplex mode with their peers.

You can set the speed of an Ethernet interface or enable it to automatically negotiate a speed with its peer. For a 100-Mbps or 1000-Mbps Layer 2 Ethernet interface, you can also set speed options for auto negotiation. The two ends can select a speed only from the available options. For more information, see "Setting speed options for auto negotiation on an Ethernet interface."

Follow these steps to configure an Ethernet interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Change the description of the interface	**description** *text*	Optional By default, the description of an interface is the interface name followed by the "Interface" string, **GigabitEthernet1/0/1 Interface** for example.
Set the duplex mode	**duplex** { **auto** | **full** | **half** }	Optional The optical port of an SFP port and the electrical port of an Ethernet port whose port rate is configured as 1000 Mbps do not support the **half** keyword. The default duplex mode of a port is auto negotiation.

To do…	Use the command…	Remarks
Set the interface speed	**speed** { **10** | **100** | **1000** | **auto** }	Optional The optical port of an SFP port does not support the **10** and **100** keywords. By default, the **auto** option is enabled.
Shut down the Ethernet interface	**shutdown**	Optional By default, an Ethernet interface is in the up state. To bring up an Ethernet interface, use the **undo shutdown** command.

Setting speed options for auto negotiation on an Ethernet interface

NOTE:

Optical interfaces do not support this feature.

As shown in Figure 1, speed auto negotiation enables an Ethernet interface to negotiate with its peer for the highest speed that both ends support by default. You can narrow down the speed option list for negotiation.

Figure 1 Speed auto negotiation application scenario

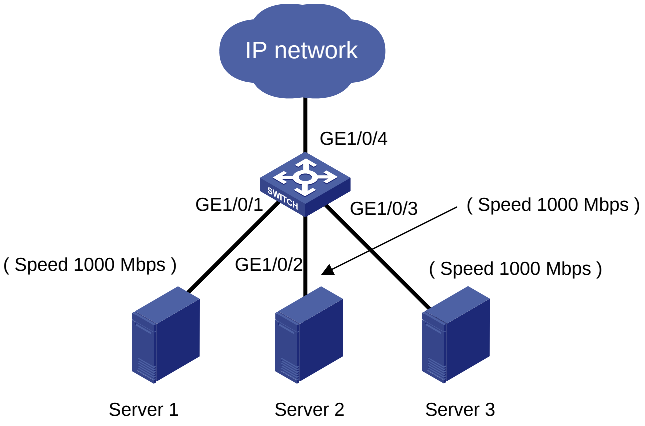

All interfaces on the switch are operating in speed auto negotiation mode, with the highest speed of 1000 Mbps. If the transmission rate of each server in the server cluster is 1000 Mbps, their total transmission rate will exceed the capability of interface GigabitEthernet 1/0/4, the interface providing access to the Internet for the servers.

To avoid congestion on GigabitEthernet 1/0/4, set 100 Mbps as the only speed option available for negotiation on interface GigabitEthernet 1/0/1, GigabitEthernet 1/0/2, and GigabitEthernet 1/0/3. As a result, the transmission rate on each interface connected to a server is limited to 100 Mbps.

Follow these steps to configure an auto-negotiation transmission rate:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Set speed options for auto negotiation	**speed auto** [**10** | **100** | **1000**] *	Optional

NOTE:

- This function is only available for Gigabit Layer-2 copper (electrical) Ethernet interfaces that support speed auto negotiation.
- The **speed** and **speed auto** commands supersede each other, and whichever is configured last takes effect.

Configuring generic flow control on an Ethernet interface

An interface implements generic flow control by sending and receiving common pause frames. The following generic flow control modes are available:

- TxRx mode enables an interface to both send and receive common pause frames.
- Rx mode enables an interface to receive but not send common pause frames.

In Figure 2, when both Port A and Port B forward packets at 1000 Mbps, Port C is congested. To avoid packet loss, enable flow control on Port A and Port B.

Figure 2 Flow control application scenario

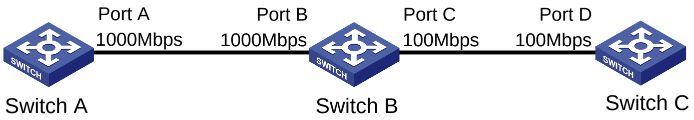

Configure Port B to operate in TxRx mode, Port A in Rx mode.

- When congestion occurs on Port C, Switch B buffers frames. When the amount of buffered frames exceeds a certain value, Switch B sends a common pause frame out of Port B to ask Port A to suspend sending packets. This pause frame also tells Port A for how long it is expected to pause.
- Upon receiving the common pause frame from Port B, Port A suspends sending packets to Port B for a period.
- If congestion persists, Port B keeps sending common pause frames to Port A until the congestion condition is removed.

Follow these steps to configure flow control on an interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Enable TxRx mode flow control	**flow-control**	Required Use either command. By default, flow control is disabled on an Ethernet interface.
Enable Rx mode flow control	**flow-control receive enable**	

Configuring link change suppression on an Ethernet interface

An Ethernet interface has two physical link states: up and down. Each time the physical link of an interface goes up or comes down, the physical layer reports the change to the upper layers, and the upper layers handle the change, resulting in increased overhead.

To prevent physical link flapping from affecting system performance, configure link change suppression to delay the reporting of physical link state changes. When the delay expires, the interface reports any detected change.

Link change suppression does not suppress administrative up or down events. When you shut down or bring up an interface by using the **shutdown** or **undo shutdown** command, the interface reports the event to the upper layers immediately.

On an A5120 EI switch, you can configure link-down suppression or link-up suppression, but not both.

Link-down suppression enables an interface to suppress link-down events and start a delay timer each time the physical link goes down. During this delay, the interface does not report the link-down event, and the **display interface brief** or **display interface** command displays the interface state as UP. If the physical link is still down when the timer expires, the interface reports the link-down event to the upper layers.

Link-up suppression enables an interface to suppress link-up events and start a delay timer each time the physical link goes up. During this delay, the interface does not report the link-up event, and the **display interface brief** or **display interface** command displays the interface state as DOWN. If the physical link is still up when the timer expires, the interface reports the link-up event to the upper layers.

Configuring link-down suppression

Follow these steps to enable an Ethernet interface to suppress link-down events:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Set a link-down suppression interval	**link-delay** *delay-time*	Required Link-down suppression is disabled by default.

Configuring link-up suppression

Follow these steps to configure link-up suppression on an Ethernet interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Set a link-up suppression interval	**link-delay** *delay-time* **mode up**	Required Link-up suppression is disabled by default.

NOTE:

The **link-delay mode up** command and the **link-delay** command supersedes each other, and whichever is configured last takes effect.

Configuring loopback testing on an Ethernet interface

You can perform loopback testing on an Ethernet interface to determine whether the interface functions properly. The Ethernet interface cannot forward data packets during the testing. Loopback testing falls into the following categories:

- Internal loopback testing, which tests all on-chip functions related to Ethernet interfaces. As shown in Figure 3, internal loopback testing is performed on Port 1. During internal loopback testing, the interface sends a certain number of test packets, which are looped back to the interface over the self-loop created on the switching chip.

Figure 3 Internal loopback testing

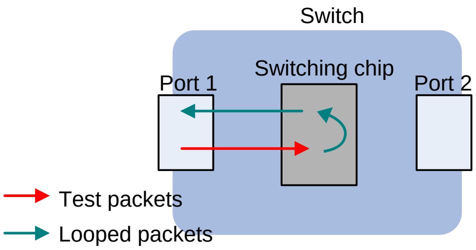

- External loopback testing, which tests the hardware of Ethernet interfaces. As shown in Figure 4, external loopback testing is performed on Port 1. To perform external loopback testing on an Ethernet interface, insert a loopback plug into the interface. During external loopback testing, the interface sends a certain number of test packets, which are looped over the plug and back to the interface. If the interface fails to receive any test packet, the hardware of the interface is faulty.

Figure 4 External loopback testing

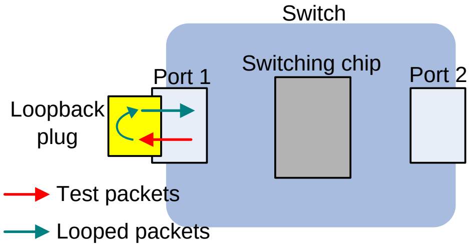

Follow these steps to perform loopback testing on an Ethernet interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Perform loopback testing	**loopback** { **external** | **internal** }	Required

NOTE:

- On an interface that is physically down, you can only perform internal loopback testing. On an interface administratively shut down, you can perform neither internal nor external loopback testing.
- The **speed**, **duplex**, **mdi**, and **shutdown** commands are not available during loopback testing.
- During loopback testing, the Ethernet interface operates in full duplex mode. When you disable loopback testing, the port returns to its duplex setting.

Configuring a port group

Some interfaces on your switch might use the same set of settings. To configure these interfaces in bulk rather than one by one, you can assign them to a port group.

You create port groups manually. All settings made for a port group apply to all the member ports of the group. For example, you can configure a traffic suppression threshold (see "Configuring traffic storm protection") for multiple interfaces in bulk by assigning these interfaces to a port group.

Even though the settings are made on the port group, they are saved on each interface basis rather than on a port group basis. You can only view the settings in the view of each interface by using the **display current-configuration** or **display this** command.

Follow these steps to configure a port group:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Create a port group and enter port group view	**port-group manual** *port-group-name*	Required
Assign Ethernet interfaces to the port group	**group-member** *interface-list*	Required
Shut down all Ethernet interfaces in the port group	**shutdown**	Optional By default, all Ethernet interfaces in a port group are up. To bring up all Ethernet interfaces shut down manually in a port group, use the **undo shutdown** command in port group view.

Configuring traffic storm protection

A traffic storm occurs when a large amount of broadcast, multicast, or unknown unicast packets congest a network. The A5120 EI switches provide the following storm protection approaches:

- Storm suppression, which you can use to limit the size of monitored traffic that passes through an Ethernet interface by setting a traffic threshold. The port discards all traffic that exceeds the threshold.
- Storm control, which you can use to shut down Ethernet interfaces or block traffic when monitored traffic exceeds the traffic threshold. Depending on your configuration, storm control can also enable an interface to send trap or log messages when monitored traffic reaches a certain traffic threshold.

For a particular type of traffic, configure either storm suppression or storm control, but not both. If you configure both of them, you might fail to achieve the expected storm control effect.

Configuring storm suppression on an Ethernet interface

You can use the following guidelines to set one suppression threshold for broadcast, multicast, and unknown unicast traffic separately on an Ethernet interface.

- Set the threshold as a percentage of the interface transmission capability.
- Set the threshold in kbps, limiting the number of kilobits of monitored traffic passing through the interface per second.
- Set the threshold in pps, limiting the number of monitored packets passing through the interface per second.

NOTE:

If one suppression threshold has been set in pps on an Ethernet interface, you must set other suppression thresholds in pps. If one suppression threshold has been set in percentage or kbps, you cannot set other suppression thresholds in pps.

Follow these steps to configure storm suppression on an Ethernet interface:

To do…		**Use the command…**	**Remarks**
Enter system view		**system-view**	—
Enter Ethernet interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Use either command. To configure storm suppression on one Ethernet interface, enter Ethernet interface view. To configure storm suppression on a group of Ethernet interfaces, enter port group view.
	Enter port group view	**port-group manual** *port-group-name*	
Set a broadcast suppression threshold		**broadcast-suppression** { *ratio* | **pps** *max-pps* | **kbps** *max-bps* }	Optional By default, all broadcast traffic is allowed to pass through an interface.
Set a multicast suppression threshold		**multicast-suppression** { *ratio* | **pps** *max-pps* | **kbps** *max-bps* }	Optional By default, all multicast traffic is allowed to pass through an interface.
Set a unknown unicast suppression threshold		**unicast-suppression** { *ratio* | **pps** *max-pps* | **kbps** *max-bps* }	Optional By default, all unknown unicast traffic is allowed to pass through an interface.

NOTE:

If you set a storm suppression threshold in both Ethernet interface view and port group view, the threshold configured last takes effect.

Configuring storm control on an Ethernet interface

Storm control compares broadcast, multicast, and unknown unicast traffic regularly with their respective traffic thresholds on an Ethernet interface. For each type of traffic, storm control provides a lower threshold and a higher threshold.

For management purposes, you can configure the interface to send threshold event traps and log messages when monitored traffic exceeds the upper threshold or falls below the lower threshold from the upper threshold.

When the traffic exceeds its higher threshold, the interface does either of the following, depending on your configuration:

- Blocks the particular type of traffic, while forwarding other types of traffic. Even though the interface does not forward the blocked traffic, it still counts the traffic. When the blocked traffic drops below the threshold, the interface begins to forward the traffic.
- Shuts down automatically. The interface shuts down automatically and stops forwarding any traffic. To bring up the interface, use the **undo shutdown** command or disable the storm control function.

Follow these steps to configure the storm control function on an Ethernet interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Set the traffic polling interval of the storm control module	**storm-constrain interval** *seconds*	Optional 10 seconds by default.
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Enable storm control, and set the lower and upper thresholds for broadcast, multicast, or unknown unicast traffic	**storm-constrain** { **broadcast** | **multicast** | **unicast** } { **pps** | **kbps** | **ratio** } *max-pps-values min-pps-values*	Required Disabled by default.
Set the control action to take when monitored traffic exceeds the upper threshold	**storm-constrain control** { **block** | **shutdown** }	Optional Disabled by default.
Enable the interface to send storm control threshold event traps.	**storm-constrain enable trap**	Optional By default, the interface sends traps when monitored traffic exceeds the upper threshold or drops below the lower threshold from the upper threshold.
Enable the interface to log storm control threshold events.	**storm-constrain enable log**	Optional By default, the interface outputs log messages when monitored traffic exceeds the upper threshold or falls below the lower threshold from the upper threshold.

NOTE:

- For network stability, use the default or set a higher traffic polling interval.
- Storm control uses a complete polling cycle to collect traffic data, and analyzes the data in the next cycle. An interface takes one to two polling intervals to take a storm control action.

Setting the statistics polling interval

Follow these steps to set the statistics polling interval on an Ethernet interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Set the statistics polling interval on the Ethernet interface	**flow-interval** *interval*	Optional The default interface statistics polling interval is 300 seconds.

To display the interface statistics collected in the last polling interval, use the **display interface** command.

To clear interface statistics, use the **reset counters interface** command.

Enabling the auto power-down function on an Ethernet interface

To save power, enable the auto power-down function on Ethernet interfaces. An interface enters the power save mode if it has not received any packet for a certain period of time (this interval depends on the specifications of the chip, and is not configurable). When a packet arrives later, the interface enters its normal state.

Follow these steps to enable auto power-down on an Ethernet interface:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter Ethernet interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Use either command. To configure auto power-down on one Ethernet interface, enter Ethernet interface view. To configure auto power-down on a group of Ethernet interfaces, enter port group view.
	Enter port group view	**port-group manual** *port-group-name*	
Enable auto power-down on an Ethernet interface		**port auto-power-down**	Required Disabled by default.

Configuring jumbo frame support

Ethernet frames longer than the standard Ethernet frame size (1536 bytes) are called "jumbo frames", which are typical of file transfer.

- If you set an Ethernet interface to accept jumbo frames, it allows frames up to 9216 bytes to pass through.
- If you disable an Ethernet interface to accept jumbo frames, it allows frames up to 1536 bytes to pass through.

Follow these steps to configure jumbo frame support in Ethernet interface view:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter Ethernet interface view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Use either command.

To do…	Use the command…	Remarks
or port group view	Enter port group view	To configure jumbo frame support on one Ethernet interface, enter Ethernet interface view. To configure jumbo frame support on a group of Ethernet interfaces, enter port group view.
Enable the interface to accept jumbo frames	**jumboframe enable**	Required By default, an Ethernet interface accepts jumbo frames (up to 9216 bytes).

Enabling single-port loopback detection on an Ethernet interface

If an interface receives a packet that it sent, a loop occurs. Loops might cause broadcast storms, degrading network performance. You can use loopback detection to detect loops on an interface and configure the protective action to take on the interface when a loop is detected, for example, to shut down the interface. In addition to the configured protective action, the switch also performs other actions to alleviate the impact of the loop condition, as described in Table 1.

Table 1 Actions to take upon detection of a loop condition

Port type	Actions	
	No protective action is configured	**A protective action is configured**
Access interface	• Put the interface in controlled mode. The interface discards all incoming packets, but still forwards outgoing traffic. • Create traps. • Delete all MAC address entries of the interface.	• Perform the configured protective action. • Create traps and log messages. • Delete all MAC address entries of the interface.
Hybrid or trunk interface	• Create traps. • If loopback detection control is enabled, set the interface in controlled mode. The interface discards all incoming packets, but still forwards outgoing packets. • Delete all MAC address entries of the interface.	• Create traps and log messages. • If loopback detection control is enabled, take the configured protective action on the interface. • Delete all MAC address entries of the interface.

Follow these steps to configure single-port loopback detection:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enable global loopback detection	**loopback-detection enable**	Required Disabled by default.
Set the loopback detection interval	**loopback-detection interval-time** *time*	Optional 30 seconds by default.

To do…		Use the command…	Remarks
Enter Ethernet interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Use either command. To configure loopback detection on one interface, enter Ethernet interface view. To configure loopback detection on a group of Ethernet interfaces, enter port group view.
	Enter port group view	**port-group manual** *port-group-name*	
Enable loopback detection on the interface		**loopback-detection enable**	Required Disabled by default.
Enable loopback detection control		**loopback-detection control enable**	Optional Disabled by default.
Enable loopback detection in all VLANs on the trunk or hybrid interface		**loopback-detection per-vlan enable**	Optional By default, a trunk or hybrid interface performs loopback detection only in its PVID.
Set the protective action to take on the interface when a loop is detected		**loopback-detection action** { **shutdown** | **semi-block** | **no-learning** }	Optional By default, a looped interface discards all incoming packets but still forwards outgoing packets. The system generates traps and deletes all MAC address entries of the looped interface. With the **shutdown** keyword used, the switch shuts down looped Ethernet interfaces and sets their physical state to Loop down. When a looped interface recovers, you must use the **undo shutdown** command to restore its forwarding capability.

NOTE:

- To use single-port loopback detection on an Ethernet interface, you must enable the function both globally and on the interface.
- To disable loopback detection on all interfaces, run the **undo loopback-detection enable** command in system view.
- To enable a hybrid or trunk interface to take the administratively specified protective action, you must use the **loopback-detection control enable** command on the interface.
- When you change the link type of an Ethernet interface by using the **port link-type** command, the switch removes the protective action configured on the interface. For more information about the **port link-type** command, see the *Layer 2—LAN Switching Command Reference*.

Enabling multi-port loopback detection

When an interface receives packets sent from another interface on the same switch, a loop occurs between the two interfaces. Such a loop is called a "multi-port loop". As shown in Figure 5, if Port 1 receives packets sent out Port 2, a multi-port loop occurs between the two interfaces, and Port 1 (the interface that receives the looped packets) is the looped interface. Multi-port loops might also cause broadcast storms.

Figure 5 Network diagram for multi-port loopback detection

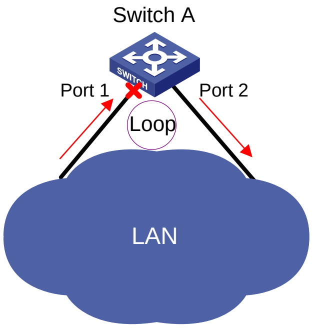

The multi-port loopback detection function detects loops among interfaces on your switch. You can use the **loopback-detection action** command to configure the protective action to take on looped interfaces— for example, to shut down the interface, eliminating the loops. In addition, the switch also takes other link type-dependant actions on the looped interface (for example, Port 1 in Figure 5) to alleviate the impact of the loop condition. For more information, see "Setting the statistics polling interval."

Multi-port loopback detection is implemented on the basis of single-port loopback detection configurations on Ethernet interfaces. To implement multi-port loopback detection, you must enable single-port loopback detection on one or multiple Ethernet interfaces on the switch.

Follow these steps to configure multi-port loopback detection:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Enable multi-port loopback detection	**loopback-detection multi-port-mode enable**	Required Disabled by default.

NOTE:

- To enable multi-port loopback detection, you must configure the **loopback-detection multi-port-mode enable** and **loopback-detection enable** commands in system view, and configure the **loopback-detection enable** command in the view of the related interfaces.
- The single-port loopback detection function is available when the switch is performing multi-port loopback detection.

Setting the MDI mode of an Ethernet interface

NOTE:

Optical interfaces do not support the MDI mode setting.

You can use both crossover and straight-through Ethernet cables to connect copper Ethernet interfaces. To accommodate these types of cables, a copper Ethernet interface can operate in one of the following Medium Dependent Interface (MDI) modes:

- Across mode
- Normal mode
- Auto mode

A copper Ethernet interface uses an RJ-45 connector, which comprises eight pins, each of which plays a dedicated role. For example, pins 1 and 2 transmit signals, and pins 3 and 6 receive signals. The pin role varies by the MDI modes as follows:

- In normal mode, pins 1 and 2 are transmit pins, and pins 3 and 6 are receive pins.
- In across mode, pins 1 and 2 are receive pins, and pins 3 and 6 are transmit pins.
- In auto mode, the interface negotiates pin roles with its peer.

To enable the interface to communicate with its peer, ensure that its transmit pins are connected to the remote receive pins. If the interface can detect the connection cable type, set the interface in auto MDI mode. If not, set its MDI mode by using the following guidelines:

- When a straight-through cable is used, set the interface to work in the MDI mode different than its peer.
- When a crossover cable is used, set the interface to work in the same MDI mode as its peer, or set either end to work in auto mode.

Follow these steps to set the MDI mode of an Ethernet interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Set the MDI mode of the Ethernet interface	**mdi** { **across** | **auto** | **normal** }	Optional By default, a copper Ethernet interface operates in auto mode to negotiate pin roles with its peer.

Enabling bridging on an Ethernet interface

When an incoming packet arrives, the switch looks up the destination MAC address of the packet in the MAC address table. If an entry is found, but the outgoing interface is the same as the receiving interface (for example, if the destination and source MAC addresses of the packet are the same), the switch discards the packet.

To enable the switch to return such packets to the sender through the receiving interface rather than drop them, enable the bridging function on the Ethernet interface.

Follow these steps to enable bridging on an Ethernet interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Enable bridging on the Ethernet interface	**port bridge enable**	Required Disabled by default.

Testing the cable connection of an Ethernet interface

NOTE:

- Optical interfaces do not support this feature.
- If the link of an Ethernet interface is up, testing its cable connection will cause the link to come down and then go up.

You can test the cable connection of an Ethernet interface for a short or open circuit. The device displays cable test results within five seconds. If any fault is detected, the test results include the length of the faulty cable segment.

Follow these steps to test the cable connection of an Ethernet interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Test the cable connected to the Ethernet interface	**virtual-cable-test**	Required

Displaying and maintaining an Ethernet interface

To do…	Use the command…	Remarks
Display the current state of an interface and the related information	**display interface** [*interface-type* [*interface-number*]] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the summary of an interface	**display interface** [*interface-type* [*interface-number*]] **brief** [**|** { **begin** | **exclude** | **include** } *regular-expression*] **display interface** [*interface-type*] **brief down** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the statistics on the packets that pass through a specific type of interfaces	**display counters** { **inbound** | **outbound** } **interface** [*interface-type*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the statistics on the rate of the packets that pass through the interfaces that are of a specific type and are in the up state in the latest sampling interval	**display counters rate** { **inbound** | **outbound** } **interface** [*interface-type*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display information about discarded packets on an interface	**display packet-drop interface** [*interface-type* [*interface-number*]] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display summary information about discarded packets on all interfaces	**display packet-drop summary** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display information about a manual port group or all manual port groups	**display port-group manual** [**all** | **name** *port-group-name*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

To do…	Use the command…	Remarks
Display information about the loopback function	**display loopback-detection** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display information about storm control on interfaces	**display storm-constrain** [**broadcast** | **multicast** | **unicast**] [**interface** *interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Clear the statistics of an interface	**reset counters interface** [*interface-type* [*interface-number*]]	Available in user view
Clear the statistics of discarded packets on an interface	**reset packet-drop interface** [*interface-type* [*interface-number*]]	Available in user view
Display the combo interfaces and the fiber and copper combo ports	**display port combo** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

Loopback and null interface configuration

Loopback interface

Introduction to loopback interface

A loopback interface is a software-only virtual interface. It delivers the following benefits.

- The physical layer state and link-layer protocols of a loopback interface are always up unless the loopback interface is manually shut down.
- To save IP address resources, you can assign an IP address with an all-F mask to a loopback interface. When you assign an IPv4 address whose mask is not 32-bit, the system automatically changes the mask into a 32-bit mask. When you assign an IPv6 address whose mask is not 128-bit, the system automatically changes the mask into a 128-bit mask.
- You can enable routing protocols on a loopback interface, and a loopback interface can send and receive routing protocol packets.

You can configure a loopback interface address as the source address of the IP packets that the switch generates. Because loopback interface addresses are stable unicast addresses, they are usually used as device identifications. When you configure a rule on an authentication or security server to permit or deny packets that a switch generates, you can simplify the rule by configuring it to permit or deny packets that carry the loopback interface address that identifies the switch. When you use a loopback interface address as the source address of IP packets, be sure to perform any necessary routing configuration to ensure that the route from the loopback interface to the peer is reachable. All data packets sent to the loopback interface are treated as packets sent to the switch itself, so the switch does not forward these packets.

Configuring a loopback interface

Follow these steps to configure a loopback interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Create a loopback interface and enter loopback interface view	**interface loopback** *interface-number*	—
Set a description for the loopback interface	**description** *text*	Optional By default, the description of an interface is the interface name followed by the "Interface" string.
Shut down the loopback interface	**shutdown**	Optional By default, a loopback interface is up after it is created.

NOTE:

You can configure settings such as IP addresses and IP routes on loopback interfaces. For more information, see the *Layer 3—IP Services Configuration Guide* and *Layer 3—IP Routing Configuration Guide*.

Null interface

Introduction to null interface

A null interface is a completely software-based logical interface, and is always up. However, you cannot use it to forward data packets or configure an IP address or link-layer protocol on it. With a null interface specified as the next hop of a static route to a specific network segment, any packets routed to the network segment are dropped. The null interface provides a simpler way to filter packets than ACL. You can filter uninteresting traffic by transmitting it to a null interface instead of applying an ACL.

For example, by executing the **ip route-static 92.101.0.0 255.255.0.0 null 0** command (which configures a static route that leads to null interface 0), you can have all the packets destined to the network segment 92.101.0.0/16 discarded.

Only one null interface, Null 0, is supported on your switch. You cannot remove or create a null interface.

Configuring null 0 interface

Follow these steps to enter null interface view:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter null interface view	**interface null 0**	Required The Null 0 interface is the default null interface on your switch. It cannot be manually created or removed.
Set a description for the null interface	**description** *text*	Optional By default, the description of an interface is the interface name followed by the "Interface" string.

Displaying and maintaining loopback and null interfaces

To do…	Use the command…	Remarks
Display information about loopback interfaces	**display interface loopback** [**brief** [**down**]] [| { **begin** | **exclude** | **include** } *regular-expression*] **display interface loopback** *interface-number* [**brief**] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

To do…	Use the command…	Remarks
Display information about the null interface	**display interface null** [**brief** [**down**]] [**|** { **begin** | **exclude** | **include** } *regular-expression*] **display interface null 0** [**brief**] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Clear the statistics on a loopback interface	**reset counters interface** [**loopback** [*interface-number*]]	Available in user view
Clear the statistics on the null interface	**reset counters interface** [**null** [**0**]]	Available in user view

MAC address table configuration

Overview

Every Ethernet switch maintains a MAC address table for forwarding frames through unicast instead of broadcast. This table describes from which port a MAC address (or host) can be reached. When forwarding a frame, the switch first looks up the MAC address of the frame in the MAC address table for a match. If the switch finds an entry, it forwards the frame out of the outgoing port in the entry. If the switch does not find an entry, it broadcasts the frame out of all but the incoming port.

How a MAC address table entry is created

The switch automatically obtains entries in the MAC address table, or you can add them manually.

MAC address learning

The switch can automatically populate its MAC address table by obtaining the source MAC addresses (called "MAC address learning") of incoming frames on each port.

When a frame arrives at a port, Port A for example, the switch performs the following tasks:

1. Verifies the source MAC address (for example, MAC-SOURCE) of the frame.
2. Looks up the MAC address in the MAC address table.
3. Updates an entry if it finds one. If the switch does not find an entry, it adds an entry for MAC-SOURCE and Port A.

The switch performs the learning process each time it receives a frame from an unknown source MAC address, until the MAC address table is fully populated.

After obtaining the source MAC address of a frame, the switch looks up the destination MAC address in the MAC address table. If the switch finds an entry for the MAC address, it forwards the frame out of the specific outgoing port, Port A in this example.

Manually configuring MAC address entries

With dynamic MAC address learning, a switch does not distinguish between illegitimate and legitimate frames, which can invite security hazards. For example, if a hacker sends frames with a forged source MAC address to a port different from the one that the real MAC address is connected to, the switch will create an entry for the forged MAC address, and forward frames destined for the legal user to the hacker instead.

To enhance the security of a port, you can bind specific user devices to the port by manually adding MAC address entries into the MAC address table of the switch. Because manually configured entries have higher priority than dynamically obtained ones, you can prevent hackers from stealing data using forged MAC addresses.

Types of MAC address table entries

A MAC address table can contain the following types of entries:

- Static entries, which are manually added and never age out.

- Dynamic entries, which can be manually added or dynamically obtained and might age out.
- Blackhole entries, which are manually configured and never age out. Blackhole entries are configured for filtering out frames with specific destination MAC addresses. For example, to block all packets destined for a specific user for security concerns, you can configure the MAC address of this user as a blackhole destination MAC address entry.

To adapt to network changes and prevent inactive entries from occupying table space, an aging mechanism is adopted for dynamic MAC address entries. Each time a dynamic MAC address entry obtained or created, an aging time starts. If the entry has not updated when the aging timer expires, the switch deletes the entry. If the entry has updated before the aging timer expires, the aging timer restarts.

NOTE:

A static or blackhole MAC address entry can overwrite a dynamic MAC address entry, but not vice versa.

MAC address table-based frame forwarding

When forwarding a frame, the switch adopts the following forwarding modes based on the MAC address table:

- Unicast mode: If an entry is available for the destination MAC address, the switch forwards the frame out of the outgoing interface indicated by the MAC address table entry.
- Broadcast mode: If the switch receives a frame with the destination address as all ones, or no entry is available for the destination MAC address, the switch broadcasts the frame to all the interfaces except the receiving interface.

Configuring the MAC address table

The MAC address table configuration tasks include:

- Manually configuring MAC address table entries
- Disabling MAC address learning
- Configuring the aging timer for dynamic MAC address entries
- Configuring the MAC learning limit on ports

These configuration tasks are all optional and can be performed in any order.

NOTE:

- The MAC address table can contain only Layer 2 Ethernet ports and Layer 2 aggregate interfaces.
- This document covers configuring static, dynamic, and blackhole unicast MAC address table entries. For more information about static multicast MAC address table entries, see the *IP Multicast Configuration Guide*.

Manually configuring MAC address table entries

To help prevent MAC address spoofing attacks and improve port security, you can manually add MAC address table entries to bind ports with MAC addresses. You can also configure blackhole MAC address entries to filter out packets with certain source or destination MAC addresses.

Follow these steps to add, modify, or remove entries in the MAC address table in system view:

To do…	Use the command…		Remarks
Enter system view	**system-view**		—
Configure MAC address table entries	Configure static or dynamic MAC address table entries	**mac-address** { **dynamic** | **static** } *mac-address* **interface** *interface-type interface-number* **vlan** *vlan-id*	Required Use either command. Make sure that you have created the VLAN and assign the interface to the VLAN.
	Configure blackhole MAC address table entries	**mac-address blackhole** *mac-address* **vlan** *vlan-id*	

Follow these steps to add or modify a MAC address table entry in interface view:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter interface view	**interface** *interface-type interface-number*	—
Configure a MAC address table entry	**mac-address** { **dynamic** | **static** } *mac-address* **vlan** *vlan-id*	Required Ensure that you have created the VLAN and assign the interface to the VLAN

Disabling MAC address learning

Sometimes, you might need to disable MAC address learning to prevent the MAC address table from being saturated, for example, when your switch is being attacked by a large amount of packets with different source MAC addresses.

Disabling global MAC address learning

Disabling global MAC address learning disables the learning function on all ports.

Follow these steps to disable MAC address learning:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Disable global MAC address learning	**mac-address mac-learning disable**	Required Enabled by default.

NOTE:

When MAC address learning is disabled, the obtained MAC addresses remain valid until they age out.

Disabling MAC address learning on ports

After enabling global MAC address learning, you can disable the function on a single port, or on all ports in a port group as needed.

Follow these steps to disable MAC address learning on an interface or a port group:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enable global MAC address learning		**undo mac-address mac-learning disable**	Optional Enabled by default.
Enter interface view or port group view	Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command. The configuration made in Layer 2 Ethernet or Layer 2 aggregate interface view takes effect on the current interface only. The configuration made in port group view takes effect on all the member ports in the port group.
	Enter port group view	**port-group manual** *port-group-name*	
Disable MAC address learning on the interface or all ports in the port group		**mac-address mac-learning disable**	Required Enabled by default.

NOTE:

- When MAC address learning is disabled, the obtained MAC addresses remain valid until they age out.
- For more information about port groups, see the chapter "Ethernet interface configuration."

Configuring the aging timer for dynamic MAC address entries

The MAC address table uses an aging timer for dynamic MAC address entries for security and efficient use of table space. If a dynamic MAC address entry has failed to update before the aging timer expires, the switch deletes the entry. This aging mechanism ensures that the MAC address table can quickly update to accommodate the latest network changes.

Set the aging timer appropriately. A long aging interval might cause the MAC address table to retain outdated entries, exhaust the MAC address table resources, and fail to update its entries to accommodate the latest network changes. A short interval might result in the removal of valid entries and unnecessary broadcasts, which might affect device performance.

Follow these steps to configure the aging timer for dynamic MAC address entries:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Configure the aging timer for dynamic MAC address entries	**mac-address timer** { **aging** *seconds* | **no-aging** }	Optional 300 seconds by default.

You can reduce broadcasts on a stable network by disabling the aging timer to prevent dynamic entries from unnecessarily aging out. By reducing broadcasts, you improve not only network performance, but also security, because you reduce the chances that a data packet will reach unintended destinations.

Configuring the MAC learning limit on ports

As the MAC address table is growing, the forwarding performance of your device might degrade. To prevent the MAC address table from getting so large that the forwarding performance degrades, you can limit the number of MAC addresses that a port can obtain.

Follow these steps to configure the MAC learning limit on a Layer 2 Ethernet interface or all ports in a port group:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter Layer 2 Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Use either command. The configuration made in Layer 2 Ethernet interface view takes effect on the current interface only. The configuration made in port group view takes effect on all the member ports in the port group.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the MAC learning limit on the interface or port group		**mac-address max-mac-count** *count*	Required No MAC learning limit is configured by default.

NOTE:

- Layer 2 aggregate interfaces do not support the **mac-address max-mac-count** command.
- Do not configure the MAC learning limit on any member ports of an aggregation group. Otherwise, the member ports cannot be selected.

Displaying and maintaining MAC address tables

To do...	Use the command...	Remarks
Display MAC address table information	**display mac-address** [*mac-address* [**vlan** *vlan-id*] | [[**dynamic** | **static**] [**interface** *interface-type interface-number*] | **blackhole**] [**vlan** *vlan-id*] [**count**]] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the aging timer for dynamic MAC address entries	**display mac-address aging-time** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the system or interface MAC address learning state	**display mac-address mac-learning** [*interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display MAC address statistics	**display mac-address statistics** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

MAC address table configuration example

Network requirements

- The MAC address of one host is 000f-e235-dc71 and belongs to VLAN 1. It is connected to GigabitEthernet 1/0/1 of the device. To prevent MAC address spoofing, add a static entry into the MAC address table of the device for the host.
- The MAC address of another host is 000f-e235-abcd and belongs to VLAN 1. Because this host once behaved suspiciously on the network, you can add a destination blackhole MAC address entry for the MAC address to drop all packets destined for the host.
- Set the aging timer for dynamic MAC address entries to 500 seconds.

Configuration procedure

# Add a static MAC address entry.

```
<Sysname> system-view
[Sysname] mac-address static 000f-e235-dc71 interface gigabitethernet 1/0/1 vlan 1
```

# Add a destination blackhole MAC address entry.

```
[Sysname] mac-address blackhole 000f-e235-abcd vlan 1
```

# Set the aging timer for dynamic MAC address entries to 500 seconds.

```
[Sysname] mac-address timer aging 500
```

# Display the MAC address entry for port GigabitEthernet 1/0/1.

```
[Sysname] display mac-address interface gigabitethernet 1/0/1
MAC ADDR          VLAN ID  STATE            PORT INDEX              AGING TIME

000f-e235-dc71   1       Config static   GigabitEthernet 1/0/1         NOAGED

  ---  1 mac address(es) found  ---
```

# Display information about the destination blackhole MAC address table.

```
[Sysname] display mac-address blackhole
MAC ADDR        VLAN ID    STATE            PORT INDEX              AGING TIME
000f-e235-abcd  1          Blackhole        N/A                     NOAGED

  ---  1 mac address(es) found  ---
```

# View the aging time of dynamic MAC address entries.

```
[Sysname] display mac-address aging-time
Mac address aging time: 500s
```

MAC Information configuration

Overview

Introduction to MAC Information

To monitor a network, you must monitor users who are joining and leaving the network. Because a MAC address uniquely identifies a network user, you can monitor users who are joining and leaving a network by monitoring their MAC addresses.

With the MAC Information function, Layer 2 Ethernet interfaces send Syslog or trap messages to the monitor end in the network when they obtain or delete MAC addresses. By analyzing these messages, the monitor end can monitor users who are accessing the network.

How MAC Information works

When a new MAC address is obtained or an existing MAC address is deleted on a device, the device writes related information about the MAC address to the buffer area used to store user information. When the timer set for sending MAC address monitoring Syslog or trap messages expires, or when the buffer reaches capacity, the device sends the Syslog or trap messages to the monitor end.

Configuring MAC Information

The MAC Information configuration tasks include:

- Enabling MAC Information globally
- Enabling MAC Information on an interface
- Configuring MAC Information mode
- Configuring the interval for sending Syslog or trap messages
- Configuring the MAC Information queue length

Enabling MAC Information globally

Follow these steps to enable MAC Information globally:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enable MAC Information globally	**mac-address information enable**	Required Disabled by default.

Enabling MAC Information on an interface

Follow these steps to enable MAC Information on an interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	—
Enable MAC Information on the interface	**mac-address information enable** { **added** | **deleted** }	Required Disabled by default.

NOTE:

To enable MAC Information on an Ethernet interface, enable MAC Information globally first.

Configuring MAC Information mode

Follow these steps to configure MAC Information mode:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Configure MAC Information mode	**mac-address information mode** { **syslog** | **trap** }	Optional **trap** by default.

Configuring the interval for sending Syslog or trap messages

To prevent Syslog or trap messages from being sent too frequently, you can set the interval for sending Syslog or trap messages.

Follow these steps to set the interval for sending Syslog or trap messages:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Set the interval for sending Syslog or trap messages	**mac-address information interval** *interval-time*	Optional One second by default.

Configuring the MAC Information queue length

To avoid losing user MAC address information, when the buffer that stores user MAC address information reaches capacity, the user MAC address information in the buffer is sent to the monitor end in the network, even if the timer set for sending MAC address monitoring Syslog or trap messages has not expired yet.

Follow these steps to configure the MAC Information queue length:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Configure the MAC Information queue length	**mac-address information queue-length** *value*	Optional 50 by default.

MAC Information configuration example

Network requirements

- Host A is connected to a remote server (Server) through Device.
- Enable MAC Information on GigabitEthernet 1/0/1 on Device. Device sends MAC address changes in Syslog messages to Host B through GigabitEthernet 1/0/3. Host B analyzes and displays the Syslog messages.

Figure 6 Network diagram for MAC Information configuration

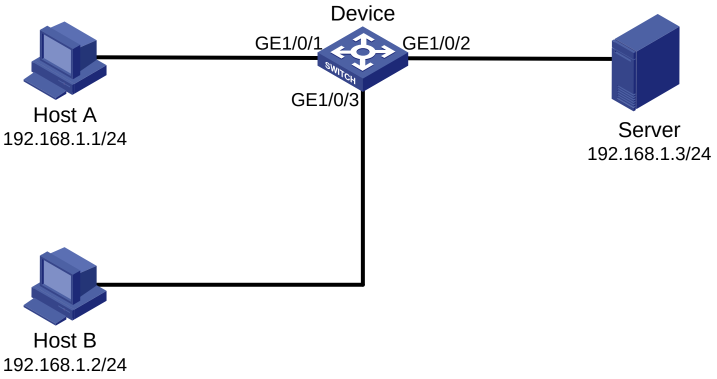

Configuration procedure

1. Configure Device to send Syslog messages to Host B.

For more information, see the *Network Management and Monitoring Configuration Guide*.

2. Enable MAC Information.

# Enable MAC Information on Device.

```
<Device> system-view
[Device] mac-address information enable
```

# Configure MAC Information mode as Syslog.

```
[Device] mac-address information mode syslog
```

# Enable MAC Information on GigabitEthernet 1/0/1.

```
[Device] interface gigabitethernet 1/0/1
[Device-GigabitEthernet1/0/1] mac-address information enable added
[Device-GigabitEthernet1/0/1] mac-address information enable deleted
[Device-GigabitEthernet1/0/1] quit
```

# Set the MAC Information queue length to 100.

```
[Device] mac-address information queue-length 100
```

# Set the interval for sending Syslog or trap messages to 20 seconds.

```
[Device] mac-address information interval 20
```

Ethernet link aggregation configuration

Overview

Ethernet link aggregation, or simply link aggregation, combines multiple physical Ethernet ports into one logical link, called an "aggregate link". Link aggregation delivers the following benefits:

- Increases bandwidth beyond the limits of any single link. In an aggregate link, traffic is distributed across the member ports.
- Improves 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 7, Device A and Device B are connected by three physical Ethernet links. These physical Ethernet links are combined into an aggregate link, Link aggregation 1. The bandwidth of this aggregate link is as high as the total bandwidth of the three physical Ethernet links. At the same time, the three Ethernet links back up each other.

Figure 7 Diagram for Ethernet link aggregation

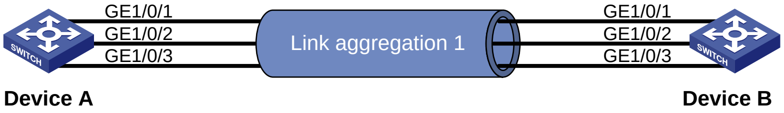

Basic concepts

Aggregation group, member port, aggregate interface

Link aggregation is implemented through link aggregation groups. An aggregation group is a group of Ethernet interfaces aggregated together, which are called "member ports" of the aggregation group. For each aggregation group, a logical interface, called an "aggregate interface", is created. To an upper layer entity that uses the link aggregation service, a link aggregation group appears to be a single logical link and data traffic is transmitted through the aggregate interface.

When you create an aggregate interface, the switch automatically creates an aggregation group of the same type and number as the aggregate interface. For example, when you create interface Bridge-aggregation 1, Layer 2 aggregation group 1 is created.

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

NOTE:

The rate of an aggregate interface equals the total rate of its member ports in the Selected state, and its duplex mode is the same as 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 either of the following aggregation states:

- Selected: A Selected port can forward user traffic.

- Unselected: An Unselected port cannot forward user traffic.

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 this operational key.

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

Configuration classes

Every configuration setting on a port might affect its aggregation state. Port configurations fall into the following classes:

- Port attribute configurations, including port rate, duplex mode, and link status (up/down). These are the most basic port configurations.
- Class-two configurations. A member port can be placed in the Selected state only if it has the same class-two configurations as the aggregate interface.

Table 2 Class-two configurations

Feature	Considerations
Port isolation	Whether the port has joined an isolation group
QinQ	QinQ enable state (enable/disable), TPID for VLAN tags, outer VLAN tags to be added, inner-to-outer VLAN priority mappings, inner-to-outer VLAN tag mappings, inner VLAN ID substitution mappings
VLAN	Permitted VLAN IDs, PVID, link type (trunk, hybrid, or access), IP subnet-based VLAN configuration, protocol-based VLAN configuration, VLAN tagging mode
MAC address learning	MAC address learning capability, MAC address learning limit, forwarding of frames with unknown destination MAC addresses after the MAC address learning limit is reached

NOTE:

- Class-two configurations made on an aggregate interface are automatically synchronized to all member ports of the interface. These configurations are retained on the member ports even after the aggregate interface is removed.
- Any class-two configuration change might affect the aggregation state of link aggregation member ports and ongoing traffic. To be sure that you are aware of the risk, the system displays a warning message every time you attempt to change a class-two configuration setting on a member port.

- Class-one configurations do not affect the aggregation state of the member port even if they are different from those on the aggregate interface. GVRP and MSTP settings are examples of class-one configurations.

Reference port

When setting the aggregation state of the ports in an aggregation group, the system automatically picks a member port as the reference port. A Selected port must have the same port attributes and class-two configurations as the reference port.

LACP

The IEEE 802.3ad Link Aggregation Control Protocol (LACP) enables dynamic aggregation of physical links. It uses link aggregation control protocol data units (LACPDUs) for exchanging aggregation information between LACP-enabled devices.

1. LACP functions

The IEEE 802.3ad 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 LACP priority, port number, and operational key. Each member port in a LACP-enabled aggregation group exchanges the preceding information with its peer. When a member port receives an LACPDU, it compares the received information with the information received on the other member ports. In this way, the two systems reach an agreement on which ports should be placed in the Selected state.
Extended LACP functions	Implemented by extending the LACPDU with new Type/Length/Value (TLV) fields. This is how the LACP multi-active detection (MAD) mechanism of the Intelligent Resilient Framework (IRF) feature is implemented. An A5120 EI Switch Series can participate in LACP MAD as either an IRF member switch or an intermediate device.

NOTE:

For more information about IRF, member switches, intermediate devices, and the LACP MAD mechanism, see the *IRF Configuration Guide*.

2. LACP priorities

LACP priorities have two types: system LACP priority and port LACP priority.

Table 4 LACP priorities

Type	Description	Remarks
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 first, and then the system that has lower priority sets the port state accordingly.	The smaller the priority value, the higher the priority.
Port LACP priority	Determines the likelihood of a member port to be selected on a system. The higher the port LACP priority, the higher the likelihood.	

3. 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 fails to receive LACPDUs from the peer within three times the LACP timeout interval, the member port assumes that the peer port has failed. You can configure the LACP timeout interval as either the short timeout interval (1 second) or the long timeout interval (30 seconds).

Marker protocol

During a session, if member ports are added to or removed from a dynamic link aggregation group, service traffic must be redistributed among all the new member ports of the link aggregation group. The Marker protocol can be employed to quickly redistribute service traffic within link aggregation groups and ensure the orderly transmission of data frames. The process is:

- The device stops transmitting service traffic and starts a timer. No data frames will be transmitted on the links until the timer expires.
- The local end uses the Marker protocol to send a Marker Protocol Data Unit (PDU).
- When a Marker Response Protocol Data Unit (PDU) is received from the peer or the timer expires, the device starts to redistribute service traffic on all the new link aggregation member ports in the Selected state.

NOTE:

The A5120 EI Switch Series supports returning Marker Response PDUs only after dynamic link aggregation member ports receive Marker PDUs.

Link aggregation modes

Link aggregation has two modes: dynamic and static. Dynamic link aggregation uses LACP and static link aggregation does not. Table 5 compares the two aggregation modes.

Table 5 A comparison between static and dynamic aggregation modes

Aggregation mode	LACP status on member ports	Pros	Cons
Static	Disabled	Aggregation is stable. Peers do not affect the aggregation state of the member ports.	The member ports do not adjust the aggregation state according to that of the peer ports. The administrator must manually maintain link aggregations.
Dynamic	Enabled	The administrator does not need to maintain link aggregations. The peer systems maintain the aggregation state of the member ports automatically.	Aggregation is unstable. The aggregation state of the member ports is susceptible to network changes.

The following points apply to a dynamic link aggregation group:

- A Selected port can receive and send LACPDUs.
- An Unselected port can receive and send LACPDUs only if it is up and has the same class-two configurations as the aggregate interface.

Aggregating links in static mode

LACP is disabled on the member ports in a static aggregation group. You must manually maintain the aggregation state of the member ports.

The static link aggregation procedure comprises:

- Selecting a reference port
- Setting the aggregation state of each member port

Selecting a reference port

The system selects a reference port from the member ports that are:

- Are in the up state and have
- Have the same class-two configurations as the aggregate interface.

The candidate ports are sorted by duplex and speed in this order: full duplex/high speed, full duplex/low speed, half duplex/high speed, and half duplex/low speed. The one at the top is selected as the reference port. If two ports have the same duplex mode and speed, the one with the lower port number wins.

Setting the aggregation state of each member port

After selecting the reference port, the static aggregation group sets the aggregation state of each member port.

Figure 8 Setting the aggregation state of a member port in a static aggregation group

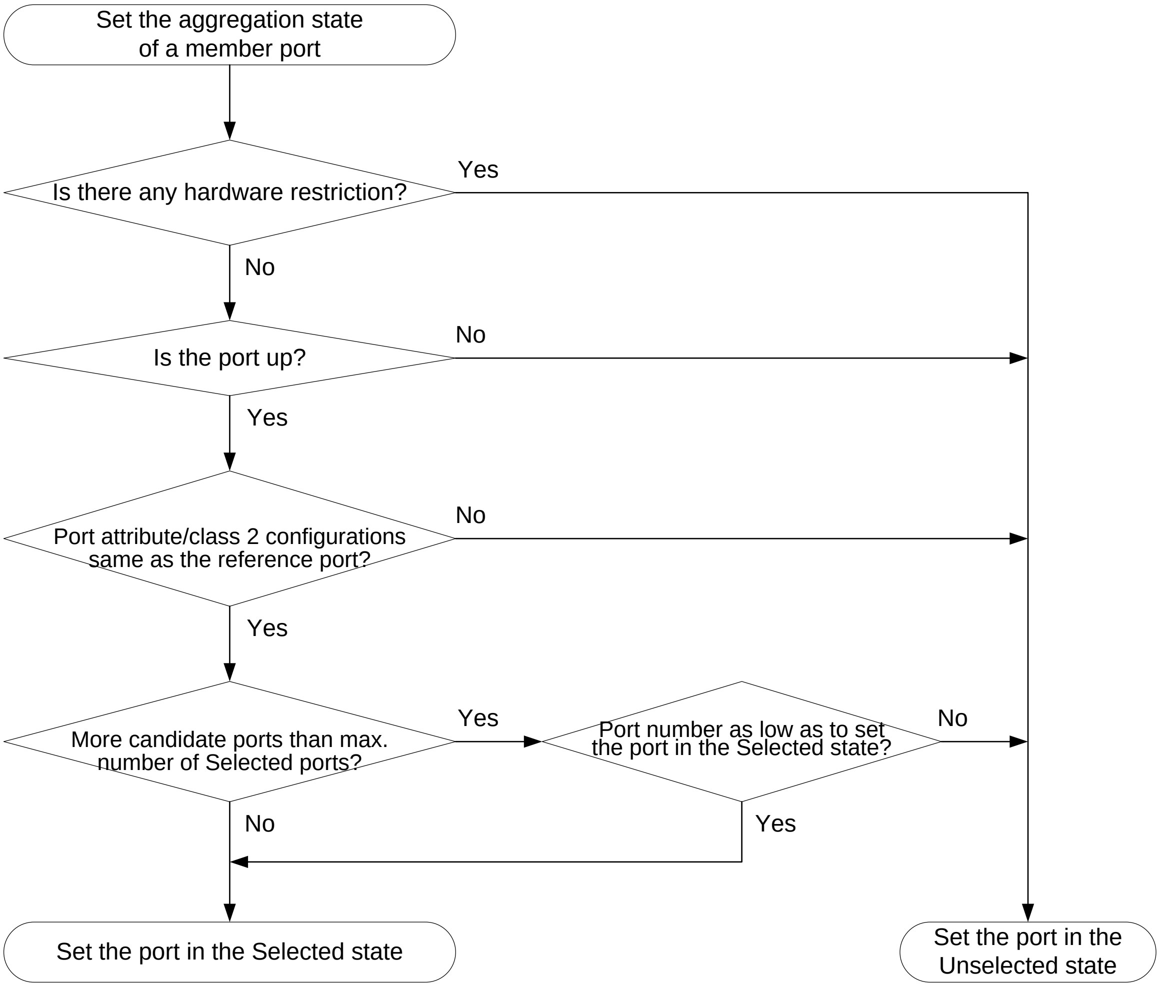
Set the aggregation state of a member port
Is there any hardware restriction?
Yes
No
Is the port up?
No
Yes
Port attribute/class 2 configurations same as the reference port?
No
Yes
More candidate ports than max. number of Selected ports?
Yes
Port number as low as to set the port in the Selected state?
No
No
Yes
Set the port in the Selected state
Set the port in the Unselected state

NOTE:

- To ensure stable aggregation state and service continuity, do not change port attributes or class-two configurations on any member port.
- If a static aggregation group has reached the limit on Selected ports, any port that joins the group is placed in the Unselected state to avoid traffic interruption on the current Selected ports. Avoid this situation, however, because it might cause the aggregation state of a port to change after a reboot.

Aggregating links in dynamic mode

LACP is automatically enabled on all member ports in a dynamic aggregation group. The protocol automatically maintains the aggregation state of ports.

The dynamic link aggregation procedure comprises:

- Selecting a reference port
- Setting the aggregation state of each member port

Selecting a reference port

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

1. The systems compare the system ID (which comprises the system LACP priority and the system MAC address). The system with the lower LACP priority value wins. If they are the same, the systems compare the system MAC addresses. The system with the lower MAC address wins.
2. The system with the smaller system ID selects the port with the smallest port ID as the reference port. A port ID comprises a port LACP priority and a port number. The port with the lower LACP priority value wins. If two ports have the same LACP priority, the system compares their port numbers. The port with the smaller port number wins.

Setting the aggregation state of each member port

After the reference port is selected, the system with the lower system ID sets the state of each member port in the dynamic aggregation group on its side.

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

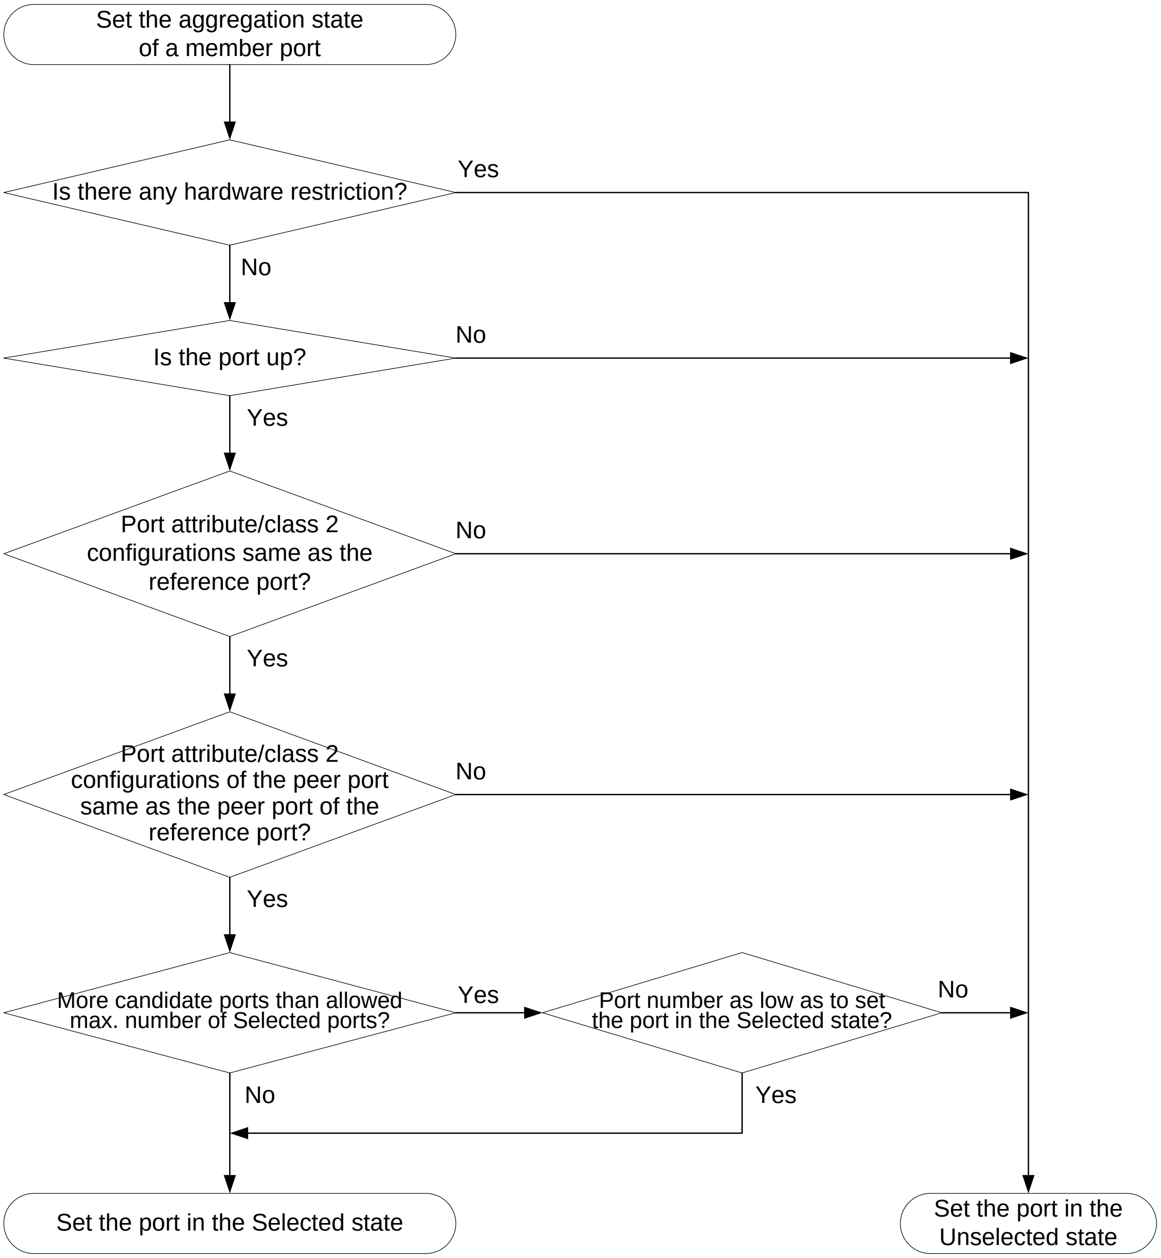
Set the aggregation state of a member port
Is there any hardware restriction?
Yes
No
Is the port up?
No
Yes
Port attribute/class 2 configurations same as the reference port?
No
Yes
Port attribute/class 2 configurations of the peer port same as the peer port of the reference port?
No
Yes
More candidate ports than allowed max. number of Selected ports?
Yes
Port number as low as to set the port in the Selected state?
No
No
Yes
Set the port in the Selected state
Set the port in the Unselected state

Meanwhile, the system with the higher system ID, which has identified the aggregation state changes on the remote system, sets the aggregation state of local member ports as the same as their peer ports.

NOTE:

- To ensure stable aggregation state and service continuity, do not change port attributes or class-two configurations on any member port.
- In a dynamic aggregation group, when the aggregation state of a local port changes, the aggregation state of the peer port changes.
- A port that joins a dynamic aggregation group after the Selected port limit has been reached will be placed in the Selected state if it is more eligible for being selected than a current member port.

Load-sharing criteria for link aggregation groups

In a link aggregation group, traffic can be load-shared across the selected member ports based on a set of criteria, depending on your configuration.

You can choose one of the following criteria or any combination of them for load sharing:

- MAC addresses
- IP addresses
- Service port numbers
- Receiving port numbers

Ethernet link aggregation configuration task list

Complete the following tasks to configure Ethernet link aggregation:

Task		Remarks
Configuring an aggregation group	Configuring a static aggregation group	Select either task
	Configuring a dynamic aggregation group	
Configuring an aggregate interface	Configuring the description of an aggregate interface	Optional
	Enabling link state traps for an aggregate interface	Optional
	Shutting down an aggregate interface	Optional
Configuring load sharing for link aggregation groups	Configuring load-sharing criteria for link aggregation groups	Optional
	Enabling local-first load sharing for link aggregation	Optional
Enabling link-aggregation traffic redirection		Optional

Configuring an aggregation group

Configuration guidelines

You cannot assign a port to a Layer 2 aggregation group if any of the features listed in Table 6 is configured on the port.

Table 6 Features incompatible with Layer 2 aggregation groups

Feature	Reference
RRPP	RRPP configuration in the *High Availability Configuration Guide*
MAC authentication	MAC authentication configuration in the *Security Configuration Guide*
Port security	Port security configuration in the *Security Configuration Guide*
IP source guard	IP source guard configuration in the *Security Configuration Guide*
802.1X	802.1X configuration in the *Security Configuration Guide*

NOTE:

To achieve better load sharing results for data traffic among the member ports of a link aggregation group, assign ports of the same type (such as all 100 Mbps ports or all GE ports and so on) to the link aggregation group.

△ CAUTION:

Removing an aggregate interface also removes the corresponding aggregation group. At the same time, all member ports leave the aggregation group.

Configuring a static aggregation group

NOTE:

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

Follow these steps to configure a Layer 2 static aggregation group:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Create a Layer 2 aggregate interface and enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	Required When you create a Layer 2 aggregate interface, the system automatically creates a Layer 2 static aggregation group numbered the same.
Exit to system view	**quit**	—
Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Required Repeat these two steps to assign multiple Layer 2 Ethernet interfaces to the aggregation group.
Assign the Ethernet interface to the aggregation group	**port link-aggregation group** *number*	

Configuring a dynamic aggregation group

NOTE:

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.

Follow these steps to configure a Layer 2 dynamic aggregation group:

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

To do...	Use the command...	Remarks
Set the system LACP priority	**lacp system-priority** *system-priority*	Optional By default, the system LACP priority is 32768. Changing the system LACP priority might affect the aggregation state of the ports in a dynamic aggregation group.
Create a Layer 2 aggregate interface and enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	Required When you create a Layer 2 aggregate interface, the system automatically creates a Layer 2 static aggregation group numbered the same.
Configure the aggregation group to work in dynamic aggregation mode	**link-aggregation mode dynamic**	Required By default, an aggregation group works in static aggregation mode.
Exit to system view	**quit**	—
Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Required Repeat these two steps to assign more Layer 2 Ethernet interfaces to the aggregation group.
Assign the Ethernet interface to the aggregation group	**port link-aggregation group** *number*	
Assign the port an LACP priority	**lacp port-priority** *port-priority*	Optional By default, the LACP priority of a port is 32768. Changing the LACP priority of a port might affect the aggregation state of the ports in the dynamic aggregation group.
Set the LACP timeout interval on the port to the short timeout interval (1 second)	**lacp period short**	Optional By default, the LACP timeout interval on a port is the long timeout interval (30 seconds).

Configuring an aggregate interface

You can perform the following configurations on an aggregate interface:

- Configuring the description of an aggregate interface
- Enabling link state traps for an aggregate interface
- Shutting down an aggregate interface

NOTE:

Most of the configurations that can be performed on Layer 2 Ethernet interfaces can also be performed on Layer 2 aggregate interfaces.

Configuring the description of an aggregate interface

You can configure the description of an aggregate interface for administration purposes such as describing the purpose of the interface.

Follow these steps to configure the description of an aggregate interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	—
Configure the description of the aggregate interface	**description** *text*	Optional By default, the description of an interface is *interface-name* **Interface**, such as **Bridge-Aggregation1 Interface**.

Enabling link state traps for an aggregate interface

You can configure an aggregate interface to generate linkUp trap messages when its link goes up and linkDown trap messages when its link goes down. For more information, see the *Network Management and Monitoring Configuration Guide*.

Follow these steps to enable link state traps on an aggregate interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enable the trap function globally	**snmp-agent trap enable** [**standard** [**linkdown** | **linkup**] *]	Optional By default, link state trapping is enabled globally and on all interfaces.
Enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	—
Enable link state traps for the aggregate interface	**enable snmp trap updown**	Optional Enabled by default.

Shutting down an aggregate interface

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

Follow these steps to shut down an aggregate interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—

To do...	Use the command...	Remarks
Enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	—
Shut down the aggregate interface	**shutdown**	Required By default, aggregate interfaces are up.

Configuring load sharing for link aggregation groups

Configuring load-sharing criteria for link aggregation groups

You can determine how traffic is load-shared across a link aggregation group by configuring load-sharing criteria. The criteria can be service port numbers, IP addresses, MAC addresses, receiving ports, or any combination.

The switch supports configuring global and group-specific aggregation load-sharing criteria. A link aggregation group preferentially uses group-specific load-sharing criteria. If no group-specific load-sharing criteria are available, the group uses the global load-sharing criteria.

Configuring the global link-aggregation load-sharing criteria

Follow these steps to configure global link-aggregation load-sharing criteria:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Configure the global link-aggregation load-sharing criteria	**link-aggregation load-sharing mode** { **destination-ip** | **destination-mac** | **destination-port** | **ingress-port** | **source-ip** | **source-mac** | **source-port** } *	Required By default, the global link-aggregation load-sharing criteria include the receiving port, source MAC address, and destination MAC address for Layer 2 packet types such as ARP, and the source and destination IP addresses for Layer 3 packet types such as IP packets.

You can set the following global aggregation load-sharing criteria:

- Source IP address
- Destination IP address
- Source MAC address
- Destination MAC address
- Source IP address and destination IP address
- Source IP address and source port number
- Destination IP address and destination port number

- Any two or all three of these elements – ingress port number, source MAC address, and destination MAC address

Configuring group-specific load-sharing criteria

Follow these steps to configure load-sharing criteria for a link aggregation group:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter aggregate interface view	**interface bridge-aggregation** *interface-number*	—
Configure the load-sharing criteria for the aggregation group	**link-aggregation load-sharing mode** { **destination-ip** | **destination-mac** | **source-ip** | **source-mac** } *	Required By default, an aggregation group uses the global link-aggregation load-sharing criteria.

You can set the following group-specific load-sharing criteria:

- Source IP address
- Destination IP address
- Source IP address and destination IP address
- Source MAC address
- Destination MAC address
- Destination MAC address and source MAC address

CAUTION:

By default, an aggregation group uses the global link-aggregation load sharing criteria. You can configure the group-specific link-aggregation load-sharing criteria to overwrite the global ones, except those specified with the **destination-port**, **source-port**, or **ingress-port** keywords.

Enabling local-first load sharing for link aggregation

Use the local-first load sharing mechanism in a cross-switch link aggregation scenario to distribute traffic preferentially across all member ports on the ingress switch rather than all member ports.

When you aggregate ports on different member switches in an IRF fabric, you can use local-first load sharing to reduce traffic on IRF links, as shown in Figure 10. For more information about IRF, see the *IRF Configuration Guide*.

Figure 10 Local-first link-aggregation load sharing

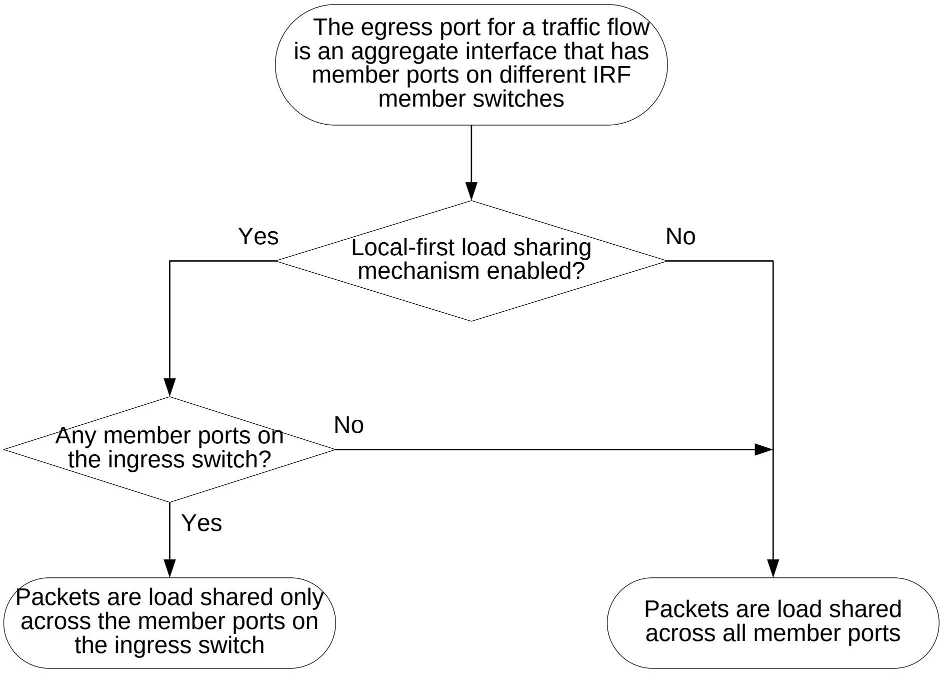
The egress port for a traffic flow is an aggregate interface that has member ports on different IRF member switches
Local-first load sharing mechanism enabled?
Yes
No
Any member ports on the ingress switch?
No
Yes
Packets are load shared only across the member ports on the ingress switch
Packets are load shared across all member ports

Follow these steps to enable local-first load sharing for link aggregation:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enable local-first load-sharing for link aggregation	**link-aggregation load-sharing mode local-first**	Optional Enabled by default.

Enabling link-aggregation traffic redirection

The link-aggregation traffic redirection function is available on IRF member switches. It can redirect traffic between IRF member switches for a cross-device link aggregation group. Link-aggregation traffic redirection prevents traffic interruption when you reboot an IRF member switch that contains link aggregation member ports. For more information about IRF, see the *IRF Configuration Guide*.

Follow these steps to enable link-aggregation traffic redirection:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enable link-aggregation traffic redirection	**link-aggregation lacp traffic-redirect-notification enable**	Optional Disabled by default.

CAUTION:

- Link-aggregation traffic redirection applies only to dynamic link aggregation groups.
- 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, disable both MSTP and link-aggregation traffic redirection.
- In an IRF fabric that adopts the ring connection, slight packet loss can occur when the IRF member switch enabled with link-aggregation traffic redirection reboots. To prevent packet loss, you can enable local-first load-sharing for link aggregation on all IRF member switches (see "Enabling local-first load sharing for link aggregation").

Displaying and maintaining Ethernet link aggregation

To do…	Use the command…	Remarks
Display information for an aggregate interface or multiple aggregate interfaces	**display interface bridge-aggregation** [**brief** [**down**]] [**|** { **begin** | **exclude** | **include** } *regular-expression*] **display interface bridge-aggregation** *interface-number* [**brief**] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the local system ID	**display lacp system-id** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the global or group-specific link-aggregation load-sharing criteria	**display link-aggregation load-sharing mode** [**interface** [**bridge-aggregation** *interface-number*]] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display detailed link aggregation information on link aggregation member ports	**display link-aggregation member-port** [*interface-list*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the summary of all aggregation groups	**display link-aggregation summary** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display detailed information about a specific or all aggregation groups	**display link-aggregation verbose** [**bridge-aggregation** [*interface-number*]] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Clear LACP statistics for a specific or all link aggregation member ports	**reset lacp statistics** [**interface** *interface-list*]	Available in user view
Clear statistics for a specific or all aggregate interfaces	**reset counters interface** [**bridge-aggregation** [*interface-number*]]	Available in user view

Ethernet link aggregation configuration examples

NOTE:

In an aggregation group, only ports that have the same port attributes and class-two configurations (see "Configuration classes") as the reference port (see "Reference port") can operate as Selected ports. You must ensure that all member ports have the same port attributes and class-two configurations as the reference port. The other settings only need to be configured on the aggregate interface, not on the member ports.

Layer 2 static aggregation configuration example

Network requirements

As shown in Figure 11:

- Device A and Device B are connected through their respective Layer 2 Ethernet interfaces GigabitEthernet 1/0/1 through GigabitEthernet 1/0/3.
- Configure a Layer 2 static link aggregation group 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, and VLAN 20 at one end to communicate with VLAN 20 at the other end.
- Enable traffic to be load-shared across aggregation group member ports based on source and destination MAC addresses.

Figure 11 Network diagram for Layer 2 static aggregation

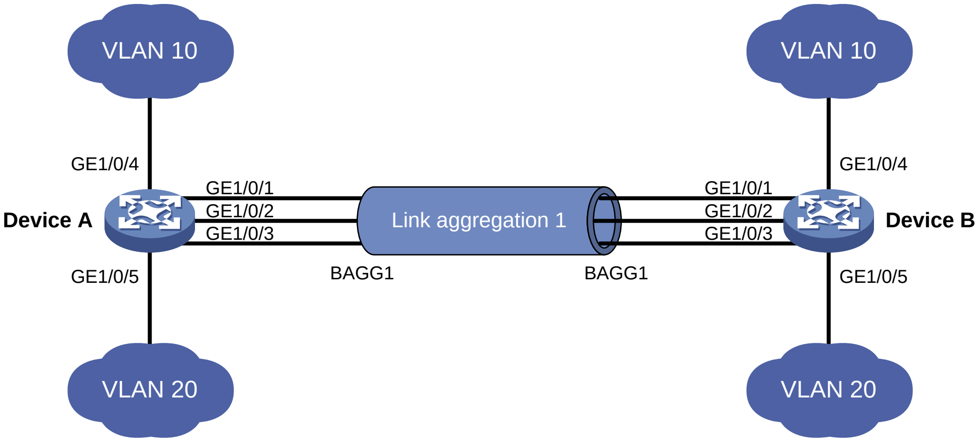

Configuration procedure

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.

NOTE:

This configuration automatically propagates to all the member ports in link aggregation group 1.

```
[DeviceA] interface bridge-aggregation 1
[DeviceA-Bridge-Aggregation1] port link-type trunk
[DeviceA-Bridge-Aggregation1] port trunk permit vlan 10 20
 Please wait... Done.
 Configuring GigabitEthernet1/0/1... Done.
 Configuring GigabitEthernet1/0/2... Done.
 Configuring GigabitEthernet1/0/3... Done.
[DeviceA-Bridge-Aggregation1] quit
```

# Configure the device to use the source and destination MAC addresses of packets as the global link-aggregation load-sharing criteria.

```
[DeviceA] link-aggregation load-sharing mode source-mac destination-mac
```

2. Configure Device B

Configure Device B as you configure Device A.

3. Verify the configurations

# Display the summary information about all aggregation groups on Device A.

```
[DeviceA] display link-aggregation summary

Aggregation Interface Type:
BAGG -- Bridge-Aggregation, RAGG -- Route-Aggregation
Aggregation Mode: S -- Static, D -- Dynamic
Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Actor System ID: 0x8000, 000f-e2ff-0001

AGG         AGG       Partner ID               Select Unselect   Share
Interface   Mode                               Ports  Ports      Type
-------------------------------------------------------------------------------
BAGG1       S         none                     3      0          Shar
```

The output shows that link aggregation group 1 is a load-shared Layer 2 static aggregation group, and it contains three Selected ports.

# Display the global link-aggregation load-sharing criteria on Device A.

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

Link-Aggregation Load-Sharing Mode:
  destination-mac address, source-mac address
```

The output shows that all link aggregation groups created on the device perform load sharing based on source and destination MAC addresses.

Layer 2 dynamic aggregation configuration example

Network requirements

As shown in Figure 12:

- Device A and Device B are connected through their respective Layer 2 Ethernet interfaces GigabitEthernet 1/0/1 through GigabitEthernet 1/0/3.
- Configure a Layer 2 dynamic link aggregation group 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, and VLAN 20 at one end to communicate with VLAN 20 at the other end.
- Enable traffic to be load-shared across aggregation group member ports based on source and destination MAC addresses.

Figure 12 Network diagram for Layer 2 dynamic aggregation

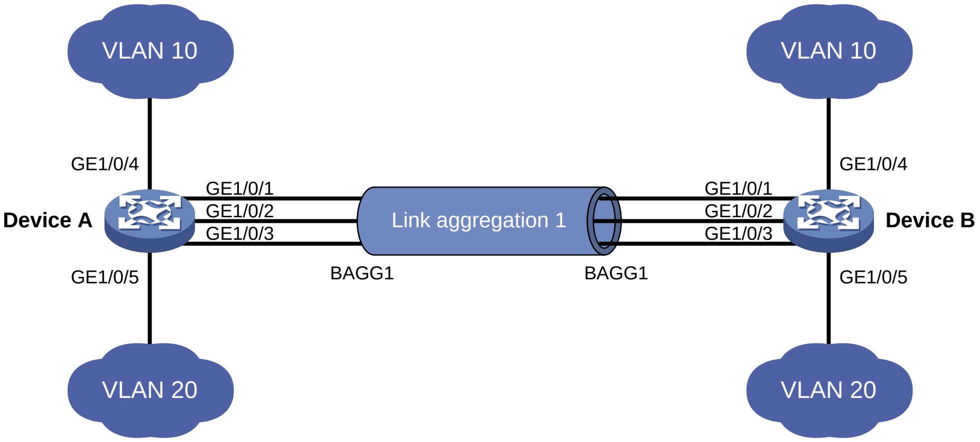

Configuration procedure

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, and configure the link aggregation mode as dynamic.

```
[DeviceA] interface bridge-aggregation 1
[DeviceA-Bridge-Aggregation1] link-aggregation mode dynamic
```

# Assign ports GigabitEthernet 1/0/1 through GigabitEthernet 1/0/3 to link aggregation group 1 one at a time.

```
[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.

NOTE:

This configuration automatically propagates to all the member ports in link aggregation group 1.

```
[DeviceA] interface bridge-aggregation 1
[DeviceA-Bridge-Aggregation1] port link-type trunk
[DeviceA-Bridge-Aggregation1] port trunk permit vlan 10 20
 Please wait... Done.
 Configuring GigabitEthernet1/0/1... Done.
 Configuring GigabitEthernet1/0/2... Done.
 Configuring GigabitEthernet1/0/3... Done.
[DeviceA-Bridge-Aggregation1] quit
```

# Configure the device to use the source and destination MAC addresses of packets as the global link-aggregation load-sharing criteria.

```
[DeviceA] link-aggregation load-sharing mode source-mac destination-mac
```

2. Configure Device B

Configure Device B as you configure Device A.

3. Verify the configurations

# Display the summary information about all aggregation groups on Device A.

```
[DeviceA] display link-aggregation summary

Aggregation Interface Type:
BAGG -- Bridge-Aggregation, RAGG -- Route-Aggregation
Aggregation Mode: S -- Static, D -- Dynamic
Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Actor System ID: 0x8000, 000f-e2ff-0001
```

```
AGG         AGG       Partner ID               Select Unselect   Share
Interface   Mode                               Ports  Ports      Type
-------------------------------------------------------------------------------
BAGG1       D         0x8000, 000f-e2ff-0002   3      0          Shar
```

The output shows that link aggregation group 1 is a load-shared Layer 2 dynamic aggregation group, and it contains three Selected ports.

# Display the global link-aggregation load-sharing criteria on Device A.

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

Link-Aggregation Load-Sharing Mode:
  destination-mac address, source-mac address
```

The output shows that all link aggregation groups created on the device perform load sharing based on source and destination MAC addresses.

Layer 2 aggregation load sharing configuration example

Network requirements

As shown in Figure 13:

- Device A and Device B are connected by their Layer 2 Ethernet interfaces GigabitEthernet 1/0/1 through GigabitEthernet 1/0/4.
- Configure two Layer 2 static link aggregation groups (1 and 2) on Device A and Device B respectively, and enable VLAN 10 at one end of the aggregate link to communicate with VLAN 10 at the other end through Bridge-Aggregation 1, and VLAN 20 at one end to communicate with VLAN 20 at the other end through Bridge-Aggregation 2.
- Configure the load sharing criterion for link aggregation group 1 as the source MAC addresses of packets and the load sharing criterion for link aggregation group 2 as the destination MAC addresses of packets to enable traffic to be load-shared across aggregation group member ports.

Figure 13 Network diagram for Layer 2 aggregation load sharing configuration

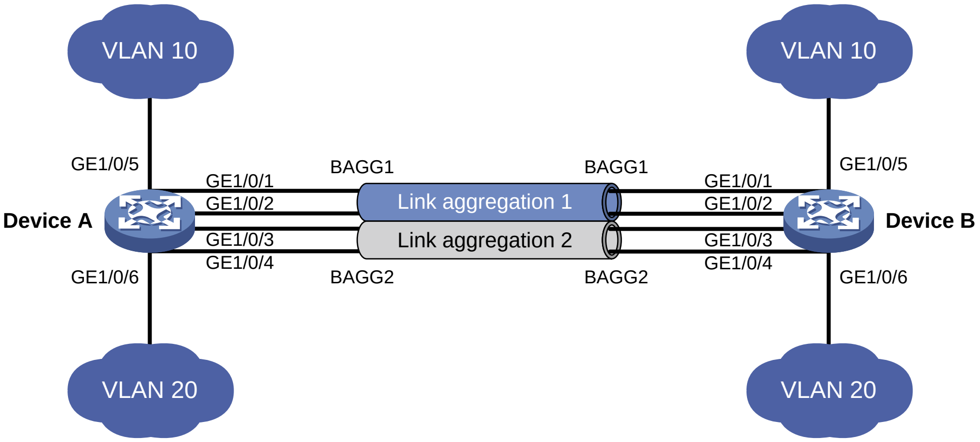

Configuration procedure

1. Configure Device A

# Create VLAN 10, and assign 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 port GigabitEthernet 1/0/6 to VLAN 20.

```
<DeviceA> system-view
[DeviceA] vlan 20
[DeviceA-vlan20] port gigabitethernet 1/0/6
[DeviceA-vlan20] quit
```

# Create Layer 2 aggregate interface Bridge-Aggregation 1, and configure the load sharing criterion for the link aggregation group as the source MAC addresses of packets.

```
[DeviceA] interface bridge-aggregation 1
[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 VLANs 10.

NOTE:

This configuration automatically propagates to all the member ports in link aggregation group 1.

```
[DeviceA] interface bridge-aggregation 1
[DeviceA-Bridge-Aggregation1] port link-type trunk
[DeviceA-Bridge-Aggregation1] port trunk permit vlan 10
 Please wait... Done.
 Configuring GigabitEthernet1/0/1... Done.
 Configuring GigabitEthernet1/0/2... Done.
[DeviceA-Bridge-Aggregation1] quit
```

# Create Layer 2 aggregate interface Bridge-Aggregation 2, and configure the load sharing criterion for the link aggregation group as the destination MAC addresses of packets.

```
[DeviceA] interface bridge-aggregation 2
[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 VLANs 20.

```
[DeviceA] interface bridge-aggregation 2
[DeviceA-Bridge-Aggregation2] port link-type trunk
[DeviceA-Bridge-Aggregation2] port trunk permit vlan 20
 Please wait... Done.
 Configuring GigabitEthernet1/0/3... Done.
 Configuring GigabitEthernet1/0/4... Done.
[DeviceA-Bridge-Aggregation2] quit
```

2. Configure Device B

Configure Device B as you configure Device A.

3. Verify the configurations

# Display the summary information about all aggregation groups on Device A.

```
[DeviceA] display link-aggregation summary

Aggregation Interface Type:
BAGG -- Bridge-Aggregation, RAGG -- Route-Aggregation
Aggregation Mode: S -- Static, D -- Dynamic
Loadsharing Type: Shar -- Loadsharing, NonS -- Non-Loadsharing
Actor System ID: 0x8000, 000f-e2ff-0001

AGG         AGG       Partner ID               Select Unselect   Share
Interface   Mode                               Ports  Ports      Type
--------------------------------------------------------------------------------
BAGG1       S         none                     2      0          Shar
BAGG2       S         none                     2      0          Shar
```

The output shows that link aggregation groups 1 and 2 are both load-sharing-capable Layer 2 static aggregation groups and each contains two Selected ports.

# Display all the group-specific load-sharing criteria 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 the load sharing criterion for link aggregation group 1 is the source MAC addresses of packets and that for link aggregation group 2 is the destination MAC addresses of packets.

Port isolation configuration

Introduction to port isolation

Assigning access ports to different VLANs is a typical way to isolate Layer 2 traffic for data privacy and security, but this approach is demanding on VLAN resources. To isolate Layer 2 traffic without using VLANs, HP introduced the port isolation feature.

To use the feature, you assign ports to a port isolation group. Ports in an isolation group are called "isolated ports." An isolated port does not forward any Layer 2 traffic to any other isolated port on the same switch, even if they are in the same VLAN. Still, an isolated port can communicate with any other port outside the isolation group, provided that they are in the same VLAN.

The A5120 EI Switch Series support one isolation group called "isolation group 1." This isolation group is automatically created and cannot be deleted. There is no limit on the number of member ports.

Configuring the isolation group

Follow these steps to assign a port to the isolation group:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Required Use one of the commands, as follows: • To assign an Ethernet port to the isolation group, enter Ethernet interface view. • To assign a Layer 2 aggregate interface to the isolation group, enter Layer 2 aggregate interface view. The subsequent configuration applies to both the Layer 2 aggregate interface and all its member ports. • To assign multiple Ethernet ports to the isolation group in bulk, enter port group view.
	Enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	
	Enter port group view	**port-group manual** *port-group-name*	
Assign the port or ports to the isolation group		**port-isolate enable**	Required The isolation group does not contain any ports by default.

NOTE:

If the switch fails to apply the **port-isolate enable** command to a Layer 2 aggregate interface, it does not assign any member port of the aggregate interface to the isolation group. If the failure occurs on a member port, the switch can still assign other member ports to the isolation group.

Displaying and maintaining isolation groups

To do…	Use the command…	Remarks
Display information about the isolation group	**display port-isolate group** [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

Port isolation configuration example

Network requirements

As shown in Figure 14:

- Hosts A, B, and C are connected to port GigabitEthernet 1/0/1, GigabitEthernet 1/0/2, and GigabitEthernet 1/0/3 of Device.
- Device is connected to the Internet through GigabitEthernet 1/0/4.
- GigabitEthernet 1/0/1, GigabitEthernet 1/0/2, GigabitEthernet 1/0/3, and GigabitEthernet 1/0/4 belong to the same VLAN.

Configure Device to enable Host A, Host B, and Host C to access the Internet when they are isolated from one another.

Figure 14 Network diagram for port isolation configuration

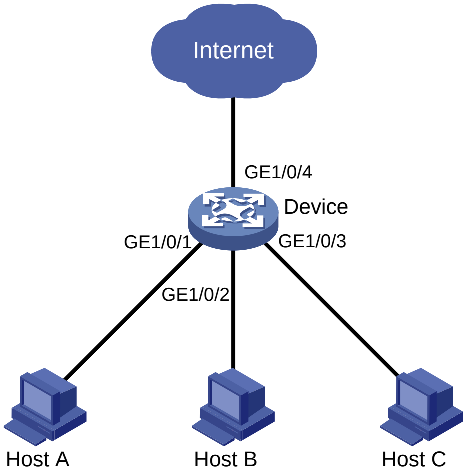

Configuration procedure

# Assign ports GigabitEthernet 1/0/1, GigabitEthernet 1/0/2, and GigabitEthernet 1/0/3 to isolation group 1.

```
<Device> system-view
[Device] interface gigabitethernet 1/0/1
[Device-GigabitEthernet1/0/1] port-isolate enable
[Device-GigabitEthernet1/0/1] quit
[Device] interface gigabitethernet 1/0/2
[Device-GigabitEthernet1/0/2] port-isolate enable
[Device-GigabitEthernet1/0/2] quit
[Device] interface gigabitethernet 1/0/3
```

```
[Device-GigabitEthernet1/0/3] port-isolate enable
```

# Display information about the isolation group.

```
<Device> display port-isolate group
 Port-isolate group information:
 Uplink port support: NO
 Group ID: 1
 Group members:
    GigabitEthernet1/0/1     GigabitEthernet1/0/2     GigabitEthernet1/0/3
```

MSTP configuration

As a Layer 2 management protocol, the Spanning Tree Protocol (STP) eliminates Layer 2 loops by selectively blocking redundant links in a network, putting them in a standby state, which still allows for link redundancy.

The recent versions of STP are the Rapid Spanning Tree Protocol (RSTP) and the Multiple Spanning Tree Protocol (MSTP). This document describes the features of STP, RSTP, and MSTP.

Introduction to STP

Why STP

STP was developed based on the 802.1d standard of IEEE to eliminate loops at the data link layer in a local area network (LAN). Networks often have redundant links as backups in case of failures, but loops are a very serious problem. Devices that run this protocol detect loops in the network by exchanging information with one another, and eliminate loops by selectively blocking certain ports to prune the loop structure into a loop-free tree structure. This avoids proliferation and infinite cycling of packets that would occur in a loop network, and prevents received duplicate packets from decreasing the performance of network devices.

In the narrow sense, STP refers to IEEE 802.1d STP. In the broad sense, STP refers to the IEEE 802.1d STP and various enhanced spanning tree protocols derived from that protocol.

Protocol packets of STP

STP uses bridge protocol data units (BPDUs), also known as "configuration messages", as its protocol packets.

STP-enabled network devices exchange BPDUs to establish a spanning tree. BPDUs contain sufficient information for the network devices to complete spanning tree calculation.

In STP, BPDUs have the following types:

- Configuration BPDUs, used by network devices to calculate a spanning tree and maintain the spanning tree topology
- Topology change notification (TCN) BPDUs, which notify network devices of the network topology changes

Configuration BPDUs contain sufficient information for the network devices to complete spanning tree calculation. Important fields in a configuration BPDU include the following:

- Root bridge ID: Consisting of the priority and MAC address of the root bridge.
- Root path cost: Cost of the path to the root bridge denoted by the root identifier from the transmitting bridge.
- Designated bridge ID: Consisting of the priority and MAC address of the designated bridge.
- Designated port ID: Consisting of the priority and global port number of the designated port.
- Message age: Age of the configuration BPDU while it propagates in the network.

- Max age: Maximum age of the configuration BPDU stored on the switch.
- Hello time: Configuration BPDU transmission interval.
- Forward delay: Delay that STP bridges use to transition port state.

Basic concepts in STP

Root bridge

A tree network must have a root bridge. The entire network contains only one root bridge. The root bridge is not permanent, but can change along with changes of the network topology.

Upon initialization of a network, each device generates and periodically sends configuration BPDUs, with itself as the root bridge. After network convergence, only the root bridge generates and periodically sends configuration BPDUs, and the other devices forward the BPDUs.

Root port

On a non-root bridge, the port nearest to the root bridge is the root port. The root port is responsible for communication with the root bridge. Each non-root bridge has only one root port. The root bridge has no root port.

Designated bridge and designated port

Table 7 Description of designated bridges and designated ports

Classification	Designated bridge	Designated port
For a device	A device directly connected to the local device and responsible for forwarding BPDUs to the local device	The port through which the designated bridge forwards BPDUs to this device
For a LAN	The device responsible for forwarding BPDUs to this LAN segment	The port through which the designated bridge forwards BPDUs to this LAN segment

As shown in Figure 15, both Device B and Device C directly connect to the LAN. If Device A forwards BPDUs to Device B through port A1, the designated bridge for Device B is Device A, and the designated port of Device B is port A1 on Device A. If Device B forwards BPDUs to the LAN, the designated bridge for the LAN is Device B, and the designated port for the LAN is port B2 on Device B.

Figure 15 A schematic diagram of designated bridges and designated ports

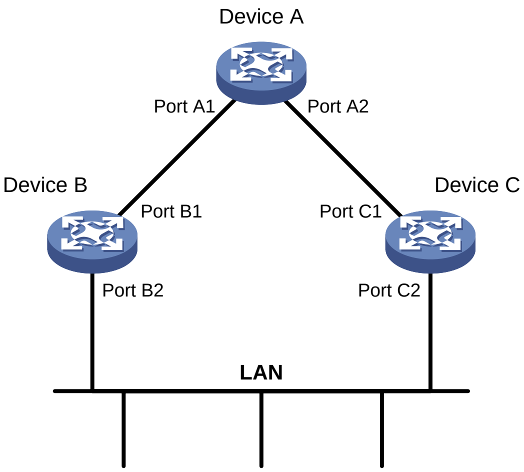

Path cost

Path cost is a reference value used for link selection in STP. STP calculates path costs to select the most robust links and block redundant links that are less robust, to prune the network into a loop-free tree.

How STP works

NOTE:

The spanning tree calculation process described in the following sections is a simplified process for example only.

STP has the following workflow:

1. Initial state

Upon initialization of a device, each port generates a BPDU with the device as the root bridge, in which the root path cost is 0, the designated bridge ID is the device ID, and the designated port is the port itself.

2. Selection of the root bridge

Initially, each STP device on the network assumes itself to be the root bridge, with its own device ID as the root bridge ID. By exchanging configuration BPDUs, the devices compare their root bridge IDs to elect the device with the smallest root bridge ID as the root bridge.

3. Selection of the root port and designated ports

Table 8 Selection of the root port and designated ports

Step	Description
1	A non-root device regards the port on which it received the optimum configuration BPDU as the root port. For the selection of the optimum configuration BPDUs, see Table 9.

Step	Description
2	Based on the configuration BPDU and the path cost of the root port, the device calculates a designated port configuration BPDU for each of the other ports. • The root bridge ID is replaced with that of the configuration BPDU of the root port. • The root path cost is replaced with that of the configuration BPDU of the root port plus the path cost of the root port. • The designated bridge ID is replaced with the ID of this device. • The designated port ID is replaced with the ID of this port.
3	The device compares the calculated configuration BPDU with the configuration BPDU on the port whose port role will be defined, and acts depending on the result of the comparison. • If the calculated configuration BPDU is superior, the device considers this port as the designated port, replaces the configuration BPDU on the port with the calculated configuration BPDU, and periodically sends the calculated configuration BPDU. • If the configuration BPDU on the port is superior, the device blocks this port without updating its configuration BPDU. The blocked port can receive BPDUs but not send BPDUs or forward data traffic.

NOTE:

When the network topology is stable, only the root port and designated ports forward traffic, and other ports are all in the blocked state in which the port receive BPDUs but do not forward BPDUs or user traffic.

Table 9 Selection of the optimum configuration BPDU

Step	Actions
1	Upon receiving a configuration BPDU on a port, the device performs the following: • If the received configuration BPDU has a lower priority than that of the configuration BPDU generated by the port, the device discards the received configuration BPDU and keeps the configuration BPDU this port generated. • If the received configuration BPDU has a higher priority than that of the configuration BPDU generated by the port, the device replaces the content of the configuration BPDU generated by the port with the content of the received configuration BPDU.
2	The device compares the configuration BPDUs of all the ports and chooses the optimum configuration BPDU.

NOTE:

The following are the principles of configuration BPDU comparison:

- The configuration BPDU with the lowest root bridge ID has the highest priority.
- If all configuration BPDUs have the same root bridge ID, their root path costs are compared. For example, the root path cost in a configuration BPDU plus the path cost of a receiving port is S. The configuration BPDU with the smallest S value has the highest priority.
- If all configuration BPDUs have the same root path cost, their designated bridge IDs, designated port IDs, and the IDs of the receiving ports are compared in sequence. The configuration BPDU that contains the smallest ID wins.

A tree topology forms upon successful election of the root bridge, the root port on each non-root bridge and the designated ports.

Figure 16 provides an example of how the STP algorithm works.

Figure 16 Network diagram for the STP algorithm

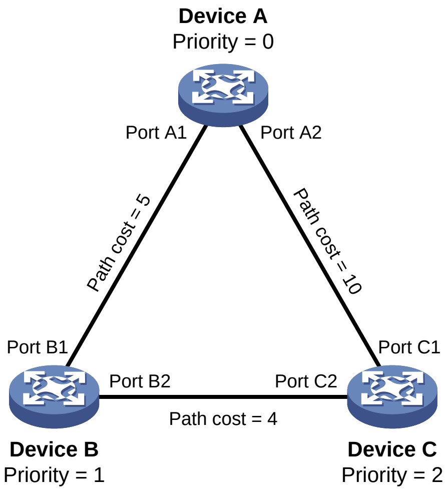

As shown in Figure 16, the priority of Device A, Device B, and Device C is 0, 1, and 2 respectively, and the path costs among these links are 5, 10, and 4 respectively.

4. Initial state of each device

Table 10 Initial state of each device

Device	Port name	Configuration BPDU on the port
Device A	Port A1	{0, 0, 0, Port A1}
	Port A2	{0, 0, 0, Port A2}
Device B	Port B1	{1, 0, 1, Port B1}
	Port B2	{1, 0, 1, Port B2}
Device C	Port C1	{2, 0, 2, Port C1}
	Port C2	{2, 0, 2, Port C2}

NOTE:

In Table 10, each configuration BPDU contains the following fields: root bridge ID, root path cost, designated bridge ID, and designated port ID.

5. Comparison process and result on each device

Table 11 Comparison process and result on each device

Device	Comparison process	Configuration BPDU on ports after comparison
Device A	• Port A1 receives the configuration BPDU of Port B1 {1, 0, 1, Port B1}, finds that its existing configuration BPDU {0, 0, 0, Port A1} is superior to the received configuration BPDU, and discards the received one. • Port A2 receives the configuration BPDU of Port C1 {2, 0, 2, Port C1}, finds that its existing configuration BPDU {0, 0, 0, Port A2} is superior to the received configuration BPDU, and discards the received one. • Device A finds that it is both the root bridge and designated bridge in the configuration BPDUs of all its ports, and considers itself as the root bridge. It does not change the configuration BPDU of any port and starts to periodically send configuration BPDUs.	• Port A1: {0, 0, 0, Port A1} • Port A2: {0, 0, 0, Port A2}
Device B	• Port B1 receives the configuration BPDU of Port A1 {0, 0, 0, Port A1}, finds that the received configuration BPDU is superior to its existing configuration BPDU {1, 0, 1, Port B1}, and updates its configuration BPDU. • Port B2 receives the configuration BPDU of Port C2 {2, 0, 2, Port C2}, finds that its existing configuration BPDU {1, 0, 1, Port B2} is superior to the received configuration BPDU, and discards the received one.	• Port B1: {0, 0, 0, Port A1} • Port B2: {1, 0, 1, Port B2}
	• Device B compares the configuration BPDUs of all its ports, decides that the configuration BPDU of Port B1 is the optimum, and selects Port B1 as the root port with the configuration BPDU unchanged. • Based on the configuration BPDU and path cost of the root port, Device B calculates a designated port configuration BPDU for Port B2 {0, 5, 1, Port B2}, and compares it with the existing configuration BPDU of Port B2 {1, 0, 1, Port B2}. Device B finds that the calculated one is superior, decides that Port B2 is the designated port, replaces the configuration BPDU on Port B2 with the calculated one, and periodically sends the calculated configuration BPDU.	• Root port (Port B1): {0, 0, 0, Port A1} • Designated port (Port B2): {0, 5, 1, Port B2}
Device C	• Port C1 receives the configuration BPDU of Port A2 {0, 0, 0, Port A2}, finds that the received configuration BPDU is superior to its existing configuration BPDU {2, 0, 2, Port C1}, and updates its configuration BPDU. • Port C2 receives the original configuration BPDU of Port B2 {1, 0, 1, Port B2}, finds that the received configuration BPDU is superior to the existing configuration BPDU {2, 0, 2, Port C2}, and updates its configuration BPDU.	• Port C1: {0, 0, 0, Port A2} • Port C2: {1, 0, 1, Port B2}

Device	Comparison process	Configuration BPDU on ports after comparison
	• Device C compares the configuration BPDUs of all its ports, decides that the configuration BPDU of Port C1 is the optimum, and selects Port C1 as the root port with the configuration BPDU unchanged. • Based on the configuration BPDU and path cost of the root port, Device C calculates the configuration BPDU of Port C2 {0, 10, 2, Port C2}, and compares it with the existing configuration BPDU of Port C2 {1, 0, 1, Port B2}. Device C finds that the calculated configuration BPDU is superior to the existing one, selects Port C2 as the designated port, and replaces the configuration BPDU of Port C2 with the calculated one.	• Root port (Port C1): {0, 0, 0, Port A2} • Designated port (Port C2): {0, 10, 2, Port C2}
	• Port C2 receives the updated configuration BPDU of Port B2 {0, 5, 1, Port B2}, finds that the received configuration BPDU is superior to its existing configuration BPDU {0, 10, 2, Port C2}, and updates its configuration BPDU. • Port C1 receives a periodic configuration BPDU {0, 0, 0, Port A2} from Port A2, finds that it is the same as the existing configuration BPDU, and discards the received one.	• Port C1: {0, 0, 0, Port A2} • Port C2: {0, 5, 1, Port B2}
	• Device C finds that the root path cost of Port C1 (10) (root path cost of the received configuration BPDU (0) plus path cost of Port C1 (10)) is larger than that of Port C2 (9) (root path cost of the received configuration BPDU (5) plus path cost of Port C2 (4)), decides that the configuration BPDU of Port C2 is the optimum, and selects Port C2 as the root port with the configuration BPDU unchanged. • Based on the configuration BPDU and path cost of the root port, Device C calculates a designated port configuration BPDU for Port C1 {0, 9, 2, Port C1} and compares it with the existing configuration BPDU of Port C1 {0, 0, 0, Port A2}. Device C finds that the existing configuration BPDU is superior to the calculated one and blocks Port C1 with the configuration BPDU unchanged. Then Port C1 does not forward data until a new event triggers a spanning tree calculation process, for example, the link between Device B and Device C is down.	• Blocked port (Port C1): {0, 0, 0, Port A2} • Root port (Port C2): {0, 5, 1, Port B2}

NOTE:

In Table 11, each configuration BPDU contains the following fields: root bridge ID, root path cost, designated bridge ID, and designated port ID.

After the comparison processes described in Table 11, a spanning tree with Device A as the root bridge is established, and the topology is shown in Figure 17.

Figure 17 Topology of the final calculated spanning tree

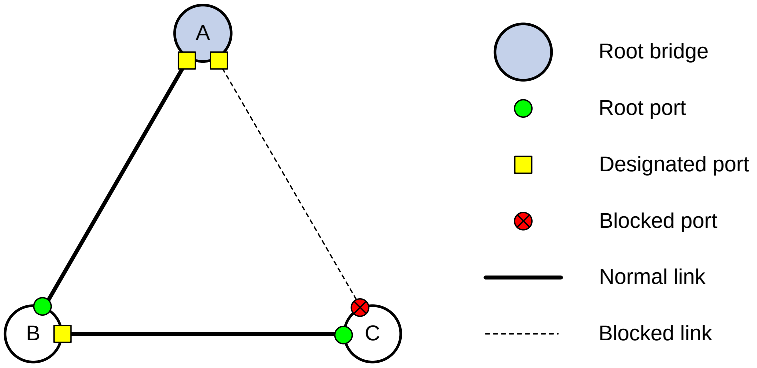

The BPDU forwarding mechanism in STP

STP forwards configuration BPDUs following these guidelines:

- Upon network initiation, every switch regards itself as the root bridge, generates configuration BPDUs with itself as the root, and sends the configuration BPDUs at a regular hello interval.
- If the root port received a configuration BPDU and the received configuration BPDU is superior to the configuration BPDU of the port, the device increases the message age carried in the configuration BPDU following a certain rule and starts a timer to time the configuration BPDU while sending this configuration BPDU through the designated port.
- If the configuration BPDU received on a designated port has a lower priority than the configuration BPDU of the local port, the port immediately sends its own configuration BPDU in response.
- If a path becomes faulty, the root port on this path no longer receives new configuration BPDUs and the old configuration BPDUs will be discarded because of timeout. The device generates a configuration BPDU with itself as the root and sends the BPDUs and TCN BPDUs. This triggers a new spanning tree calculation process to establish a new path to restore the network connectivity.

However, the newly calculated configuration BPDU cannot be propagated throughout the network immediately, so the old root ports and designated ports that have not detected the topology change continue forwarding data along the old path. If the new root ports and designated ports begin to forward data as soon as they are elected, a temporary loop might occur.

STP timers

The most important timing parameters in STP calculation are forward delay, hello time, and max age.

- Forward delay: Specifies the delay time for port state transition. A path failure can cause spanning tree re-calculation to adapt the spanning tree structure to the change. However, the resulting new configuration BPDU cannot propagate throughout the network immediately. If the newly elected root ports and designated ports start to forward data immediately, a temporary loop will likely occur. For this reason, as a mechanism for state transition in STP, the newly elected root ports or designated ports require twice the forward delay time before they transit to the forwarding state to ensure that the new configuration BPDU has propagated throughout the network.
- Hello time: Specifies the time interval at which a device sends hello packets to the surrounding devices to ensure that the paths are fault-free.
- Max age: Determines whether a configuration BPDU held by the device has expired. A configuration BPDU beyond the max age is discarded.

Introduction to RSTP

Developed based on the 802.1w standard of IEEE, RSTP is an optimized version of STP. It achieves rapid network convergence by allowing a newly elected root port or designated port to enter the forwarding state much faster under certain conditions than STP.

NOTE:

- In RSTP, a newly elected root port can enter the forwarding state rapidly if the old root port on the device has stopped forwarding data and the upstream designated port has started forwarding data.
- In RSTP, a newly elected designated port can enter the forwarding state rapidly if the designated port is an edge port (a port that directly connects to a user terminal rather than to another device or a shared LAN segment) or a port connected to a point-to-point link. If the designated port is an edge port, it can enter the forwarding state directly. If the designated port is connected to a point-to-point link, it can enter the forwarding state immediately after the device undergoes handshake with the downstream device and gets a response.

Introduction to MSTP

Why MSTP

Limitations of STP and RSTP

STP does not support rapid state transition of ports. A newly elected root port or designated port must wait twice the forward delay time before it transits to the forwarding state, even if it is a port on a point-to-point link or an edge port.

Although RSTP supports rapid network convergence, it has the same drawback as STP—All bridges within a LAN share the same spanning tree, so redundant links cannot be blocked based on VLAN, and the packets of all VLANs are forwarded along the same spanning tree.

Features of MSTP

Developed based on IEEE 802.1s, MSTP overcomes the limitations of STP and RSTP. In addition to supporting for rapid network convergence, it provides a better load sharing mechanism for redundant links by allowing data flows of different VLANs to be forwarded along separate paths. For more information about VLANs, see the chapter “VLAN configuration.”

MSTP provides the following features:

- MSTP supports mapping VLANs to spanning tree instances by means of a VLAN-to-instance mapping table. MSTP can reduce communication overheads and resource usage by mapping multiple VLANs to one instance.
- MSTP divides a switched network into multiple regions, each of which contains multiple spanning trees that are independent of one another.
- MSTP prunes a loop network into a loop-free tree, which avoids proliferation and endless cycling of packets in a loop network. In addition, it supports load balancing of VLAN data by providing multiple redundant paths for data forwarding.
- MSTP is compatible with STP and RSTP.

Basic concepts in MSTP

Figure 18 Basic concepts in MSTP

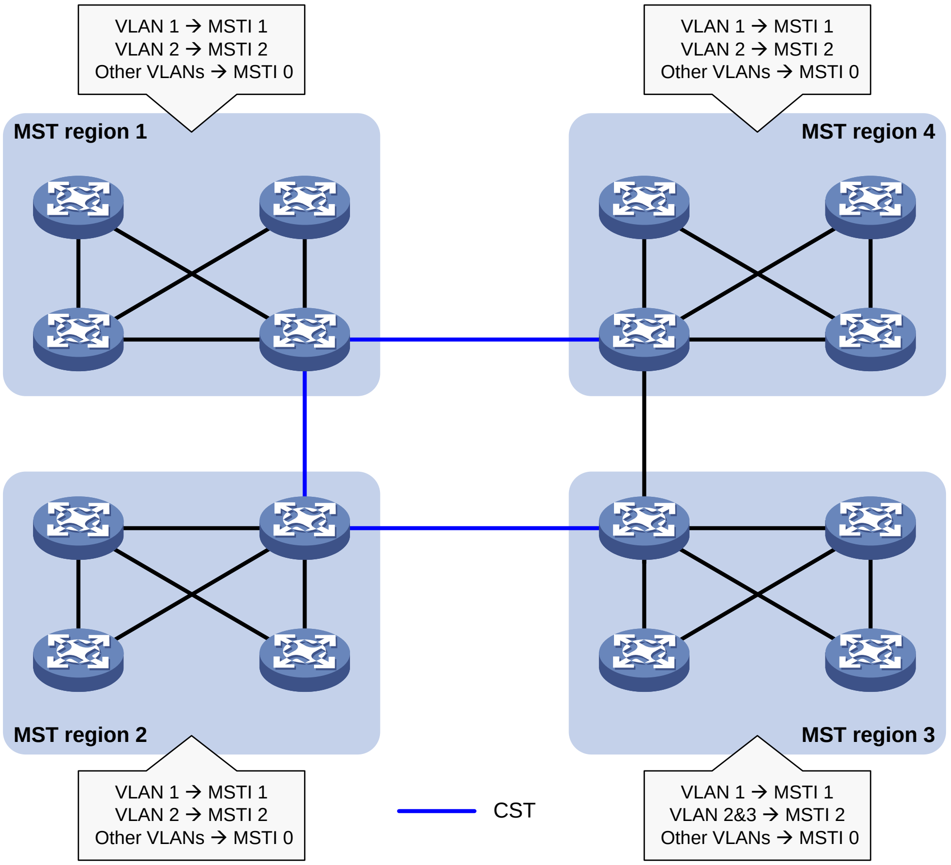

Figure 19 Network diagram and topology of MST region 3

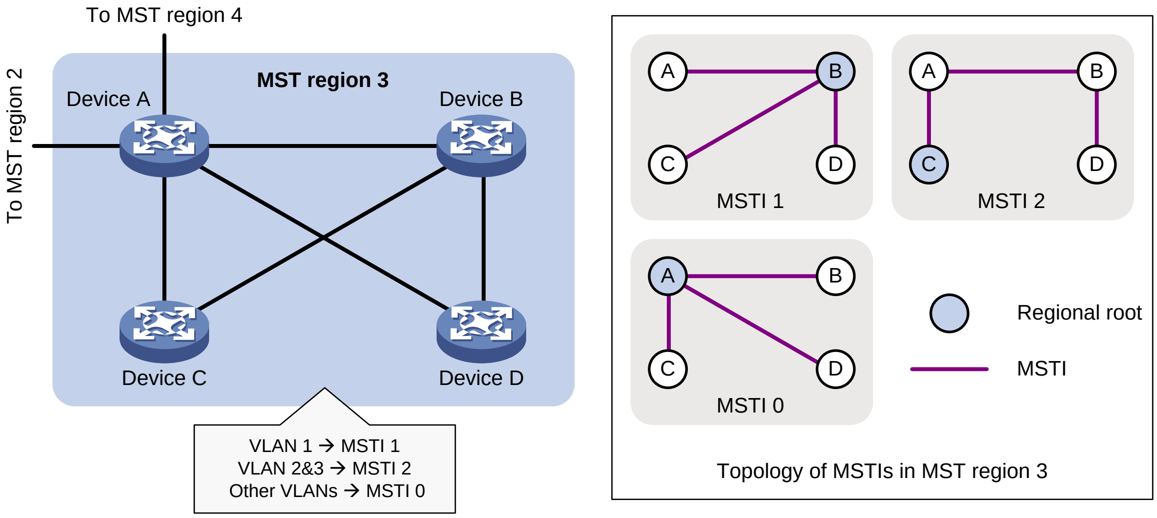

As shown in Figure 18, a switched network comprises four MST regions, and each MST region comprises four devices running MSTP. Figure 19 shows the networking topology of MST region 3.

MST region

A multiple spanning tree region (MST region) consists of multiple devices in a switched network and the network segments among them. All these devices have the following characteristics:

- MSTP-enabled
- Same region name
- Same VLAN-to-instance mapping configuration
- Same MSTP revision level configuration
- Physically linked with one another

Multiple MST regions can exist in a switched network. You can assign multiple devices to the same MST region. In Figure 18, the switched network comprises four MST regions, MST region 1 through MST region 4, and all devices in each MST region have the same MST region configuration.

MSTI

MSTP can generate multiple independent spanning trees in an MST region, and each spanning tree is mapped to the specific VLANs. Each spanning tree is referred to as a "multiple spanning tree instance (MSTI)."

In Figure 19, for example, MST region 3 comprises three MSTIs, MSTI 1, MSTI 2, and MSTI 0.

VLAN-to-instance mapping table

As an attribute of an MST region, the VLAN-to-instance mapping table describes the mapping relationships between VLANs and MSTIs.

In Figure 19, for example, the VLAN-to-instance mapping table of MST region 3 is: VLAN 1 to MSTI 1, VLAN 2 and VLAN 3 to MSTI 2, and other VLANs to MSTI 0. MSTP achieves load balancing by means of the VLAN-to-instance mapping table.

CST

The common spanning tree (CST) is a single spanning tree that connects all MST regions in a switched network. If you regard each MST region as a device, the CST is a spanning tree calculated by these devices through STP or RSTP.

For example, the blue lines in Figure 18 represent the CST.

IST

An internal spanning tree (IST) is a spanning tree that runs in an MST region. It is a special MSTI, and is also called "MSTI 0." All VLANs are mapped to MSTI 0 by default. As shown in Figure 18, MSTI 0 is the IST in MST region 3.

CIST

Jointly constituted by ISTs and the CST, the common and internal spanning tree (CIST) is a single spanning tree that connects all devices in a switched network. ISTs in all MST regions and the CST jointly constitute the CIST of the entire network. In Figure 18, for example, the ISTs in all MST regions plus the inter-region CST constitute the CIST of the entire network.

Regional root

The root bridge of the IST or an MSTI within an MST region is the regional root of the IST or MSTI. Based on the topology, different spanning trees in an MST region might have different regional roots.

For example, in MST region 3 in Figure 19, the regional root of MSTI 1 is Device B, the regional root of MSTI 2 is Device C, and the regional root of MSTI 0 (also known as the IST) is Device A.

Common root bridge

The common root bridge is the root bridge of the CIST.

In Figure 18, for example, the common root bridge is a device in MST region 1.

Roles of ports

A port can play different roles in different MSTIs. As shown in Figure 20, an MST region comprises Device A, Device B, Device C, and Device D. Port A1 and port A2 of Device A connect to the common root bridge. Port B2 and Port B3 of Device B form a loop. Port C3 and Port C4 of Device C connect to other MST regions. Port D3 of Device D directly connects to a host.

Figure 20 Port roles

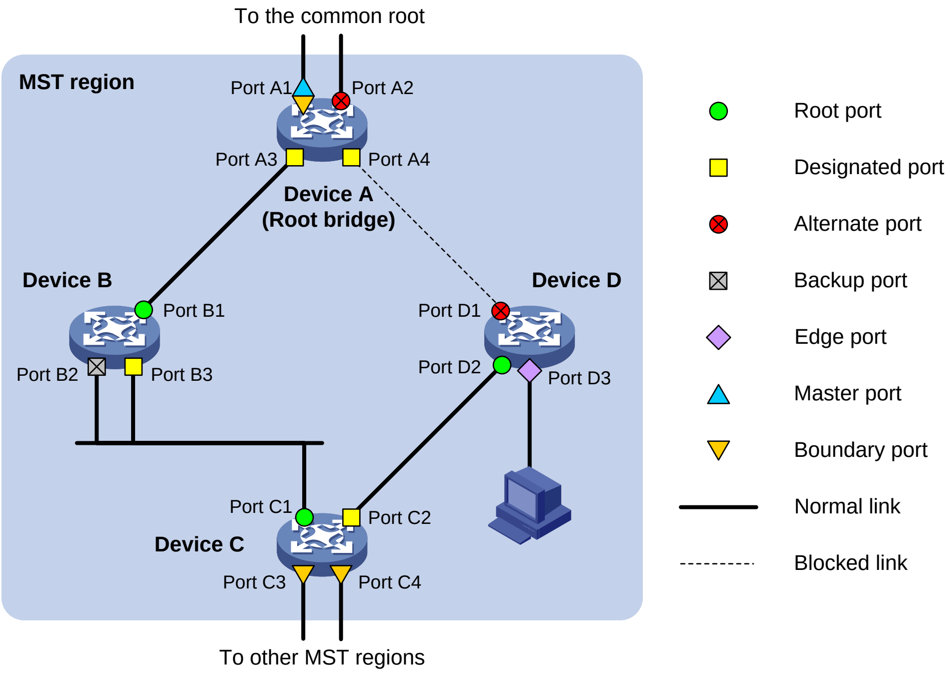

MSTP calculation involves the following port roles:

- Root port: Forwards data for a non-root bridge to the root bridge. The root bridge does not have any root port.
- Designated port: Forwards data to the downstream network segment or device.
- Alternate port: The backup port for a root port or master port. When the root port or master port is blocked, the alternate port takes over.
- Backup port: The backup port of a designated port. When the designated port fails, the backup port takes over. When a loop occurs because of the interconnection of two ports of the same MSTP device, the device blocks either of the two ports, and the blocked port is the backup port.
- Edge port: An edge port does not connect to any network device or network segment, but directly connects to a user host.
- Master port: A port on the shortest path from the local MST region to the common root bridge. The master port is a root port on the IST or CIST and still a master port on the other MSTIs.
- Boundary port: Connects an MST region to another MST region or to an STP/RSTP-running device. In MSTP calculation, a boundary port's role on an MSTI is consistent with its role on the CIST. But that is not true with master ports. A master port on MSTIs is a root port on the CIST.

Port states

In MSTP, a port can be in one of the following states:

- Forwarding: The port receives and sends BPDUs, obtains MAC addresses, and forwards user traffic.
- Learning: The port receives and sends BPDUs, obtains MAC addresses, but does not forward user traffic. Learning is an intermediate port state.
- Discarding: The port receives and sends BPDUs, but does not obtain MAC addresses or forward user traffic.

NOTE:

When in different MSTIs, a port can be in different states.

A port state is not exclusively associated with a port role. Table 12 lists the port states that each port role supports. (A check mark [√] indicates that the port supports this state, while a dash [—] indicates that the port does not support this state.)

Table 12 Port states that different port roles support

Port role (right) **Port state (below)**	**Root port/master port**	**Designated port**	**Alternate port**	**Backup port**
Forwarding	√	√	—	—
Learning	√	√	—	—
Discarding	√	√	√	√

How MSTP works

MSTP divides an entire Layer 2 network into multiple MST regions, which are connected by a calculated CST. Inside an MST region, multiple spanning trees are calculated. Each spanning tree is an MSTI. Among these MSTIs, MSTI 0 is the IST. Like STP, MSTP uses configuration BPDUs to calculate spanning trees. An important difference is that an MSTP BPDU carries the MSTP configuration on the bridge from which the BPDU is sent.

CIST calculation

The calculation of a CIST tree is also the process of configuration BPDU comparison. During this process, the device with the highest priority is elected as the root bridge of the CIST. MSTP generates an IST within each MST region through calculation. At the same time, MSTP regards each MST region as a single device and generates a CST among these MST regions through calculation. The CST and ISTs constitute the CIST of the entire network.

MSTI calculation

Within an MST region, MSTP generates different MSTIs for different VLANs based on the VLAN-to-instance mappings. For each spanning tree, MSTP performs a separate calculation process, which is similar to spanning tree calculation in STP. For more information, see "How STP works."

In MSTP, a VLAN packet is forwarded along the following paths:

- Within an MST region, the packet is forwarded along the corresponding MSTI.
- Between two MST regions, the packet is forwarded along the CST.

Implementation of MSTP on devices

MSTP is compatible with STP and RSTP. Devices that are running MSTP and that are used for spanning tree calculation can identify STP and RSTP protocol packets.

In addition to basic MSTP functions, the following functions are provided for ease of management:

- Root bridge hold
- Root bridge backup
- Root guard
- BPDU guard
- Loop guard
- TC-BPDU guard
- BPDU drop

Protocols and standards

MSTP is documented in the following protocols and standards:

- IEEE 802.1d: *Media Access Control (MAC) Bridges*
- IEEE 802.1w: *Part 3: Media Access Control (MAC) Bridges—Amendment 2: Rapid Reconfiguration*
- IEEE 802.1s: *Virtual Bridged Local Area Networks—Amendment 3: Multiple Spanning Trees*

MSTP configuration task list

Before configuring MSTP, you must plan the role of each device in each MSTI, root bridge or leaf node, and then configure the devices as planned. In each MSTI, only one device acts as the root bridge, and all others act as leaf nodes.

Complete these tasks to configure MSTP:

Task		Remarks
Configuring the root bridge	Configuring an MST region	Required
	Configuring the root bridge or a secondary root bridge	Optional
	Configuring the work mode of an MSTP device	Optional
	Configuring the priority of a device	Optional
	Configuring the maximum hops of an MST region	Optional
	Configuring the network diameter of a switched network	Optional
	Configuring timers of MSTP	Optional
	Configuring the timeout factor	Optional
	Configuring the maximum port rate	Optional
	Configuring ports as edge ports	Optional
	Configuring the link type of ports	Optional
	Configuring the mode a port uses to recognize/send MSTP packets	Optional

Task		Remarks
	Enabling the output of port state transition information	Optional
	Enabling the MSTP feature	Required
Configuring the leaf nodes	Configuring an MST region	Required
	Configuring the work mode of an MSTP device	Optional
	Configuring the timeout factor	Optional
	Configuring the maximum port rate	Optional
	Configuring ports as edge ports	Optional
	Configuring path costs of ports	Optional
	Configuring port priority	Optional
	Configuring the link type of ports	Optional
	Configuring the mode a port uses to recognize/send MSTP packets	Optional
	Enabling the output of port state transition information	Optional
	Enabling the MSTP feature	Required
Performing mCheck		Optional
Configuring Digest Snooping		Optional
Configuring No Agreement Check		Optional
Configuring protection functions		Optional

NOTE:

- If GVRP and MSTP are enabled on a device at the same time, GVRP packets are forwarded along the CIST. To advertise a certain VLAN within the network through GVRP, be sure that this VLAN is mapped to the CIST (MSTI 0) when you configure the VLAN-to-instance mapping table. For more information about GVRP, see the chapter "GVRP configuration."
- MSTP is mutually exclusive with any of the following functions on a port: RRPP, Smart Link, and BPDU tunneling.
- Configurations made in system view take effect globally. Configurations made in Ethernet interface view take effect on the current interface only. Configurations made in port group view take effect on all member ports in the port group. Configurations made in Layer 2 aggregate interface view take effect only on the aggregate interface. Configurations made on an aggregation member port can take effect only after the port is removed from the aggregation group.
- After you enable MSTP on a Layer 2 aggregate interface, the system performs MSTP calculation on the Layer 2 aggregate interface but not on the aggregation member ports. The MSTP enable state and forwarding state of each selected port in an aggregation group is consistent with those of the corresponding Layer 2 aggregate interface.
- Though the member ports of an aggregation group do not participate in MSTP calculation, the ports still reserve their MSTP configurations for participating in MSTP calculation after leaving the aggregation group.

Configuring MSTP

Configuring an MST region

Make the following configurations on the root bridge and on the leaf nodes separately.

Follow these steps to configure an MST region:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter MST region view	**stp region-configuration**	—
Configure the MST region name	**region-name** *name*	Optional The MST region name is the MAC address by default.
Configure the VLAN-to-instance mapping table	**instance** *instance-id* **vlan** *vlan-list*	Optional Use either command. All VLANs in an MST region are mapped to the CIST (or MSTI 0) by default.
	vlan-mapping modulo *modulo*	
Configure the MSTP revision level of the MST region	**revision-level** *level*	Optional 0 by default.
Display the MST region configurations that are not activated yet	**check region-configuration**	Optional
Activate MST region configuration manually	**active region-configuration**	Required
Display the activated configuration information of the MST region	**display stp region-configuration** [| { **begin** | **exclude** | **include** } *regular-expression*]	Optional Available in any view

NOTE:

- Two or more MSTP-enabled devices belong to the same MST region only if they are configured to have the same format selector (0 by default, not configurable), MST region name, VLAN-to-instance mapping entries in the MST region, and MST region revision level, and they are connected via a physical link.
- The configuration of MST region–related parameters, especially the VLAN-to-instance mapping table, will cause MSTP to begin a new spanning tree calculation process, which might result in network topology instability. To reduce the possibility of topology instability caused by configuration, MSTP does not immediately begin a new spanning tree calculation process when it is processing MST region–related configurations. Instead, such configurations takes effect only after you activate the MST region–related parameters by using the **active region-configuration** command, or enable MSTP by using the **stp enable** command if MSTP is disabled.

Configuring the root bridge or a secondary root bridge

You can have MSTP determine the root bridge of a spanning tree through MSTP calculation, or you can specify the current device as the root bridge or as a secondary root bridge using the commands that the system provides.

Note the following rules:

- A device has independent roles in different MSTIs. It can act as the root bridge or a secondary root bridge of one MSTI and the root bridge or a secondary root bridge of another MSTI. However, one device cannot be the root bridge and a secondary root bridge in the same MSTI at the same time.
- There is only one root bridge in effect in a spanning tree instance. If two or more devices have been designated as root bridges of the same spanning tree instance, MSTP selects the device with the lowest MAC address as the root bridge.
- When the root bridge of an instance fails or is shut down, the secondary root bridge (if you have specified one) can take over the role of the primary root bridge. However, if you specify a new primary root bridge for the instance then, the one you specify, not the secondary root bridge will become the root bridge. If you have specified multiple secondary root bridges for an instance, when the root bridge fails, MSTP will select the secondary root bridge with the lowest MAC address as the new root bridge.

Configuring the current device as the root bridge of a specific spanning tree

Follow these steps to configure the current device as the root bridge of a specific spanning tree:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Configure the current device as the root bridge of a specific spanning tree	**stp** [**instance** *instance-id*] **root primary**	Required By default, a device does not function as the root bridge of any spanning tree.

Configuring the current device as a secondary root bridge of a specific spanning tree

Follow these steps to configure the current device as a secondary root bridge of a specific spanning tree:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Configure the current device as a secondary root bridge of a specific spanning tree	**stp** [**instance** *instance-id*] **root secondary**	Required By default, a device does not function as a secondary root bridge.

NOTE:

- After specifying the current device as the root bridge or a secondary root bridge, you cannot change the priority of the device.
- Alternatively, you can configure the current device as the root bridge by setting the priority of the device to 0. For the device priority configuration, see "Configuring the priority of a device."

Configuring the work mode of an MSTP device

MSTP and RSTP are mutually compatible and can recognize each other's protocol packets. However, STP cannot recognize MSTP packets. For hybrid networking with legacy STP devices, and for full interoperability with RSTP-enabled devices, MSTP supports the following work modes: STP-compatible mode, RSTP mode, and MSTP mode.

- In STP-compatible mode, all ports of the device send STP BPDUs,

- In RSTP mode, all ports of the device send RSTP BPDUs. If the device detects that it is connected to a legacy STP device, the port that connects to the legacy STP device will automatically migrate to STP-compatible mode.
- In MSTP mode, all ports of the device send MSTP BPDUs. If the device detects that it is connected to a legacy STP device, the port that connects to the legacy STP device will automatically migrate to STP-compatible mode.

Make this configuration on the root bridge and on the leaf nodes separately.

Follow these steps to configure the MSTP work mode:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Configure the work mode of MSTP	**stp mode** { **stp** | **rstp** | **mstp** }	Required MSTP mode by default.

Configuring the priority of a device

Priority is a factor in spanning tree calculation. The priority of a device determines whether it can be elected as the root bridge of a spanning tree. A lower numeric value indicates a higher priority. You can set the priority of a device to a low value to specify the device as the root bridge of the spanning tree. An MSTP-enabled device can have different priorities in different MSTIs.

Make this configuration on the root bridge only.

Follow these steps to configure the priority of a device in a specified MSTI:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Configure the priority of the current device in a specified MSTI	**stp** [**instance** *instance-id*] **priority** *priority*	Required 32768 by default.

△ CAUTION:

- You cannot change the priority of a device after it is configured as the root bridge or as a secondary root bridge.
- During root bridge selection, if all devices in a spanning tree have the same priority, the one with the lowest MAC address will be selected as the root bridge of the spanning tree.

Configuring the maximum hops of an MST region

By setting the maximum hops of an MST region, you can restrict the region size. The maximum hops configured on the regional root bridge will be used as the maximum hops of the MST region.

Configuration BPDUs sent by the regional root bridge always have a hop count set to the maximum value. When a switch receives this configuration BPDU, it decrements the hop count by 1, and uses the new hop count in the BPDUs that it propagates. When the hop count of a BPDU reaches 0, it is discarded by the device that received it. Devices beyond the reach of the maximum hop can no longer participate in spanning tree calculation, so the size of the MST region is limited.

Make this configuration on the root bridge only. All devices other than the root bridge in the MST region use the maximum hop value set for the root bridge.

Follow these steps to configure the maximum number of hops of an MST region:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Configure the maximum hops of the MST region	**stp max-hops** *hops*	Required 20 by default.

Configuring the network diameter of a switched network

Any two terminal devices in a switched network are connected through a specific path composed of a series of devices. The network diameter is the number of devices on the path composed of the most devices. The network diameter is a parameter that indicates the network size. A bigger network diameter indicates a larger network size.

Make this configuration on the root bridge only.

Follow these steps to configure the network diameter of a switched network:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Configure the network diameter of the switched network	**stp bridge-diameter** *diameter*	Required 7 by default.

NOTE:

- Based on the network diameter you configured, MSTP automatically sets an optimal hello time, forward delay, and max age for the device.
- In MSTP mode, each MST region is considered as a device. The network diameter configuration is effective only for the CIST (or the common root bridge), but not for MSTIs.

Configuring timers of MSTP

STP calculation involves the following timing parameters.

- Forward delay: Determines the time interval of port state transition. To prevent temporary loops, a port must go through an intermediate state, the learning state, before it transitions from the discarding state to the forwarding state, and must wait a certain period of time (forward delay) before it transitions from one state to another to keep synchronized with the remote device during state transition.
- Hello time: Used to detect link failures. STP sends configuration BPDUs at the interval of hello time. If a device fails to receive configuration BPDUs within the hello time, a new spanning tree calculation process will be triggered because of configuration BPDU timeout.
- Max age: Used to detect configuration BPDU timeout. In the CIST, the device uses the max age parameter to determine whether a configuration BPDU received on a port has expired. If a port receives a configuration BPDU that has expired, that MSTI must be re-calculated. The max age is meaningless for MSTIs.

To avoid frequent network changes, be sure that the settings of the hello time, forward delay and max age timers meet the following formulas:

- 2 × (forward delay – 1 second) ƒ max age

- Max age ƒ 2 × (hello time + 1 second)

HP does not recommend you to manually set the timers. Instead, you can use the **stp bridge-diameter** command to set the network diameter, and let the network automatically adjust the three timers according to the network size. When the network diameter is the default value, the three timers are also set to their defaults.

Make this configuration on the common root bridge only, and then this configuration applies to all devices on the entire switched network.

Follow these steps to configure the timers of MSTP:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Configure the forward delay timer	**stp timer forward-delay** *time*	Optional 1500 centiseconds (15 seconds) by default.
Configure the hello timer	**stp timer hello** *time*	Optional 200 centiseconds (2 seconds) by default.
Configure the max age timer	**stp timer max-age** *time*	Optional 2000 centiseconds (20 seconds) by default.

NOTE:

- The length of the forward delay is related to the network diameter of the switched network. The larger the network diameter is, the longer the forward delay should be. If the forward delay is too short, temporary redundant paths might occur. If the forward delay is too long, network convergence might take a long time. HP recommends that you use the default setting.
- An appropriate hello time enables the device to quickly detect link failures on the network without using excessive network resources. If the hello time is set too long, the device will mistake packet loss as a link failure and trigger a new spanning tree calculation process. If the hello time is set too short, the device will frequently send repeated configuration BPDUs, which adds to the device burden and wastes network resources. HP recommends that you use the default setting.
- If the max age time is too short, the network devices will frequently begin spanning tree calculations and might mistake network congestion as a link failure. If the max age is too long, the network might fail to quickly detect link failures and fail to quickly begin spanning tree calculations, reducing the auto-sensing capability of the network. HP recommends that you use the default setting.

Configuring the timeout factor

The timeout factor is a parameter used to decide the timeout time in the following formula: Timeout time = timeout factor × 3 × hello time.

After the network topology is stabilized, each non-root-bridge device forwards configuration BPDUs to the downstream devices at the interval of hello time to determine whether any link is faulty. If a device does not receive a BPDU from the upstream device within nine times the hello time, it assumes that the upstream device has failed and starts a new spanning tree calculation process.

Sometimes a device might fail to receive a BPDU from the upstream device because the upstream device is busy. If a spanning tree calculation occurs, the calculation can fail and also waste network resources.

In a stable network, you can avoid such unwanted spanning tree calculations by setting the timeout factor to 5, 6, or 7.

Follow these steps to configure the timeout factor:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Configure the timeout factor of the device	**stp timer-factor** *factor*	Required 3 by default.

Configuring the maximum port rate

The maximum rate of a port refers to the maximum number of BPDUs the port can send within each hello time. The maximum rate of a port is related to the physical status of the port and the network structure.

Make this configuration on the root bridge and on the leaf nodes separately.

Follow these steps to configure the maximum rate of a port or a group of ports:

<table>
<tr><th colspan="2">To do...</th><th>Use the command...</th><th>Remarks</th></tr>
<tr><td colspan="2">Enter system view</td><td><b>system-view</b></td><td>—</td></tr>
<tr><td rowspan="2">Enter interface view or port group view</td><td>Enter Ethernet interface view or Layer 2 aggregate interface view</td><td><b>interface</b> <i>interface-type interface-number</i></td><td rowspan="2">Required
Use either command.</td></tr>
<tr><td>Enter port group view</td><td><b>port-group manual</b> <i>port-group-name</i></td></tr>
<tr><td colspan="2">Configure the maximum rate of the ports</td><td><b>stp transmit-limit</b> <i>limit</i></td><td>Required
10 by default.</td></tr>
</table>

NOTE:

The higher the maximum port rate is, the more BPDUs will be sent within each hello time, and the more system resources will be used. By setting an appropriate maximum port rate, you can limit the rate at which the port sends BPDUs and prevent MSTP from using excessive network resources when the network becomes unstable. HP recommends that you use the default setting.

Configuring ports as edge ports

If a port directly connects to a user terminal rather than another device or a shared LAN segment, this port is regarded as an edge port. When a network topology change occurs, an edge port will not cause a temporary loop. Because a device does not determine whether a port is directly connected to a terminal, you must manually configure the port as an edge port. After that, the port can transition rapidly from the blocked state to the forwarding state.

Make this configuration on the root bridge and on the leaf nodes separately.

Follow these steps to specify a port or a group of ports as edge port or ports:

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

To do...		Use the command...	Remarks
Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the current ports as edge ports		**stp edged-port enable**	Required All ports are non-edge ports by default.

NOTE:

- If BPDU guard is disabled, a port set as an edge port will become a non-edge port again if it receives a BPDU from another port. To restore the edge port, re-enable it.
- If a port directly connects to a user terminal, configure it as an edge port and enable BPDU guard for it. This enables the port to transition to the forwarding state quickly while ensuring network security.
- Among loop guard, root guard and edge port settings, only one function (whichever is configured the earliest) can at a time take effect on a port.

Configuring path costs of ports

Path cost is a parameter related to the rate of a port. On an MSTP-enabled device, a port can have different path costs in different MSTIs. Setting appropriate path costs allows VLAN traffic flows to be forwarded along different physical links, achieving VLAN-based load balancing.

You can have the device automatically calculate the default path cost, or you can configure the path cost for ports.

Make the following configurations on the leaf nodes only.

Specifying a standard that the device uses when it calculates the default path cost

You can specify a standard for the device to use in automatic calculation for the default path cost. The device supports the following standards:

- **dot1d-1998**—The device calculates the default path cost for ports based on IEEE 802.1d-1998.
- **dot1t**—The device calculates the default path cost for ports based on IEEE 802.1t.
- **legacy**—The device calculates the default path cost for ports based on a private standard.

Follow these steps to specify a standard for the device to use when it calculates the default path cost:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Specify a standard for the device to use when it calculates the default path costs of its ports	**stp pathcost-standard** { **dot1d-1998** | **dot1t** | **legacy** }	Optional By default, the device calculates the default path cost for ports based on a private standard.

CAUTION:

If you change the standard that the device uses to calculate the default path costs, you restore the path costs to the default.

Table 13 shows the mappings between the link speed and the path cost.

Table 13 Mappings between the link speed and the path cost

Link speed	Port type	Path cost		
		IEEE 802.1d-1998	IEEE 802.1t	Private standard
0	—	65535	200,000,000	200,000
10 Mbps	Single Port	100	2,000,000	2,000
	Aggregate interface containing 2 selected ports		1,000,000	1,800
	Aggregate interface containing 3 selected ports		666,666	1,600
	Aggregate interface containing 4 selected ports		500,000	1,400
100 Mbps	Single Port	19	200,000	200
	Aggregate interface containing 2 selected ports		100,000	180
	Aggregate interface containing 3 selected ports		66,666	160
	Aggregate interface containing 4 selected ports		50,000	140
1000 Mbps	Single Port	4	20,000	20
	Aggregate interface containing 2 selected ports		10,000	18
	Aggregate interface containing 3 selected ports		6666	16
	Aggregate interface containing 4 selected ports		5000	14
10 Gbps	Single Port	2	2000	2
	Aggregate interface containing 2 selected ports		1000	1
	Aggregate interface containing 3 selected ports		666	1

Link speed	Port type	Path cost		
		IEEE 802.1d-1998	IEEE 802.1t	Private standard
	Aggregate interface containing 4 selected ports		500	1

NOTE:

When calculating path cost for an aggregate interface, IEEE 802.1d-1998 does not take into account the number of selected ports in its aggregation group as IEEE 802.1t does. The calculation formula of IEEE 802.1t is: Path cost = 200,000,000/link speed (in 100 kbps), where link speed is the sum of the link speed values of the selected ports in the aggregation group.

Configuring path costs of ports

Follow these steps to configure the path cost of ports:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the path cost of the ports		**stp** [**instance** *instance-id*] **cost** *cost*	Required By default, MSTP automatically calculates the path cost of each port.

CAUTION:

When the path cost of a port changes, MSTP re-calculates the role of the port and initiates a state transition.

Configuration example

# Specify that the device uses IEEE 802.1d-1998 to calculate the default path costs of its ports.

```
<Sysname> system-view
[Sysname] stp pathcost-standard dot1d-1998
```

# Set the path cost of GigabitEthernet 1/0/3 on MSTI 2 to 200.

```
<Sysname> system-view
[Sysname] interface gigabitethernet 1/0/3
[Sysname-GigabitEthernet1/0/3] stp instance 2 cost 200
```

Configuring port priority

The priority of a port is an important factor in determining whether the port can be elected as the root port of a device. If all other conditions are the same, the port with the highest priority will be elected as the root port.

On an MSTP-enabled device, a port can have different priorities in different MSTIs, and the same port can play different roles in different MSTIs, so that data of different VLANs can be propagated along different physical paths, implementing per-VLAN load balancing. You can set port priority values based on the actual networking requirements.

Make this configuration on the leaf nodes only.

Follow these steps to configure the priority of a port or a group of ports:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the port priority		**stp** [**instance** *instance-id*] **port priority** *priority*	Required 128 for all ports by default.

NOTE:

- When the priority of a port changes, MSTP re-calculates the role of the port and initiates a state transition.
- A lower priority value indicates a higher priority. If you configure the same priority value for all the ports on a device, the specific priority of a port depends on the index number of the port. A lower index number means a higher priority. Changing the priority of a port triggers a new spanning tree calculation process.

Configuring the link type of ports

A point-to-point link is a link that directly connects two devices. If the two ports across a point-to-point link are root ports or designated ports, the ports can rapidly transition to the forwarding state after a proposal-agreement handshake process.

Make this configuration on the root bridge and on the leaf nodes separately.

Follow these steps to configure the link type of a port or a group of ports:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the link type of ports		**stp point-to-point** { **auto** | **force-false** | **force-true** }	Required By default, the port automatically detects whether its link is point-to-point.

NOTE:

- If the current port is a Layer 2 aggregate interface or if it works in full duplex mode, you can configure the link to which the current port connects as a point-to-point link. HP recommends that you use the default setting, and let MSTP detect the link status automatically.
- If you configure a port as connecting to a point-to-point link or a non-point-to-point link, the setting takes effect for the port in all MSTIs.
- If the physical link to which the port connects is not a point-to-point link and you manually set it to be one, your configuration might cause temporary loops.

Configuring the mode a port uses to recognize/send MSTP packets

A port can receive/send MSTP packets in the following formats:

- **dot1s**—802.1s-compliant standard format, and
- **legacy**—Compatible format

By default, the packet format recognition mode of a port is **auto**. The port automatically distinguishes the two MSTP packet formats, and determines the format of packets that it will send based on the recognized format.

You can configure the MSTP packet format on a port. When working in MSTP mode after the configuration, the port sends and receives only MSTP packets of the format that you have configured to communicate with devices that send packets of the same format.

Make this configuration on the root bridge and on the leaf nodes separately.

Follow these steps to configure the MSTP packet format to be supported on a port or a group of ports:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the mode that the port uses to recognize/send MSTP packets		**stp compliance** { **auto** | **dot1s** | **legacy** }	Required **auto** by default.

NOTE:

- MSTP provides the MSTP packet format incompatibility guard function. In MSTP mode, if a port is configured to recognize/send MSTP packets in a mode other than **auto**, and if it receives a packet in a format different from the specified type, the port becomes a designated port and remains in the discarding state to prevent the occurrence of a loop.
- MSTP provides the MSTP packet format frequent change guard function. If a port receives MSTP packets of different formats frequently, the MSTP packet format configuration can contain errors. If the port is working in MSTP mode, it will be disabled for protection. Only network administrators can restore those closed ports.

Enabling the output of port state transition information

A large-scale, MSTP-enabled network can have many MSTIs, and ports might frequently transition from one state to another. In this situation, you can enable devices to output the port state transition information of all MSTIs or the specified MSTI in order to monitor the port states in real time.

Make this configuration separately on the root bridge and on the leaf nodes.

Follow these steps to enable output of port state transition information:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Enable output of port state transition information	**stp port-log** { **all** | **instance** *instance-id* }	Required Enabled by default.

Enabling the MSTP feature

You must enable MSTP for the device before any other MSTP-related configurations can take effect.

Make this configuration on the root bridge and on the leaf nodes separately.

Follow these steps to enable the MSTP feature:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enable the MSTP feature globally		**stp enable**	Required MSTP is globally disabled by default.
Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Enable the MSTP feature for the ports		**stp enable**	Optional By default, MSTP is enabled for all ports after it is enabled for the device globally.

NOTE:

- In system view, you can use the **stp enable** or **undo stp enable** command to enable or disable STP globally.
- You can use the **undo stp enable** command to disable the MSTP feature for certain ports so that they will not participate in spanning tree calculation to save the CPU resources of the device.

Performing mCheck

MSTP has three working modes: STP compatible mode, RSTP mode, and MSTP mode.

If a port on a device that is running MSTP (or RSTP) connects to a device that is running STP, this port automatically migrates to the STP-compatible mode. However, it will not be able to automatically migrate back to the MSTP (or RSTP) mode, but will remain working in the STP-compatible mode under the following circumstances:

- The device that is running STP is shut down or removed.
- The device that is running STP migrates to the MSTP (or RSTP) mode.

You can perform an mCheck operation to force the port to migrate to the MSTP (or RSTP) mode.

The following two methods for performing mCheck produce the same results.

Performing mCheck globally

Follow these steps to perform global mCheck:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Perform mCheck	**stp mcheck**	Required

Performing mCheck in interface view

Follow these steps to perform mCheck in interface view:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	—
Perform mCheck	**stp mcheck**	Required

NOTE:

An mCheck operation takes effect on a device only when MSTP operates in RSTP or MSTP mode.

Configuring Digest Snooping

As defined in IEEE 802.1s, connected devices are in the same region only when their MST region-related configurations (region name, revision level, VLAN-to-instance mappings) are identical. An MSTP-enabled device identifies devices in the same MST region by determining the configuration ID in BPDU packets. The configuration ID includes the region name, revision level, configuration digest, which is in 16-byte length and is the result calculated via the HMAC-MD5 algorithm based on VLAN-to-instance mappings.

Because MSTP implementations vary with vendors, the configuration digests calculated via private keys are different. The different vendors' devices in the same MST region can not communicate with each other.

Enabling the Digest Snooping feature on the port that connects the local device to a third-party device in the same MST region can make the two devices communicate with each other.

NOTE:

Before you enable Digest Snooping, ensure that associated devices of different vendors are connected and run MSTP.

Configuring the Digest Snooping feature

You can enable Digest Snooping only on a device that is connected to a third-party device that uses its private key to calculate the configuration digest.

Follow these steps to configure Digest Snooping:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Enable Digest Snooping on the interface or port group		**stp config-digest-snooping**	Required Disabled by default.
Return to system view		**quit**	—
Enable global Digest Snooping		**stp config-digest-snooping**	Required Disabled by default.

CAUTION:

- With digest snooping enabled, in-the-same-region verification does not require comparison of configuration digest, so the VLAN-to-instance mappings must be the same on associated ports.
- With global Digest Snooping enabled, modification of VLAN-to-instance mappings and removal of the current region configuration via the **undo stp region-configuration** command are not allowed. You can modify only the region name and revision level.
- To make Digest Snooping take effect, you must enable Digest Snooping both globally and on associated ports. HP recommends that you enable Digest Snooping on all associated ports first and then enable it globally. This will make the configuration take effect on all configured ports and reduce impact on the network.
- To avoid loops, do not enable Digest Snooping on MST region edge ports.
- HP recommends that you enable Digest Snooping first and then MSTP. To avoid causing traffic interruption, do not configure Digest Snooping when the network is already working well.

Digest Snooping configuration example

1. Network requirements

As shown in Figure 21:

- Device A and Device B connect to Device C, which is a third-party device. All these devices are in the same region.
- Enable Digest Snooping on the ports of Device A and Device B that connect Device C, so that the three devices can communicate with one another.

Figure 21 Digest Snooping configuration

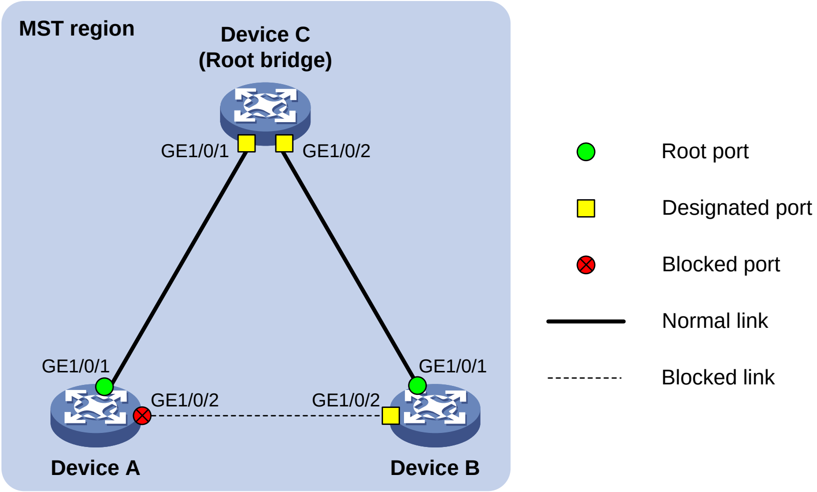

2. Configuration procedure

# Enable Digest Snooping on GigabitEthernet 1/0/1 of Device A and enable global Digest Snooping on Device A.

```
<DeviceA> system-view
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] stp config-digest-snooping
[DeviceA-GigabitEthernet1/0/1] quit
[DeviceA] stp config-digest-snooping
```

# Enable Digest Snooping on GigabitEthernet 1/0/1 of Device B and enable global Digest Snooping on Device B.

```
<DeviceB> system-view
[DeviceB] interface gigabitethernet 1/0/1
[DeviceB-GigabitEthernet1/0/1] stp config-digest-snooping
[DeviceB-GigabitEthernet1/0/1] quit
[DeviceB] stp config-digest-snooping
```

Configuring No Agreement Check

In RSTP and MSTP, the following types of messages are used for rapid state transition on designated ports:

- Proposal—Sent by designated ports to request rapid transition
- Agreement—Used to acknowledge rapid transition requests

Both RSTP and MSTP devices can perform rapid transition on a designated port only when the port receives an agreement packet from the downstream device. RSTP and MSTP devices have the following differences:

- For MSTP, the root port of the downstream device sends an agreement packet only after it receives an agreement packet from the upstream device.
- For RSTP, the downstream device sends an agreement packet regardless of whether an agreement packet from the upstream device is received.

Figure 22 shows the rapid state transition mechanism on MSTP designated ports.

Figure 22 Rapid state transition of an MSTP designated port

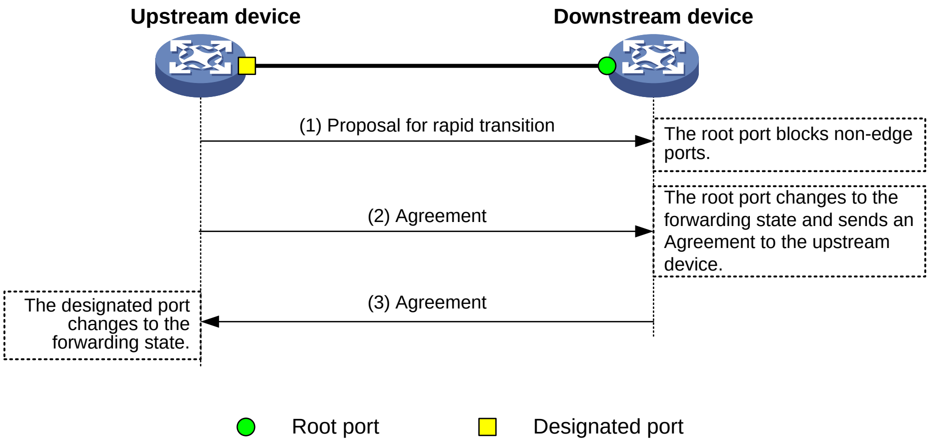
Upstream device
Downstream device
(1) Proposal for rapid transition
The root port blocks non-edge ports.
(2) Agreement
The root port changes to the forwarding state and sends an Agreement to the upstream device.
(3) Agreement
The designated port changes to the forwarding state.
Root port
Designated port

Figure 23 shows rapid state transition of an RSTP designated port.

Figure 23 Rapid state transition of an RSTP designated port

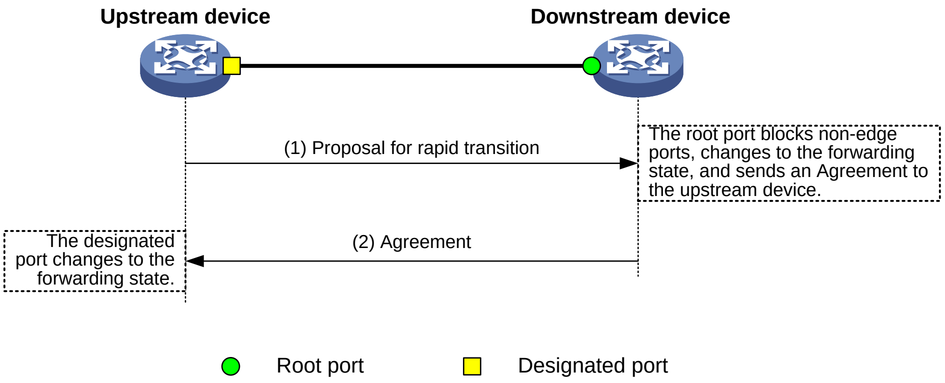
Upstream device
Downstream device
(1) Proposal for rapid transition
The root port blocks non-edge ports, changes to the forwarding state, and sends an Agreement to the upstream device.
(2) Agreement
The designated port changes to the forwarding state.
Root port
Designated port

If the upstream device is a third-party device, the rapid state transition implementation might be limited. For example, when the upstream device uses a rapid transition mechanism similar to that of RSTP, and the downstream device adopts MSTP and does not work in RSTP mode, the root port on the downstream device receives no agreement packet from the upstream device and sends no agreement packets to the upstream device. As a result, the designated port of the upstream device fails to transit rapidly, and can only change to the forwarding state after a period twice the Forward Delay.

You can enable the No Agreement Check feature on the downstream device's port to enable the designated port of the upstream device to transit its state rapidly.

Configuration prerequisites

Before you configure the No Agreement Check function, complete the following tasks:

- Connect a device is to a third-party upstream device that supports MSTP via a point-to-point link.
- Configure the same region name, revision level and VLAN-to-instance mappings on the two devices, assigning them to the same region.

Configuring the No Agreement Check function

To make the No Agreement Check feature take effect, enable it on the root port.

Follow these steps to configure No Agreement Check:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter interface or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Enable No Agreement Check		**stp no-agreement-check**	Required Disabled by default.

No Agreement Check configuration example

1. Network requirements

As shown in Figure 24:

- Device A connects to Device B, a third-party device that has different MSTP implementation. Both devices are in the same region.
- Device B is the regional root bridge, and Device A is the downstream device.

Figure 24 No Agreement Check configuration

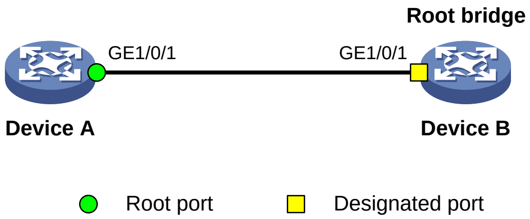

2. Configuration procedure

# Enable No Agreement Check on GigabitEthernet 1/0/1 of Device A.

```
<DeviceA> system-view
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] stp no-agreement-check
```

Configuring protection functions

An MSTP-enabled device supports the following protection functions:

- BPDU guard
- Root guard
- Loop guard
- TC-BPDU guard
- BPDU drop

Configuration prerequisites

MSTP has been correctly configured on the device.

Enabling BPDU guard

For access layer devices, the access ports can directly connect to the user terminals (such as PCs) or file servers. The access ports are configured as edge ports to allow rapid transition. When these ports receive configuration BPDUs, the system automatically sets these ports as non-edge ports and starts a new spanning tree calculation process. This causes a change of network topology. Under normal conditions, these ports should not receive configuration BPDUs. However, if someone forges configuration BPDUs maliciously to attack the devices, the network will become unstable.

MSTP provides the BPDU guard function to protect the system against such attacks. With the BPDU guard function enabled on the devices, when edge ports receive configuration BPDUs, MSTP closes these ports and notifies the NMS that these ports have been closed by MSTP. The device will reactivate the closed ports after a detection interval. For more information about this detection interval, see the *Fundamentals Configuration Guide*.

Make this configuration on a device with edge ports configured.

Follow these steps to enable BPDU guard:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Enable the BPDU guard function for the device	**stp bpdu-protection**	Required Disabled by default.

NOTE:

BPDU guard does not take effect on loopback testing-enabled ports. For more information about loopback testing, see the chapter "Ethernet interface configuration."

Enabling root guard

The root bridge and secondary root bridge of a spanning tree should be located in the same MST region. Especially for the CIST, the root bridge and secondary root bridge are put in a high-bandwidth core region during network design. However, because of possible configuration errors or malicious attacks in the network, the legal root bridge might receive a configuration BPDU with a higher priority. Another device will supersede the current legal root bridge, causing an undesired change of the network topology. The traffic that should go over high-speed links is switched to low-speed links, resulting in network congestion.

To prevent this situation, MSTP provides the root guard function. If the root guard function is enabled on a port of a root bridge, this port will keep playing the role of designated port on all MSTIs. After this port receives a configuration BPDU with a higher priority from an MSTI, it immediately sets that port to the listening state in the MSTI, without forwarding the packet. This is equivalent to disconnecting the link connected to this port in the MSTI. If the port receives no BPDUs with a higher priority within twice the forwarding delay, it reverts to its original state.

Make this configuration on a designated port.

Follow these steps to enable root guard:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.

To do...		Use the command...	Remarks
	Enter port group view	**port-group manual** *port-group-name*	
Enable the root guard function for the port(s)		**stp root-protection**	Required Disabled by default.

NOTE:

Among loop guard, root guard and edge port settings, only one function (whichever is configured the earliest) can at a time take effect on a port.

Enabling loop guard

By continuing to receive BPDUs from the upstream device, a device can maintain the state of the root port and blocked ports. However, because of link congestion or unidirectional link failures, these ports might fail to receive BPDUs from the upstream devices. The device will reselect the port roles: Those ports in forwarding state that failed to receive upstream BPDUs will become designated ports, and the blocked ports will transition to the forwarding state, resulting in loops in the switched network. The loop guard function can suppress the occurrence of such loops.

The initial state of a loop guard-enabled port is discarding in every MSTI. When the port receives BPDUs, its state transitions normally. Otherwise, it stays in the discarding state to prevent temporary loops.

Make this configuration on the root port and alternate ports of a device.

Follow these steps to enable loop guard:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Enable the loop guard function for the ports		**stp loop-protection**	Required Disabled by default.

NOTE:

- Do not enable loop guard on a port that connects user terminals. Otherwise, the port will stay in the discarding state in all MSTIs because it cannot receive BPDUs.
- Among loop guard, root guard and edge port settings, only one function (whichever is configured the earliest) can at a time take effect on a port.

Enabling TC-BPDU guard

When a switch receives topology change (TC) BPDUs (the BPDUs that notify devices of topology changes), the switch flushes its forwarding address entries. If someone forges TC-BPDUs to attack the switch, the switch will receive a large number of TC-BPDUs within a short time and be busy with forwarding address entry flushing. This affects network stability.

With the TC-BPDU guard function, you can set the maximum number of immediate forwarding address entry flushes that the switch can perform within a specified period of time after it receives the first TC-BPDU. For TC-BPDUs received in excess of the limit, the switch performs forwarding address entry flush only when the time period expires. This prevents frequent flushing of forwarding address entries.

Follow these steps to enable TC-BPDU guard:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enable the TC-BPDU guard function	**stp tc-protection enable**	Optional Enabled by default.
Configure the maximum number of forwarding address entry flushes that the device can perform within a specific time period after it receives the first TC-BPDU	**stp tc-protection threshold** *number*	Optional 6 by default.

NOTE:

HP does not recommend you to disable this feature.

Enabling BPDU drop

In an STP-enabled network, after receiving BPDUs, a device performs STP calculation according to the received BPDUs and forwards received BPDUs to other devices in the network. This allows malicious attackers to attack the network by forging BPDUs. By continuously sending forged BPDUs, they can make all the devices in the network perform STP calculations all the time. As a result, problems such as CPU overload and BPDU protocol status errors occur.

To avoid this problem, you can enable BPDU drop on ports. A BPDU drop-enabled port does not receive any BPDUs and is invulnerable to forged BPDU attacks.

Follow these steps to enable BPDU drop on an Ethernet interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Enable BPDU drop on the current interface	**bpdu-drop any**	Required Disabled by default.

Displaying and maintaining MSTP

To do…	Use the command…	Remarks
Display information about abnormally blocked ports	**display stp abnormal-port** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display BPDU statistics on ports	**display stp bpdu-statistics** [**interface** *interface-type interface-number* [**instance** *instance-id*]] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

To do...	Use the command...	Remarks
Display information about ports blocked by STP protection functions	**display stp down-port** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the historical information of port role calculation for the specified MSTI or all MSTIs	**display stp** [**instance** *instance-id*] **history** [**slot** *slot-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the statistics of TC/TCN BPDUs sent and received by all ports in the specified MSTI or all MSTIs	**display stp** [**instance** *instance-id*] **tc** [**slot** *slot-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the status and statistics of MSTP	**display stp** [**instance** *instance-id*] [**interface** *interface-list* | **slot** *slot-number*] [**brief**] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the MST region configuration information that has taken effect	**display stp region-configuration** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the root bridge information of all MSTIs	**display stp root** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Clear the statistics of MSTP	**reset stp** [**interface** *interface-list*]	Available in user view

MSTP configuration example

Network requirements

As shown in Figure 25:

- All devices on the network are in the same MST region. Device A and Device B work on the distribution layer. Device C and Device D work on the access layer.
- Configure MSTP so that packets of different VLANs are forwarded along different spanning trees: Packets of VLAN 10 are forwarded along MSTI 1, those of VLAN 30 are forwarded along MSTI 3, those of VLAN 40 are forwarded along MSTI 4, and those of VLAN 20 are forwarded along MSTI 0.
- VLAN 10 and VLAN 30 are terminated on the distribution layer devices, and VLAN 40 is terminated on the access layer devices. The root bridges of MSTI 1 and MSTI 3 are Device A and Device B respectively, and the root bridge of MSTI 4 is Device C.

Figure 25 Network diagram for MSTP configuration

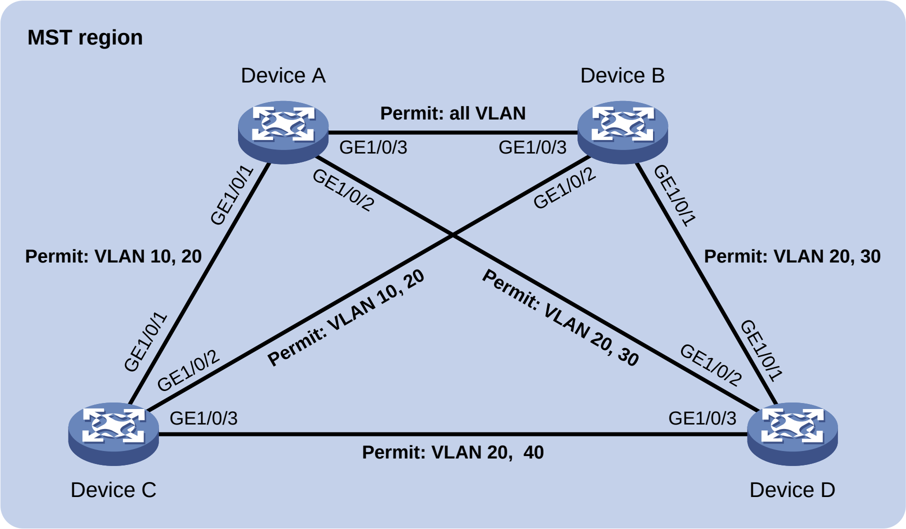

Configuration procedure

1. Configure VLANs and VLAN member ports (details not shown)

Create VLAN 10, VLAN 20, and VLAN 30 on Device A and Device B respectively, create VLAN 10, VLAN 20, and VLAN 40 on Device C, and create VLAN 20, VLAN 30, and VLAN 40 on Device D. Configure the ports on these devices as trunk ports and assign them to related VLANs.

2. Configure Device A

# Enter MST region view; configure the MST region name as **example**; map VLAN 10, VLAN 30, and VLAN 40 to MSTI 1, MSTI 3, and MSTI 4 respectively; configure the revision level of the MST region as 0.

```
<DeviceA> system-view
[DeviceA] stp region-configuration
[DeviceA-mst-region] region-name example
[DeviceA-mst-region] instance 1 vlan 10
[DeviceA-mst-region] instance 3 vlan 30
[DeviceA-mst-region] instance 4 vlan 40
[DeviceA-mst-region] revision-level 0
```

# Activate MST region configuration.

```
[DeviceA-mst-region] active region-configuration
[DeviceA-mst-region] quit
```

# Specify the current device as the root bridge of MSTI 1.

```
[DeviceA] stp instance 1 root primary
```

# Enable MSTP globally.

```
[DeviceA] stp enable
```

3. Configure Device B

# Enter MST region view, configure the MST region name as **example**, map VLAN 10, VLAN 30, and VLAN 40 to MSTI 1, MSTI 3, and MSTI 4 respectively, and configure the revision level of the MST region as 0.

```
<DeviceB> system-view
[DeviceB] stp region-configuration
```

```
[DeviceB-mst-region] region-name example
[DeviceB-mst-region] instance 1 vlan 10
[DeviceB-mst-region] instance 3 vlan 30
[DeviceB-mst-region] instance 4 vlan 40
[DeviceB-mst-region] revision-level 0
```

# Activate MST region configuration.

```
[DeviceB-mst-region] active region-configuration
[DeviceB-mst-region] quit
```

# Specify the current device as the root bridge of MSTI 3.

```
[DeviceB] stp instance 3 root primary
```

# Enable MSTP globally.

```
[DeviceB] stp enable
```

4. Configure Device C

# Enter MST region view, configure the MST region name as **example**, map VLAN 10, VLAN 30, and VLAN 40 to MSTI 1, MSTI 3, and MSTI 4 respectively, and configure the revision level of the MST region as 0.

```
<DeviceC> system-view
[DeviceC] stp region-configuration
[DeviceC-mst-region] region-name example
[DeviceC-mst-region] instance 1 vlan 10
[DeviceC-mst-region] instance 3 vlan 30
[DeviceC-mst-region] instance 4 vlan 40
[DeviceC-mst-region] revision-level 0
```

# Activate MST region configuration.

```
[DeviceC-mst-region] active region-configuration
[DeviceC-mst-region] quit
```

# Specify the current device as the root bridge of MSTI 4.

```
[DeviceC] stp instance 4 root primary
```

# Enable MSTP globally.

```
[DeviceC] stp enable
```

5. Configure on Device D

# Enter MST region view, configure the MST region name as **example**, map VLAN 10, VLAN 30, and VLAN 40 to MSTI 1, MSTI 3, and MSTI 4 respectively, and configure the revision level of the MST region as 0.

```
<DeviceD> system-view
[DeviceD] stp region-configuration
[DeviceD-mst-region] region-name example
[DeviceD-mst-region] instance 1 vlan 10
[DeviceD-mst-region] instance 3 vlan 30
[DeviceD-mst-region] instance 4 vlan 40
[DeviceD-mst-region] revision-level 0
```

# Activate MST region configuration.

```
[DeviceD-mst-region] active region-configuration
[DeviceD-mst-region] quit
```

# Enable MSTP globally.

```
[DeviceD] stp enable
```

6. Verify the configurations

You can use the **display stp brief** command to display brief spanning tree information on each device after the network is stable.

# Display brief spanning tree information on Device A.

```
[DeviceA] display stp brief
 MSTID      Port                         Role  STP State     Protection
   0        GigabitEthernet1/0/1         ALTE  DISCARDING    NONE
   0        GigabitEthernet1/0/2         DESI  FORWARDING    NONE
   0        GigabitEthernet1/0/3         ROOT  FORWARDING    NONE
   1        GigabitEthernet1/0/1         DESI  FORWARDING    NONE
   1        GigabitEthernet1/0/3         DESI  FORWARDING    NONE
   3        GigabitEthernet1/0/2         DESI  FORWARDING    NONE
   3        GigabitEthernet1/0/3         ROOT  FORWARDING    NONE
```

# Display brief spanning tree information on Device B.

```
[DeviceB] display stp brief
 MSTID      Port                         Role  STP State     Protection
   0        GigabitEthernet1/0/1         DESI  FORWARDING    NONE
   0        GigabitEthernet1/0/2         DESI  FORWARDING    NONE
   0        GigabitEthernet1/0/3         DESI  FORWARDING    NONE
   1        GigabitEthernet1/0/2         DESI  FORWARDING    NONE
   1        GigabitEthernet1/0/3         ROOT  FORWARDING    NONE
   3        GigabitEthernet1/0/1         DESI  FORWARDING    NONE
   3        GigabitEthernet1/0/3         DESI  FORWARDING    NONE
```

# Display brief spanning tree information on Device C.

```
[DeviceC] display stp brief
 MSTID      Port                         Role  STP State     Protection
   0        GigabitEthernet1/0/1         DESI  FORWARDING    NONE
   0        GigabitEthernet1/0/2         ROOT  FORWARDING    NONE
   0        GigabitEthernet1/0/3         DESI  FORWARDING    NONE
   1        GigabitEthernet1/0/1         ROOT  FORWARDING    NONE
   1        GigabitEthernet1/0/2         ALTE  DISCARDING    NONE
   4        GigabitEthernet1/0/3         DESI  FORWARDING    NONE
```

# Display brief spanning tree information on Device D.

```
[DeviceD] display stp brief
 MSTID      Port                         Role  STP State     Protection
   0        GigabitEthernet1/0/1         ROOT  FORWARDING    NONE
   0        GigabitEthernet1/0/2         ALTE  DISCARDING    NONE
   0        GigabitEthernet1/0/3         ALTE  DISCARDING    NONE
   3        GigabitEthernet1/0/1         ROOT  FORWARDING    NONE
   3        GigabitEthernet1/0/2         ALTE  DISCARDING    NONE
   4        GigabitEthernet1/0/3         ROOT  FORWARDING    NONE
```

Based on the output, you can draw the MSTI mapped to each VLAN, as shown in Figure 26.

Figure 26 MSTIs mapped to different VLANs

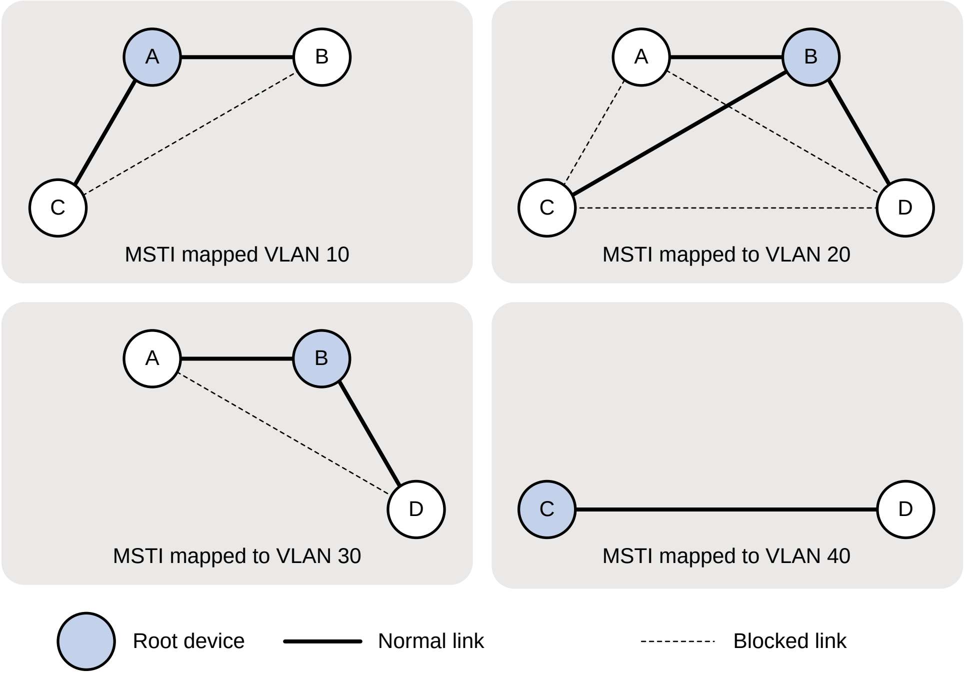

BPDU tunneling configuration

Introduction to BPDU tunneling

As a Layer 2 tunneling technology, BPDU tunneling enables Layer 2 protocol packets from geographically dispersed customer networks to be transparently transmitted over specific tunnels across a service provider network.

Background

Customers usually use dedicated lines in a service provider network to build their own Layer 2 networks. As a result, often a customer network consists of parts located at different sides of the service provider network. As shown in Figure 27, the devices for User A are CE 1 and CE 2, both of which belong to VLAN 100. User A's network is divided into network 1 and network 2, which are connected by the service provider network. When Layer 2 protocol packets cannot be transparently transmitted in the service provider network, User A's network cannot implement independent Layer 2 protocol calculation (for example, STP spanning tree calculation). The Layer 2 protocol calculation in User A's network is mixed with that in the service provider network.

Figure 27 BPDU tunneling application scenario

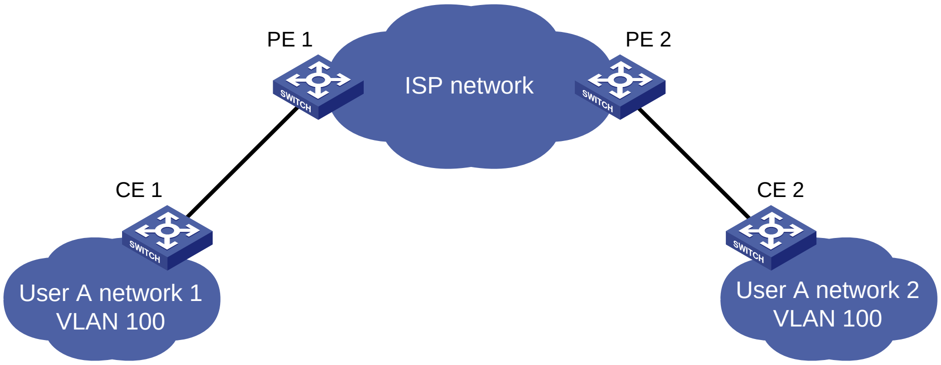

BPDU tunneling addresses this problem. With BPDU tunneling, Layer 2 protocol packets from customer networks can be transparently transmitted in the service provider network, as follows:

1. After receiving a Layer 2 protocol packet from User A network 1, PE 1 in the service provider network encapsulates the packet, replaces its destination MAC address with a specific multicast MAC address, and then forwards the packet in the service provider network.
2. The encapsulated Layer 2 protocol packet (called bridge protocol data unit, BPDU for short) is forwarded to PE 2 at the other end of the service provider network, which de-encapsulates the packet, restores the original destination MAC address of the packet, and then sends the packet to User A network 2.

NOTE:

Depending on the switch models, HP devices support BPDU tunneling for the following protocols:

- Cisco Discovery Protocol (CDP)
- Device Link Detection Protocol (DLDP)
- Ethernet Operation, Administration and Maintenance (EOAM)
- GARP VLAN Registration Protocol (GVRP)
- HW Group Management Protocol (HGMP)
- Link Aggregation Control Protocol (LACP)
- Link Layer Discovery Protocol (LLDP)
- Port Aggregation Protocol (PAGP)
- Per VLAN Spanning Tree (PVST)
- Spanning tree protocol (STP)
- Uni-directional Link Direction (UDLD)
- VLAN Trunking Protocol (VTP)

BPDU tunneling implementation

The BPDU tunneling implementations for different protocols are all similar. This section uses the Spanning Tree Protocol (STP) as an example to describe how to implement BPDU tunneling.

NOTE:

- This document uses the term *STP* in a broad sense. It includes STP, RSTP, and MSTP.
- STP calculates the topology of a network by transmitting BPDUs among devices in the network. For more information, see the chapter "MSTP configuration."

To avoid loops in your network, you can enable STP on your devices. When the topology changes at one side of the customer network, devices at that side of the customer network send BPDUs to devices on the other side of the customer network to ensure consistent spanning tree calculation in the entire customer network. However, because BPDUs are Layer 2 multicast frames, all STP-enabled devices, both in the customer network and in the service provider network, can receive and process these BPDUs. As a result, neither the service provider network nor the customer network can correctly calculate its independent spanning tree.

BPDU tunneling allows each network to calculate an independent spanning tree with STP.

BPDU tunneling delivers the following benefits:

- BPDUs can be transparently transmitted. BPDUs of the same customer network can be broadcast in a specific VLAN across the service provider network, so that the geographically dispersed networks of the same customer can implement consistent spanning tree calculation across the service provider network.
- BPDUs of different customer networks can be confined within different VLANs for transmission on the service provider network. This enables each customer network to perform independent spanning tree calculation.

Figure 28 Network diagram for BPDU tunneling implementation

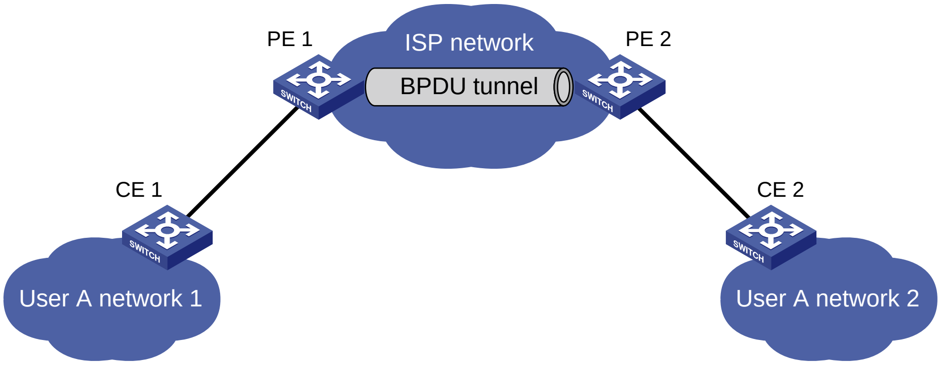

As shown in Figure 28, the upper part is the service provider network (ISP network), and the lower part represents two geographically dispersed segments of a customer network: User A network 1 and User A network 2. Enabling the BPDU tunneling function on the edge devices (PE 1 and PE 2) in the service provider network allows BPDUs of User A network 1 and User A network 2 to be transparently transmitted in the service provider network. This ensures consistent spanning tree calculation throughout User A network, without affecting the spanning tree calculation of the service provider network.

Assume that a BPDU is sent from User A network 1 to User A network 2. The BPDU is sent by using the following workflow.

1. At the ingress of the service provider network, PE 1 changes the destination MAC address of the BPDU from 0x0180-C200-0000 to a special multicast MAC address, 0x010F-E200-0003 (the default multicast MAC address), for example. In the service provider network, the modified BPDU is forwarded as a data packet in the VLAN assigned to User A.
2. At the egress of the service provider network, PE 2 recognizes the BPDU with the destination MAC address 0x010F-E200-0003, restores its original destination MAC address 0x0180-C200-0000, and then sends the BPDU to User A network 2.

NOTE:

Be sure, through configuration, that the VLAN tags carried in BPDUs are neither changed nor removed during the transparent transmission in the service provider network. Otherwise, the devices in the service provider network will fail to transparently transmit the customer network BPDUs correctly.

Configuring BPDU tunneling

Configuration prerequisites

Before you configure BPDU tunneling for a protocol, complete the following tasks:

- Enable the protocol in the customer network.
- Assign the port on which you want to enable BPDU tunneling on the PE device and the connected port on the CE device to the same VLAN.
- Configure ports that connect network devices in the service provider network as trunk ports that allow packets of any VLAN to pass through.

Enabling BPDU tunneling

You can enable BPDU tunneling for different protocols in different views.

NOTE:

- Settings made in Ethernet interface view or Layer 2 aggregate interface view take effect only on the current port. Settings made in port group view take effect on all ports in the port group.
- Before you enable BPDU tunneling for DLDP, EOAM, GVRP, HGMP, LLDP, or STP on a port, disable the protocol on the port.
- Because PVST is a special STP protocol, you must do two things before you enable BPDU tunneling for PVST on a port: first, disable STP; second, enable BPDU tunneling for STP on the port.
- Before you enable BPDU tunneling for LACP on a member port of a link aggregation group, remove the port from the link aggregation group.

Enabling BPDU tunneling for a protocol in Ethernet interface view or port group view

Follow these steps to enable BPDU tunneling for a protocol in Ethernet interface view or port group view:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter Ethernet interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Enable BPDU tunneling for a protocol		**bpdu-tunnel dot1q** { **cdp** | **dldp** | **eoam** | **gvrp** | **hgmp** | **lacp** | **lldp** | **pagp** | **pvst** | **stp** | **udld** | **vtp** }	Required Disabled by default.

Enabling BPDU tunneling for a protocol in Layer 2 aggregate interface view

Follow these steps to enable BPDU tunneling for a protocol in Layer 2 aggregate interface view:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	—
Enable BPDU tunneling for a protocol on the Layer 2 aggregate interface	**bpdu-tunnel dot1q** { **cdp** | **gvrp** | **hgmp** | **pvst** | **stp** | **vtp** }	Required Disabled by default.

Configuring destination multicast MAC address for BPDUs

By default, the destination multicast MAC address for BPDUs is 0x010F-E200-0003. You can change it to 0x0100-0CCD-CDD0, 0x0100-0CCD-CDD1, or 0x0100-0CCD-CDD2.

Follow these steps to configure destination multicast MAC address for BPDUs:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Configure the destination multicast MAC address for BPDUs	**bpdu-tunnel tunnel-dmac** *mac-address*	Optional 0x010F-E200-0003 by default.

NOTE:

For BPDUs to be recognized, the destination multicast MAC addresses configured for BPDU tunneling must be the same on the edge devices on the service provider network.

BPDU tunneling configuration examples

BPDU tunneling for STP configuration example

Network requirements

As shown in Figure 29:

- CE 1 and CE 2 are edges devices on the geographically dispersed network of User A; PE 1 and PE 2 are edge devices on the service provider network.
- All ports that connect service provider devices and customer devices are access ports and belong to VLAN 2. All ports that interconnect service provider devices are trunk ports and allow packets of any VLAN to pass through.
- MSTP is enabled on User A's network.

After the configuration, CE 1 and CE 2 must implement consistent spanning tree calculation across the service provider network, and the destination multicast MAC address carried in BPDUs must be 0x0100-0CCD-CDD0.

Figure 29 Network diagram for configuring BPDU tunneling for STP

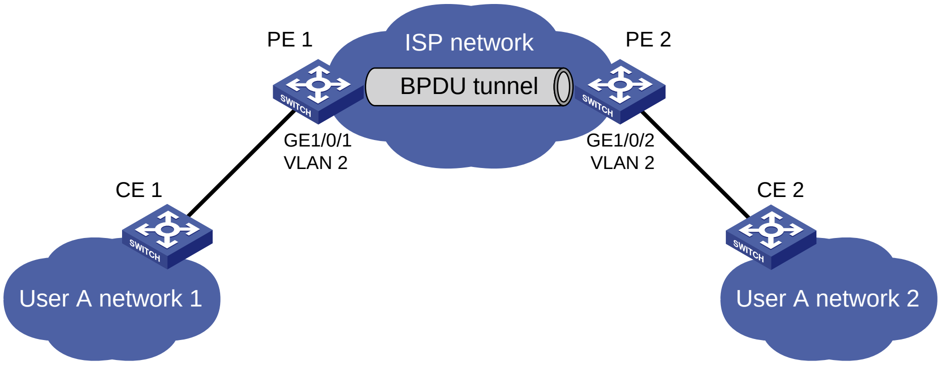

Configuration procedure

1. Configure PE 1.

# Configure the destination multicast MAC address for BPDUs as 0x0100-0CCD-CDD0.

```
<PE1> system-view
[PE1] bpdu-tunnel tunnel-dmac 0100-0ccd-cdd0
```

# Create VLAN 2 and assign GigabitEthernet 1/0/1 to VLAN 2.

```
[PE1] vlan 2
[PE1-vlan2] quit
[PE1] interface gigabitethernet 1/0/1
[PE1-GigabitEthernet1/0/1] port access vlan 2
```

# Disable STP on GigabitEthernet 1/0/1, and then enable BPDU tunneling for STP on it.

```
[PE1-GigabitEthernet1/0/1] undo stp enable
[PE1-GigabitEthernet1/0/1] bpdu-tunnel dot1q stp
```

2. Configure PE 2.

# Configure the destination multicast MAC address for BPDUs as 0x0100-0CCD-CDD0.

```
<PE2> system-view
[PE2] bpdu-tunnel tunnel-dmac 0100-0ccd-cdd0
```

# Create VLAN 2 and assign GigabitEthernet 1/0/2 to VLAN 2.

```
[PE2] vlan 2
[PE2-vlan2] quit
[PE2] interface gigabitethernet 1/0/2
[PE2-GigabitEthernet1/0/2] port access vlan 2
```

# Disable STP on GigabitEthernet 1/0/2, and then enable BPDU tunneling for STP on it.

```
[PE2-GigabitEthernet1/0/2] undo stp enable
[PE2-GigabitEthernet1/0/2] bpdu-tunnel dot1q stp
```

BPDU tunneling for PVST configuration example

Network requirements

As shown in Figure 30:

- CE 1 and CE 2 are edge devices on the geographically dispersed network of User A. PE 1 and PE 2 are edge devices on the service provider network.
- All ports that connect service provider devices and customer devices and those that interconnect service provider devices are trunk ports and allow packets of any VLAN to pass through.
- PVST is enabled for VLANs 1 through 4094 on User A's network.

After the configuration, CE 1 and CE 2 must implement consistent PVST calculation across the service provider network, and the destination multicast MAC address carried in BPDUs must be 0x0100-0CCD-CDD0.

Figure 30 Network diagram for configuring BPDU tunneling for PVST

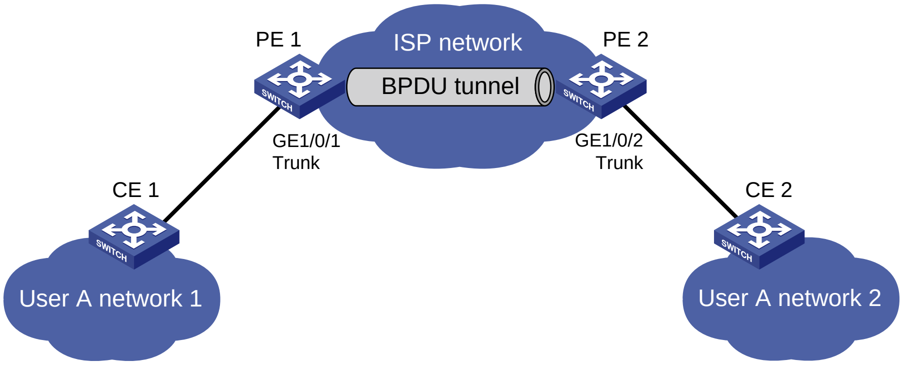

Configuration procedure

1. Configure PE 1.

# Configure the destination multicast MAC address for BPDUs as 0x0100-0CCD-CDD0.

```
<PE1> system-view
[PE1] bpdu-tunnel tunnel-dmac 0100-0ccd-cdd0
```

# Configure GigabitEthernet 1/0/1 as a trunk port and assign it to all VLANs.

```
[PE1] interface gigabitethernet 1/0/1
[PE1-GigabitEthernet1/0/1] port link-type trunk
[PE1-GigabitEthernet1/0/1] port trunk permit vlan all
```

# Disable STP on GigabitEthernet 1/0/1, and then enable BPDU tunneling for STP and PVST on it.

```
[PE1-GigabitEthernet1/0/1] undo stp enable
[PE1-GigabitEthernet1/0/1] bpdu-tunnel dot1q stp
[PE1-GigabitEthernet1/0/1] bpdu-tunnel dot1q pvst
```

2. Configure PE 2.

# Configure the destination multicast MAC address for BPDUs as 0x0100-0CCD-CDD0.

```
<PE2> system-view
[PE2] bpdu-tunnel tunnel-dmac 0100-0ccd-cdd0
```

# Configure GigabitEthernet 1/0/2 as a trunk port and assign it to all VLANs.

```
[PE2] interface gigabitethernet 1/0/2
[PE2-GigabitEthernet1/0/2] port link-type trunk
[PE2-GigabitEthernet1/0/2] port trunk permit vlan all
```

# Disable STP on GigabitEthernet 1/0/2, and then enable BPDU tunneling for STP and PVST on it.

```
[PE2-GigabitEthernet1/0/2] undo stp enable
[PE2-GigabitEthernet1/0/2] bpdu-tunnel dot1q stp
[PE2-GigabitEthernet1/0/2] bpdu-tunnel dot1q pvst
```

VLAN configuration

Introduction to VLAN

VLAN overview

Ethernet is a network technology based on the Carrier Sense Multiple Access/Collision Detect (CSMA/CD) mechanism. Because the medium is shared, collisions and excessive broadcasts are common on Ethernet networks. To address the issue, virtual LAN (VLAN) was introduced to break a LAN down into separate VLANs. VLANs are isolated from each other at Layer 2. A VLAN is a bridging domain, and contains all broadcast traffic within it.

Figure 31 A VLAN diagram

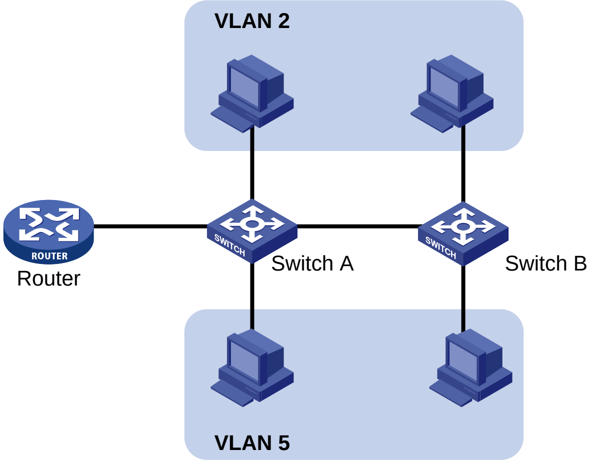

A VLAN is logically divided on an organizational basis rather than on a physical basis. For example, using VLAN, all workstations and servers that a particular workgroup uses can be assigned to the same VLAN, regardless of their physical locations.

VLAN technology delivers the following benefits:

- Confining broadcast traffic within individual VLANs. This reduces bandwidth waste and improves network performance.
- Improving LAN security. By assigning user groups to different VLANs, you can isolate them at Layer 2. To enable communication between VLANs, routers or Layer 3 switches are required.
- Creating flexible virtual workgroups. Because users from the same workgroup can be assigned to the same VLAN regardless of their physical locations, network construction and maintenance are much easier and more flexible.

VLAN fundamentals

To enable a network device to identify frames of different VLANs, a VLAN tag field is inserted into the data link layer encapsulation.

The format of VLAN-tagged frames is defined in IEEE 802.1Q issued by the Institute of Electrical and Electronics Engineers (IEEE) in 1999.

In the header of a traditional Ethernet data frame, the field after the destination MAC address and the source MAC address is the Type field, which indicates the upper layer protocol type, as shown in Figure 32.

Figure 32 Format of a traditional Ethernet frame

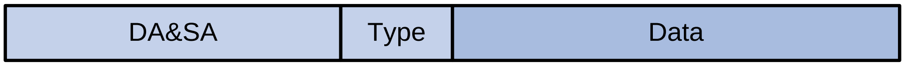

IEEE 802.1Q inserts a four-byte VLAN tag after the DA&SA field.

Figure 33 Position and format of VLAN tag

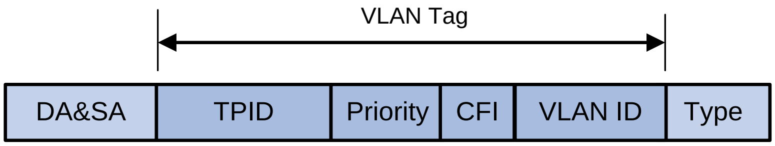

The fields of a VLAN tag are tag protocol identifier (TPID), priority, canonical format indicator (CFI), and VLAN ID.

- The 16-bit TPID field with a value of 0x8100 indicates that the frame is VLAN-tagged.
- The 3-bit priority field indicates the 802.1p priority of the frame. For more information about frame priorities, see the *ACL and QoS Configuration Guide*.
- The 1-bit CFI field specifies whether the MAC addresses are encapsulated in the standard format when packets are transmitted across different media. A value of 0 indicates that MAC addresses are encapsulated in the standard format. A value of 1 indicates that MAC addresses are encapsulated in a non-standard format. The value of the field is 0 by default.
- The 12-bit VLAN ID field identifies the VLAN that the frame belongs to. The VLAN ID range is 0 to 4095. Because 0 and 4095 are reserved, a VLAN ID actually ranges from 1 to 4094.

A network device handles an incoming frame depending on whether the frame is VLAN tagged, and the value of the VLAN tag, if any. For more information, see "Introduction to port-based VLAN."

NOTE:

- The Ethernet II encapsulation format is used here. Besides the Ethernet II encapsulation format, other encapsulation formats, including 802.2 LLC, 802.2 SNAP, and 802.3 raw, are also supported by Ethernet. The VLAN tag fields are also added to frames encapsulated in these formats for VLAN identification.
- When a frame carrying multiple VLAN tags passes through, the switch processes the frame according to its outer VLAN tag, and transmits the inner tags as payload.

Types of VLANs

You can implement VLAN based on the following criteria:

- Port
- MAC address
- Protocol
- IP subnet
- Policy

- Other criteria

NOTE:

- The A5120 EI Switch Series supports port-based VLAN, MAC-based VLAN, protocol-based VLAN, and IP-based VLAN. The port-based VLAN implementation is the basis of all other VLAN implementations. To use any other VLAN implementations, you must configure port-based VLAN settings.
- You can configure all these types of VLANs on a port at the same time. When the switch is determining which VLAN a packet that passes through the port should be assigned to, it looks up the VLANs in the default order of MAC-based VLAN, IP-based VLAN, protocol-based VLAN, and port-based VLAN.

Configuring basic VLAN settings

Follow these steps to configure basic VLAN settings:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Create VLANs	**vlan** { *vlan-id1* [**to** *vlan-id2*] | **all** }	Optional Use this command to create VLANs in bulk.
Enter VLAN view	**vlan** *vlan-id*	Required If the specified VLAN does not exist, this command creates the VLAN first. By default, only the default VLAN (VLAN 1) exists in the system.
Configure a name for the current VLAN	**name** *text*	Optional By default, the name of a VLAN is its VLAN ID (**VLAN 0001**, for example).
Configure the description of the current VLAN	**description** *text*	Optional VLAN ID is used by default (**VLAN 0001**, for example).

NOTE:

- As the default VLAN, VLAN 1 cannot be created or removed.
- You cannot manually create or remove VLANs reserved for special purposes.
- You cannot remove dynamic VLANs by using the **undo vlan** command.
- You cannot remove a VLAN with a QoS policy applied.
- After associating an isolate-user-VLAN with a secondary VLAN, you cannot add ports to, remove ports from, or remove the VLANs. To do that, remove the association first.
- You cannot remove a VLAN that is operating as a probe VLAN for remote port mirroring, or an RRPP-protected VLAN by using the **undo vlan** command. To do that, remove the remote probe VLAN or RRPP-protected VLAN configuration from it first.

Configuring basic settings of a VLAN interface

For hosts of different VLANs to communicate, you must use a router or Layer 3 switch to perform layer 3 forwarding. You use VLAN interfaces to achieve this.

VLAN interfaces are virtual interfaces used for Layer 3 communication between different VLANs. They do not exist as physical entities on devices. For each VLAN, you can create one VLAN interface. You can assign the VLAN interface an IP address and specify it as the gateway of the VLAN to forward traffic destined for an IP network segment different from that of the VLAN.

Follow these steps to configure basic settings of a VLAN interface:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Create a VLAN interface and enter VLAN interface view	**interface vlan-interface** *vlan-interface-id*	Required If the VLAN interface already exists, you enter its view directly.
Assign an IP address to the VLAN interface	**ip address** *ip-address* { *mask* | *mask-length* } [**sub**]	Optional No IP address is assigned to any VLAN interface by default.
Configure the description of the VLAN interface	**description** *text*	Optional VLAN interface name is used by default, for example, **Vlan-interface1 Interface.**
Bring up the VLAN interface	**undo shutdown**	Optional By default, a VLAN interface is in the up state. The VLAN interface is up as long as one port in the VLAN is up and goes down if all ports in the VLAN go down. An administratively shut down VLAN interface is in the down state until you bring it up, regardless of how the state of the ports in the VLAN changes.

NOTE:

Before you create a VLAN interface for a VLAN, create the VLAN.

Port-based VLAN configuration

Introduction to port-based VLAN

Port-based VLANs group VLAN members by port. A port forwards traffic for a VLAN only after it is assigned to the VLAN.

Port link type

You can configure the link type of a port as access, trunk, or hybrid. The link types use the following VLAN tag handling methods:

- An access port belongs to only one VLAN and sends traffic untagged. It is usually used to connect a terminal device unable to identify VLAN tagged-packets or when separating different VLAN members is unnecessary.
- A trunk port can carry multiple VLANs to receive and send traffic for them. Except traffic of the port VLAN ID (PVID), traffic sent through a trunk port will be VLAN tagged. Usually, ports that connect network devices are configured as trunk ports.
- Like a trunk port, a hybrid port can carry multiple VLANs to receive and send traffic for them. Unlike a trunk port, a hybrid port allows traffic of all VLANs to pass through VLAN untagged.

Port VLAN ID (PVID)

By default, VLAN 1 is the PVID for all ports. You can configure the PVID for a port as required.

When configuring the PVID on a port, use the following guidelines:

- Because an access port can join only one VLAN, its PVID is the VLAN to which it belongs, and it cannot be configured.
- Because a trunk or hybrid port can join multiple VLANs, you can configure a PVID for the port.
- You can use a nonexistent VLAN as the PVID for a hybrid or trunk port but not for an access port. After you use the **undo vlan** command to remove the VLAN that an access port resides in, the PVID of the port changes to VLAN 1. The removal of the VLAN specified as the PVID of a trunk or hybrid port, however, does not affect the PVID setting on the port.

NOTE:

- Do not set the voice VLAN as the PVID of a port in automatic voice VLAN assignment mode. For information about voice VLAN, see the chapter "Voice VLAN configuration."
- HP recommends that you set the same PVID for the local and remote ports.
- Make sure that a port is assigned to its PVID. Otherwise, when the port receives frames tagged with the PVID or untagged frames (including protocol packets such as MSTP BPDUs), the port filters out these frames.

The following table shows how ports of different link types handle frames:

Port type	Actions (in the inbound direction)		Actions (in the outbound direction)
	Untagged frame	**Tagged frame**	
Access	Tag the frame with the PVID.	• Receive the frame if its VLAN ID is the same as the PVID. • Drop the frame if its VLAN ID is different from the PVID.	Remove the VLAN tag and send the frame.
Trunk	Determine whether the PVID is permitted on the port, as follows: • If yes, tag the frame with the PVID. • If not, drop the frame.	• Receive the frame if its VLAN is carried on the port. • Drop the frame if its VLAN is not carried on the port.	• Remove the tag and send the frame if the frame carries the PVID and the port belongs to the PVID. • Send the frame without removing the tag if its VLAN is carried on the port but is different from the default one.

Port type	Actions (in the inbound direction)		Actions (in the outbound direction)
	Untagged frame	Tagged frame	
Hybrid			Send the frame if its VLAN is carried on the port. The frame is sent with the VLAN tag removed or intact depending on your configuration via the **port hybrid vlan** command. This is true of the PVID.

Assigning an access port to a VLAN

You can assign an access port to a VLAN in VLAN view, interface view (including Ethernet interface view and Layer 2 aggregate interface view), or port group view.

Follow these steps to assign one or multiple access ports to a VLAN in VLAN view:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Enter VLAN view	**vlan** *vlan-id*	Required If the specified VLAN does not exist, this command creates the VLAN first.
Assign one or a group of access ports to the current VLAN	**port** *interface-list*	Required By default, all ports belong to VLAN 1.

Follow these steps to assign an access port (in interface view) or multiple access ports (in port group view) to a VLAN:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter interface view (including Ethernet interface view, Layer 2 aggregate interface view) or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command. • In Ethernet interface view, the subsequent configurations apply to the current port. • In port group view, the subsequent configurations apply to all ports in the port group. • In Layer 2 aggregate interface view, the subsequent configurations apply to the Layer 2 aggregate interface and all its member ports.
	Enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	
	Enter port group view	**port-group manual** *port-group-name*	
Configure the link type of the port or ports as access		**port link-type access**	Optional The link type of a port is access by default.

To do…	Use the command…	Remarks
Assign the current access port(s) to a VLAN	**port access vlan** *vlan-id*	Optional By default, all access ports belong to VLAN 1.

NOTE:

- Before you assign an access port to a VLAN, create the VLAN.
- In VLAN view, you can assign only Layer 2 Ethernet interfaces to the current VLAN.
- After you configure a command on a Layer 2 aggregate interface, the system starts applying the configuration to the aggregate interface and its aggregation member ports. If the system fails to do that on the aggregate interface, it stops applying the configuration to the aggregation member ports. If it fails to do that on an aggregation member port, it skips the port and moves to the next port.

Assigning a trunk port to a VLAN

A trunk port can carry multiple VLANs. You can assign it to a VLAN in interface view (including Ethernet interface view, Layer 2 aggregate interface view) or port group view.

Follow these steps to assign a trunk port to one or multiple VLANs:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter interface view (including Ethernet interface view, Layer 2 aggregate interface view) or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command. • In Ethernet interface view, the subsequent configurations apply to the current port. • In port group view, the subsequent configurations apply to all ports in the port group. • In Layer 2 aggregate interface view, the subsequent configurations apply to the Layer 2 aggregate interface and all its member ports.
	Enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	
	Enter port group view	**port-group manual** *port-group-name*	
Configure the link type of the port or ports as trunk		**port link-type trunk**	Required
Assign the trunk port(s) to the specified VLAN(s)		**port trunk permit vlan** { *vlan-id-list* | **all** }	Required By default, a trunk port carries only VLAN 1.
Configure the PVID of the trunk port(s)		**port trunk pvid vlan** *vlan-id*	Optional VLAN 1 is the PVID by default.

NOTE:

- To change the link type of a port from trunk to hybrid or vice versa, you must set the link type to access first.
- After configuring the PVID for a trunk port, you must use the **port trunk permit vlan** command to configure the trunk port to allow packets from the PVID to pass through, so that the egress port can forward packets from the PVID.
- After you use the **port link-type** { **access** | **hybrid** | **trunk** } command to change the link type of an interface, the loopback detection action configured on the interface by using the **loopback-detection action** command will be restored to the default. For more information about the **loopback-detection action** command, see the *Layer 2—LAN Switching Command Reference*.
- After you configure a command on a Layer 2 aggregate interface, the system starts applying the configuration to the aggregate interface and its aggregation member ports. If the system fails to do that on the aggregate interface, it stops applying the configuration to the aggregation member ports. If it fails to do that on an aggregation member port, it skips the port and moves to the next port.

Assigning a hybrid port to a VLAN

A hybrid port can carry multiple VLANs. You can assign it to a VLAN in interface view (including Ethernet interface view, Layer 2 aggregate interface view) or port group view.

Follow these steps to assign a hybrid port to one or multiple VLANs:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter interface view (including Ethernet interface view, Layer 2 aggregate interface view) or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command. • In Ethernet interface view, the subsequent configurations apply to the current port. • In port group view, the subsequent configurations apply to all ports in the port group. • In Layer 2 aggregate interface view, the subsequent configurations apply to the Layer 2 aggregate interface and all its member ports.
	Enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	
	Enter port group view	**port-group manual** *port-group-name*	
Configure the link type of the port(s) as hybrid		**port link-type hybrid**	Required
Assign the hybrid port(s) to the specified VLAN(s)		**port hybrid vlan** *vlan-id-list* { **tagged** | **untagged** }	Required By default, a hybrid port allows only packets of VLAN 1 to pass through untagged.
Configure the PVID of the hybrid port		**port hybrid pvid vlan** *vlan-id*	Optional VLAN 1 is the PVID by default.

NOTE:

- To change the link type of a port from trunk to hybrid or vice versa, you must set the link type to access first.
- After you use the **port link-type** { **access** | **hybrid** | **trunk** } command to change the link type of an interface, the loopback detection action configured on the interface by using the **loopback-detection action** command will be restored to the default. For more information about the **loopback-detection action** command, see the *Layer 2—LAN Switching Command Reference*.
- Before you assign a hybrid port to a VLAN, create the VLAN.
- After configuring the PVID for a hybrid port, you must use the **port hybrid vlan** command to configure the hybrid port to allow packets from the PVID to pass through, so that the egress port can forward packets from the PVID.
- After you configure a command on a Layer 2 aggregate interface, the system starts applying the configuration to the aggregate interface and its aggregation member ports. If the system fails to do that on the aggregate interface, it stops applying the configuration to the aggregation member ports. If it fails to do that on an aggregation member port, it skips the port and moves to the next port.

Port-based VLAN configuration example

Network requirements

As shown in Figure 34:

- Host A and Host C belong to Department A, and access the enterprise network through different devices. Host B and Host D belong to Department B. They also access the enterprise network through different devices.
- To ensure communication security and avoid broadcast storms, VLANs are configured in the enterprise network to isolate Layer 2 traffic of different departments. VLAN 100 is assigned to Department A, and VLAN 200 is assigned to Department B.
- Ensure that hosts within the same VLAN can communicate with each other. Host A can communicate with Host C, and Host B can communicate with Host D.

Figure 34 Network diagram for port-based VLAN configuration

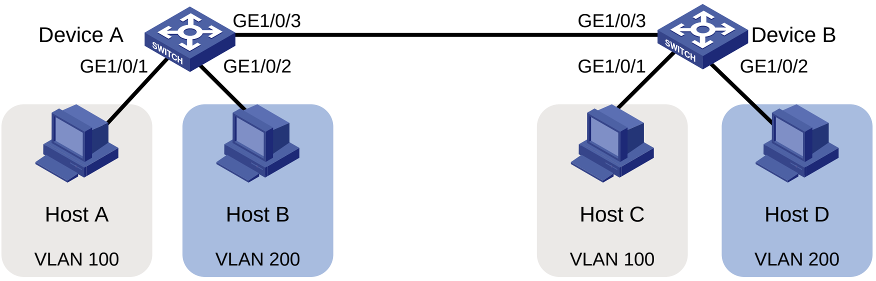

Configuration procedure

1. Configure Device A.

# Create VLAN 100, and assign port GigabitEthernet 1/0/1 to VLAN 100.

```
<DeviceA> system-view
[DeviceA] vlan 100
[DeviceA-vlan100] port gigabitethernet 1/0/1
[DeviceA-vlan100] quit
```

# Create VLAN 200, and assign port GigabitEthernet 1/0/2 to VLAN 200.

```
[DeviceA] vlan 200
```

```
[DeviceA-vlan200] port gigabitethernet 1/0/2
[DeviceA-vlan200] quit
```

# Configure port GigabitEthernet 1/0/3 as a trunk port, and assign it to VLANs 100 and 200, to enable GigabitEthernet 1/0/3 to forward traffic of VLANs 100 and 200 to Device B.

```
[DeviceA] interface gigabitethernet 1/0/3
[DeviceA-GigabitEthernet1/0/3] port link-type trunk
[DeviceA-GigabitEthernet1/0/3] port trunk permit vlan 100 200
Please wait... Done.
```

2. Configure Device B as you configure Device A.
3. Configure Host A and Host C to be on the same network segment, 192.168.100.0/24, for example. Configure Host B and Host D to be on the same network segment, 192.168.200.0/24, for example.

Verification

1. Host A and Host C and ping each other successfully, but they both fail to ping Host B. Host B and Host D and ping each other successfully, but they both fail to ping Host A.
2. Determine whether the configuration is successful by displaying relevant VLAN information.

# Display information about VLANs 100 and 200 on Device A:

```
[DeviceA-GigabitEthernet1/0/3] display vlan 100
 VLAN ID: 100
 VLAN Type: static
 Route Interface: not configured
 Description: protocol VLAN for IPv4
 Name: VLAN 0100
 Tagged   Ports:
    GigabitEthernet1/0/3
 Untagged Ports:
    GigabitEthernet1/0/1
[DeviceA-GigabitEthernet1/0/3] display vlan 200
 VLAN ID: 200
 VLAN Type: static
 Route Interface: not configured
 Description: protocol VLAN for IPv6
 Name: VLAN 0200
 Tagged   Ports:
    GigabitEthernet1/0/3
 Untagged Ports:
    GigabitEthernet1/0/2
```

MAC-based VLAN configuration

Introduction to MAC-based VLAN

The MAC-based VLAN feature assigns hosts to a VLAN based on their MAC addresses. The following approaches are available for configuring MAC-based VLANs:

Approach 1: Static MAC-based VLAN assignment

Static MAC-based VLAN assignment applies to networks containing a small number of VLAN users. In such a network, you can create a MAC address-to-VLAN map containing multiple MAC address-to-VLAN entries on a port, enable the MAC-based VLAN feature on the port, and assign the port to MAC-based VLANs.

With static MAC-based VLAN assignment configured on a port, the switch processes received frames by using the following guidelines:

- When the port receives an untagged frame, the switch looks up the MAC address-to-VLAN map based on the source MAC address of the frame for a match. If the MAC address of a MAC address-to-VLAN entry matches the source MAC address of the untagged frame, the switch tags the frame with the corresponding VLAN ID. If no match is found, the switch assigns a VLAN to the frame by using the following criteria in turn: IP addresses, protocols, and ports.
- When the port receives a tagged frame, the port forwards the frame if the VLAN ID of the frame is permitted by the port, or otherwise drops the frame.

Approach 2: Dynamic MAC-based VLAN assignment

When you cannot determine the target MAC-based VLANs of a port, you can use dynamic MAC-based VLAN assignment on the port. To do that, you can create a MAC address-to-VLAN map containing multiple MAC address-to-VLAN entries, enable the MAC-based VLAN feature and dynamic MAC-based VLAN assignment on the port. When the port receives a frame that matches a MAC address-to-VLAN entry configured on the port, the port dynamically joins the corresponding MAC-based VLAN.

The following workflows apply:

- When the port receives an untagged frame, it processes the frame by using the flowchart as shown in Figure 35.

Figure 35 Flowchart for processing an untagged frame in dynamic MAC-based VLAN assignment

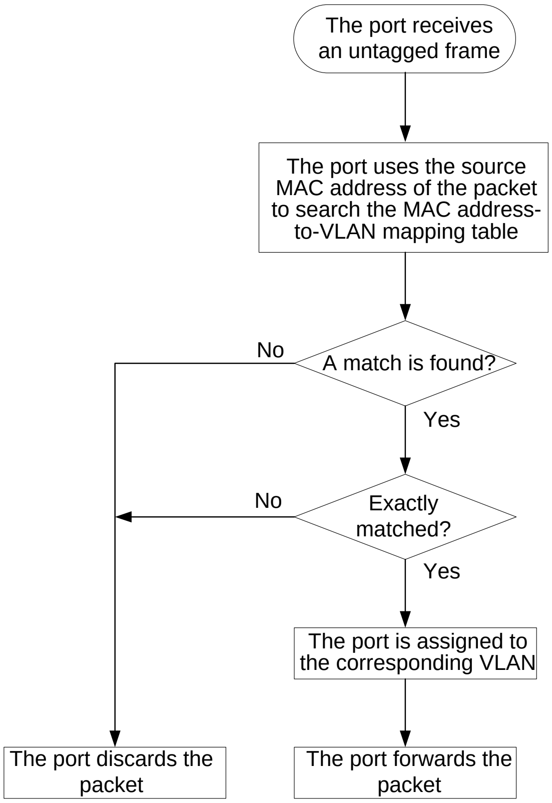
The port receives an untagged frame
The port uses the source MAC address of the packet to search the MAC address-to-VLAN mapping table
No
A match is found?
Yes
No
Exactly matched?
Yes
The port is assigned to the corresponding VLAN
The port discards the packet
The port forwards the packet

- When the port receives a tagged frame, the port forwards the frame if the VLAN ID of the frame is permitted by the port, or otherwise drops the frame.

NOTE:

If you configure both static and dynamic MAC-based VLAN assignment on the same port, dynamic MAC-based VLAN assignment applies, and the port drops the frames that do not exactly match any MAC address-to-VLAN entry.

Approach 3: Dynamic MAC-based VLAN

You can use dynamic MAC-based VLAN with access authentication (such as 802.1X authentication based on MAC addresses) to implement secure, flexible terminal access. After configuring dynamic MAC-based VLAN on the switch, you must configure the MAC address-to-VLAN entries on the access authentication server.

When a user passes authentication of the access authentication server, the switch obtains VLAN information from the server, generates a MAC address-to-VLAN entry by using the source MAC address of the user packet and the VLAN information, and assigns the port to the MAC-based VLAN. When the user goes offline, the switch automatically deletes the MAC address-to-VLAN entry, and removes the port from the MAC-based VLAN.

Configuring MAC-based VLAN

NOTE:

- MAC-based VLANs are available only on hybrid ports.
- The MAC-based VLAN feature is mainly configured on the downlink ports of the user access devices. Do not enable this function together with link aggregation.
- After associating MAC addresses with a VLAN, if you specify the 802.1p priority value corresponding to the specified MAC addresses, you must use the **qos trust dot1p** command in interface view to configure the interface to use the 802.1p priority in incoming packets for priority mapping. For more information about this command, see the *ACL and QoS Configuration Guide.*

Configuring static MAC-based VLAN assignment

Follow these steps to configure static MAC-based VLAN assignment:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Associate MAC addresses with a VLAN		**mac-vlan mac-address** *mac-address* **vlan** *vlan-id* [**priority** *priority*]	Required
Enter Ethernet interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Use either command. • The configuration made in Ethernet interface view applies only to the current port. • The configuration made in port group view applies to all ports in the port group.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the link type of the port(s) as hybrid		**port link-type hybrid**	Required
Configure the hybrid port(s) to permit packets of specific MAC-based VLANs to pass through		**port hybrid vlan** *vlan-id-list* { **tagged** | **untagged** }	Required By default, a hybrid port only permits the packets of VLAN 1 to pass through.
Enable MAC-based VLAN		**mac-vlan enable**	Required Disabled by default
Configure VLAN matching precedence		**vlan precedence** { **mac-vlan** | **ip-subnet-vlan** }	Optional By default, VLANs are preferentially matched based on MAC addresses.

Configuring dynamic MAC-based VLAN assignment

NOTE:

- With dynamic MAC-based VLAN assignment enabled, packets are delivered to the CPU for processing. The packet processing mode has the highest priority and overrides the configuration of MAC learning limit and disabling of MAC address learning. When dynamic MAC-based VLAN assignment is enabled, do not configure the MAC learning limit or disable MAC address learning.
- Do not use dynamic MAC-based VLAN assignment together with 802.X and MAC authentication.
- In dynamic MAC-based VLAN assignment, the port that receives a packet with an unknown source MAC address can be successfully assigned to the matched VLAN only when the matched VLAN is a static VLAN.
- With MSTP enabled, if a port is blocked in the MST instance (MSTI) of the target MAC-based VLAN, the port drops the received packets, instead of delivering them to the CPU. As a result, the receiving port will not be dynamically assigned to the corresponding VLAN. Do not configure dynamic MAC-based VLAN assignment together with MSTP, because the former is mainly configured on the access side.
- When a MAC address ages, the receiving port automatically leaves the VLAN to which it was dynamically assigned to. For more information about MAC address aging, see the chapter "MAC address table configuration."

Follow these steps to configure dynamic MAC-based VLAN assignment:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Associate MAC addresses with a VLAN		**mac-vlan mac-address** *mac-address* **vlan** *vlan-id* [**priority** *priority*]	Required
Enter Ethernet interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Use either command. • The configuration made in Ethernet interface view applies only to the current port. • The configuration made in port group view applies to all ports in the port group.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the link type of the port(s) as hybrid		**port link-type hybrid**	Required
Enable MAC-based VLAN		**mac-vlan enable**	Required Disabled by default
Configure VLAN matching precedence		**vlan precedence** { **mac-vlan** | **ip-subnet-vlan** }	Optional By default, VLANs are preferably matched based on MAC addresses.
Enable dynamic MAC-based VLAN assignment		**mac-vlan trigger enable**	Required Disabled by default
Disable the PVID of the port from forwarding source-unknown packets that do not match any MAC address-to-VLAN mapping		**port pvid disable**	Optional By default, source MAC unknown packets are forwarded in the PVID of the incoming port if they do not match any MAC address-to-VLAN mapping.

Configuring dynamic MAC-based VLAN

NOTE:

After enabling MAC-based VLAN on the switch, you must configure related authentication settings on the access authentication server. For more information about access authentication, see the *Security Configuration Guide*.

Follow these steps to configure dynamic MAC-based VLAN:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter Ethernet interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Use either command. • The configuration made in Ethernet interface view applies only to the current port. • The configuration made in port group view applies to all ports in the port group.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the link type of the port(s) as hybrid		**port link-type hybrid**	Required
Enable MAC-based VLAN		**mac-vlan enable**	Required Disabled by default

MAC-based VLAN configuration example

Network requirements

- As shown in Figure 36, GigabitEthernet 1/0/1 of Device A and Device C are each connected to a meeting room. Laptop 1 and Laptop 2 are used for meetings and might be used in either of the two meeting rooms.
- Different departments own Laptop 1 and Laptop 2. The two departments use VLAN 100 and VLAN 200 respectively. Each laptop must be able to access only its own department server, no matter which meeting room it is used in.
- The MAC address of Laptop 1 is 000d-88f8-4e71, and that of Laptop 2 is 0014-222c-aa69.

Figure 36 Network diagram for MAC-based VLAN configuration

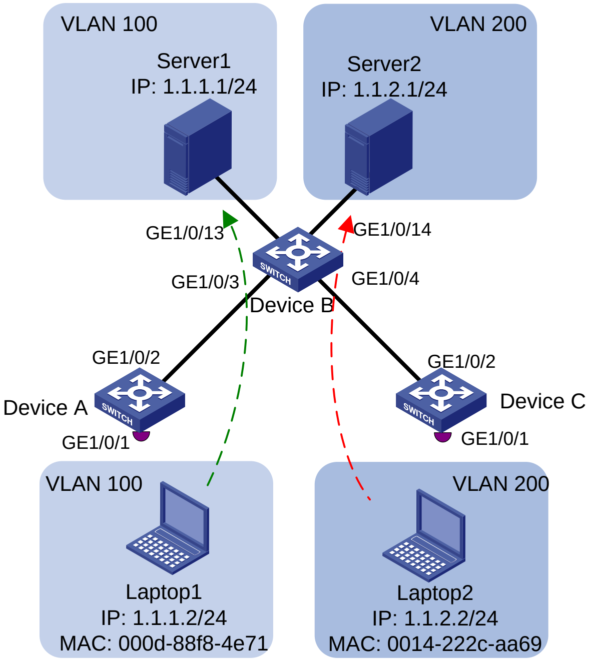

Configuration consideration

- Create VLANs 100 and 200.
- Configure the uplink ports of Device A and Device C as trunk ports, and assign them to VLANs 100 and 200.
- Configure the downlink ports of Device B as trunk ports, and assign them to VLANs 100 and 200. Configure the uplink ports of Device B as access ports connecting to the servers respectively, and assign them to VLANs 100 and 200 respectively.
- Associate the MAC address of Laptop 1 with VLAN 100, and associate the MAC address of Laptop 2 with VLAN 200.

Configuration procedure

1. Configure Device A.

# Create VLANs 100 and 200.

```
<DeviceA> system-view
[DeviceA] vlan 100
[DeviceA-vlan100] quit
[DeviceA] vlan 200
[DeviceA-vlan200] quit
```

# Associate the MAC address of Laptop 1 with VLAN 100, and associate the MAC address of Laptop 2 with VLAN 200.

```
[DeviceA] mac-vlan mac-address 000d-88f8-4e71 vlan 100
[DeviceA] mac-vlan mac-address 0014-222c-aa69 vlan 200
```

# Configure Laptop 1 and Laptop 2 to access the network through GigabitEthernet 1/0/1. Configure GigabitEthernet 1/0/1 as a hybrid port that sends packets of VLANs 100 and 200 untagged, and enable MAC-based VLAN on it.

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

```
[DeviceA-GigabitEthernet1/0/1] port link-type hybrid
[DeviceA-GigabitEthernet1/0/1] port hybrid vlan 100 200 untagged
 Please wait... Done.
[DeviceA-GigabitEthernet1/0/1] mac-vlan enable
[DeviceA-GigabitEthernet1/0/1] quit
```

# To enable the laptops to access Server 1 and Server 2, configure the uplink port GigabitEthernet 1/0/2 as a trunk port, and assign it to VLANs 100 and 200.

```
[DeviceA] interface gigabitethernet 1/0/2
[DeviceA-GigabitEthernet1/0/2] port link-type trunk
[DeviceA-GigabitEthernet1/0/2] port trunk permit vlan 100 200
[DeviceA-GigabitEthernet1/0/2] quit
```

2. Configure Device B.

# Create VLANs 100 and 200. Assign GigabitEthernet 1/0/13 to VLAN 100, and assign GigabitEthernet 1/0/14 to VLAN 200.

```
<DeviceB> system-view
[DeviceB] vlan 100
[DeviceB-vlan100] port gigabitethernet 1/0/13
[DeviceB-vlan100] quit
[DeviceB] vlan 200
[DeviceB-vlan200] port gigabitethernet 1/0/14
[DeviceB-vlan200] quit
```

# Configure GigabitEthernet 1/0/3 and GigabitEthernet 1/0/4 as trunk ports, and assign them to VLANs 100 and 200.

```
[DeviceB] interface gigabitethernet 1/0/3
[DeviceB-GigabitEthernet1/0/3] port link-type trunk
[DeviceB-GigabitEthernet1/0/3] port trunk permit vlan 100 200
[DeviceB-GigabitEthernet1/0/3] quit
[DeviceB] interface gigabitethernet 1/0/4
[DeviceB-GigabitEthernet1/0/4] port link-type trunk
[DeviceB-GigabitEthernet1/0/4] port trunk permit vlan 100 200
[DeviceB-GigabitEthernet1/0/4] quit
```

3. Configure Device C.

Configure Device C as you configure Device A.

Verification

1. Laptop 1 can access Server 1 only, and Laptop 2 can access Server 2 only.
2. On Device A and Device C, you can see that VLAN 100 is associated with the MAC address of Laptop 1, and VLAN 200 is associated with the MAC address of Laptop 2.

```
[DeviceA] display mac-vlan all
  The following MAC VLAN addresses exist:
  S:Static  D:Dynamic
  MAC ADDR         MASK             VLAN ID   PRIO   STATE
  --------------------------------------------------------
  000d-88f8-4e71   ffff-ffff-ffff   100       0      S
  0014-222c-aa69   ffff-ffff-ffff   200       0      S
```

```
Total MAC VLAN address count:2
```

Configuration guidelines

1. MAC-based VLAN can be configured only on hybrid ports.
2. MAC-based VLAN is typically configured on the downlink ports of access layer devices, and cannot be configured together with the link aggregation function.

Protocol-based VLAN configuration

Introduction to protocol-based VLAN

NOTE:

Protocol-based VLAN configuration applies only to hybrid ports.

In this approach, inbound packets are assigned to different VLANs based on their protocol types and encapsulation formats. The protocols that can be used for VLAN assignment include IP, IPX, and AppleTalk (AT). The encapsulation formats include Ethernet II, 802.3 raw, 802.2 LLC, and 802.2 SNAP.

A protocol type and an encapsulation format compose a protocol template. You can create multiple protocol templates for a protocol-based VLAN, and different protocol templates are assigned different *protocol-index* values. A protocol-based VLAN ID and a protocol index, combined, can uniquely identify a protocol template. When you use commands to associate protocol templates with ports, use *protocol-based vlan-id + protocol index* to specify the protocol templates. An untagged packet that reaches a port associated with protocol templates will be processed using the following workflow:

- If the protocol type and encapsulation format carried in the packet matches a protocol template, the packet will be tagged with the VLAN tag that corresponds to the protocol template.
- If the packet matches no protocol templates, the packet will be tagged with the PVID of the port.

The port processes a tagged packet as it processes tagged packets of a port-based VLAN.

- If the port is assigned to the VLAN that corresponds to the VLAN tag carried in the packet, it forwards the packet.
- If not, it drops the packet.

This feature is mainly used to assign packets of the specific service type to a specific VLAN.

Configuring a protocol-based VLAN

Follow these steps to configure a protocol-based VLAN:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Enter VLAN view	**vlan** *vlan-id*	Required If the specified VLAN does not exist, this command creates the VLAN first.

To do…		Use the command…	Remarks
Create a protocol template for the VLAN		**protocol-vlan** [*protocol-index*] { **at** | **ipv4** | **ipv6** | **ipx** { **ethernetii** | **llc** | **raw** | **snap** } | **mode** { **ethernetii etype** *etype-id* | **llc** { **dsap** *dsap-id* [**ssap** *ssap-id*] | **ssap** *ssap-id* } | **snap etype** *etype-id* } }	Required
Exit VLAN view		**quit**	Required
Enter interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command. - In Ethernet interface view, the subsequent configurations apply to the current port. - In port group view, the subsequent configurations apply to all ports in the port group. - In Layer 2 aggregate interface view, the subsequent configurations apply to the Layer 2 aggregate interface and all its member ports.
	Enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	
	Enter port group view	**port-group manual** *port-group-name*	
Configure the port link type as hybrid		**port link-type hybrid**	Required
Configure current hybrid port(s) to permit the packets of the specified protocol-based VLANs to pass through		**port hybrid vlan** *vlan-id-list* { **tagged** | **untagged** }	Required By default, all hybrid ports permit only packets of VLAN 1 to pass through.
Associate the hybrid port(s) with the specified protocol-based VLAN		**port hybrid protocol-vlan vlan** *vlan-id* { *protocol-index* [**to** *protocol-end*] | **all** }	Required

△ CAUTION:

- Do not configure both the *dsap-id* and *ssap-id* arguments in the **protocol-vlan** command as 0xe0 or 0xff when configuring the user-defined template for **llc** encapsulation. Otherwise, the encapsulation format of the matching packets will be the same as that of the **ipx llc** or **ipx raw** packets respectively.
- When you use the **mode** keyword to configure a user-defined protocol template, do not set *etype-id* in **ethernetii etype** *etype-id* to 0x0800, 0x8137, 0x809b, or 0x86dd. Otherwise, the encapsulation format of the matching packets will be the same as that of the IPv4, IPX, AppleTalk, and IPv6 packets respectively.
- A protocol-based VLAN on a hybrid port can process only untagged inbound packets, whereas the voice VLAN in automatic mode on a hybrid port can process only tagged voice traffic. Do not configure a VLAN as both a protocol-based VLAN and a voice VLAN. For more information, see the chapter "Voice VLAN configuration."
- After you configure a command on a Layer 2 aggregate interface, the system starts applying the configuration to the aggregate interface and its aggregation member ports. If the system fails to do that on the aggregate interface, it stops applying the configuration to the aggregation member ports. If it fails to do that on an aggregation member port, it skips the port and moves to the next port.
- If a protocol template using SNAP encapsulation is configured for a specific VLAN, the SNAP packets with an all-zero organizationally unique identifier (OUI) can be assigned to the VLAN, whereas the SNAP packets with non-all-zero OUIs cannot be assigned to the VLAN.

Protocol-based VLAN configuration example

Network requirements

In a lab environment, most hosts run the IPv4 protocol, and the rest of the hosts run the IPv6 protocol for teaching purposes. To avoid interference, isolate IPv4 traffic and IPv6 traffic at Layer 2.

Figure 37 Network diagram for protocol-based VLAN configuration

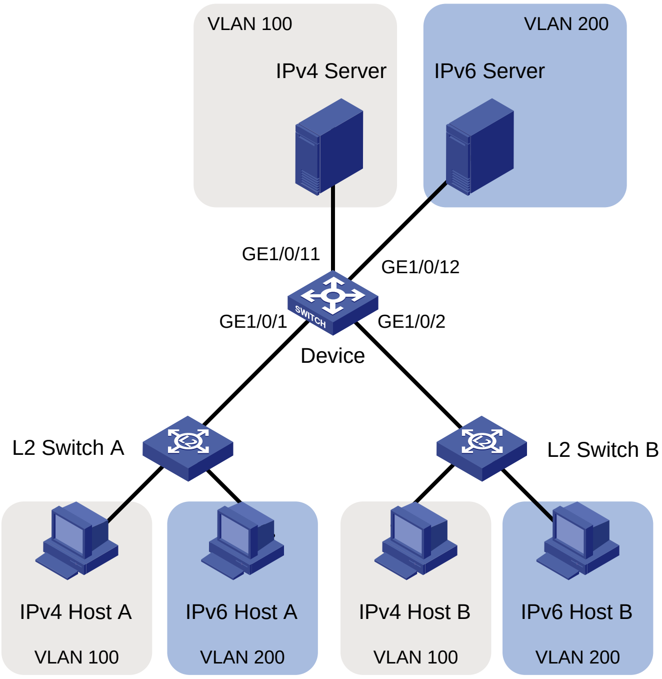

Configuration consideration

Create VLANs 100 and 200. Associate VLAN 100 with IPv4, and associate VLAN 200 with IPv6. Configure protocol-based VLANs to isolate IPv4 traffic and IPv6 traffic at Layer 2.

Configuration procedure

1. Configure Device.

# Create VLAN 100, and assign port GigabitEthernet 1/0/11 to VLAN 100.

```
<Device> system-view
[Device] vlan 100
[Device-vlan100] description protocol VLAN for IPv4
[Device-vlan100] port gigabitethernet 1/0/11
[Device-vlan100] quit
```

# Create VLAN 200, and assign port GigabitEthernet 1/0/12 to VLAN 200.

```
[Device] vlan 200
[Device-vlan200] description protocol VLAN for IPv6
[Device-vlan200] port gigabitethernet 1/0/12
```

# Create an IPv6 protocol template in the view of VLAN 200, and create an IPv4 protocol template in the view of VLAN 100.

```
[Device-vlan200] protocol-vlan 1 ipv6
[Device-vlan200] quit
[Device] vlan 100
[Device-vlan100] protocol-vlan 1 ipv4
[Device-vlan100] quit
```

# Configure port GigabitEthernet 1/0/1 as a hybrid port that forwards packets of VLANs 100 and 200 untagged.

```
[Device] interface gigabitethernet 1/0/1
[Device-GigabitEthernet1/0/1] port link-type hybrid
[Device-GigabitEthernet1/0/1] port hybrid vlan 100 200 untagged
 Please wait... Done.
```

# Associate port GigabitEthernet 1/0/1 with the IPv4 protocol template of VLAN 100 and the IPv6 protocol template of VLAN 200.

```
[Device-GigabitEthernet1/0/1] port hybrid protocol-vlan vlan 100 1
[Device-GigabitEthernet1/0/1] port hybrid protocol-vlan vlan 200 1
[Device-GigabitEthernet1/0/1] quit
```

# Configure GigabitEthernet 1/0/2 as a hybrid port that forwards packets of VLANs 100 and 200 untagged, and associate GigabitEthernet 1/0/2 with the IPv4 protocol template of VLAN 100 and the IPv6 protocol template of VLAN 200.

```
[Device] interface gigabitethernet 1/0/2
[Device-GigabitEthernet1/0/2] port link-type hybrid
[Device-GigabitEthernet1/0/2] port hybrid vlan 100 200 untagged
 Please wait... Done.
[Device-GigabitEthernet1/0/2] port hybrid protocol-vlan vlan 100 1
[Device-GigabitEthernet1/0/2] port hybrid protocol-vlan vlan 200 1
```

2. Keep the default settings of L2 Switch A and L2 Switch B.
3. Configure IPv4 Host A, IPv4 Host B, and IPv4 Server to be on the same network segment (192.168.100.0/24, for example), and configure IPv6 Host A, IPv6 Host B, and IPv6 Server to be on the same network segment (192.168.200.0/24, for example).

Verification

1. The hosts and the server in VLAN 100 can ping one another successfully. The hosts and the server in VLAN 200 can ping one another successfully. The hosts/server in VLAN 100 cannot ping the hosts and the server in VLAN 200, and vice versa.
2. Display protocol-based VLAN information on Device to determine whether the configurations have become valid.

# Display protocol-based VLAN configuration on Device.

```
[Device-GigabitEthernet1/0/2] display protocol-vlan vlan all
 VLAN ID:100
    Protocol Index      Protocol Type
 ======================================================
          1              ipv4
 VLAN ID:200
    Protocol Index      Protocol Type
 ======================================================
          1              ipv6
```

# Display protocol-based VLAN information on the ports of Device.

```
[Device-GigabitEthernet1/0/2] display protocol-vlan interface all
 Interface: GigabitEthernet 1/0/1
   VLAN ID   Protocol Index      Protocol Type
 ======================================================
     100          1               ipv4
     200          1               ipv6
 Interface: GigabitEthernet 1/0/2
   VLAN ID   Protocol Index      Protocol Type
 ======================================================
     100          1               ipv4
     200          1               ipv6
```

Configuration guidelines

Protocol-based VLAN configuration applies to hybrid ports only.

IP Subnet-based VLAN configuration

Introduction

In this approach, packets are assigned to VLANs based on their source IP addresses and subnet masks. A port configured with IP subnet-based VLANs assigns a received untagged packet to a VLAN based on the source address of the packet.

This feature is used to assign packets from the specified network segment or IP address to a specific VLAN.

Configuring an IP subnet-based VLAN

NOTE:

This feature is applicable only on hybrid ports.

Follow these steps to configure an IP subnet-based VLAN:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter VLAN view		**vlan** *vlan-id*	—
Associate an IP subnet with the current VLAN		**ip-subnet-vlan** [*ip-subnet-index*] **ip** *ip-address* [*mask*]	Required The IP network segment or IP address to be associated with a VLAN cannot be a multicast network segment or a multicast address.
Return to system view		**quit**	—
Enter interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command. • In Ethernet interface view, the subsequent configurations apply to the current port. • In port group view, the subsequent configurations apply to all ports in the port group. • In Layer 2 aggregate interface view, the subsequent configurations apply to the Layer 2 aggregate interface and all its member ports.
	Enter Layer 2 aggregate interface view	**interface bridge-aggregation** *interface-number*	
	Enter port group view	**port-group manual** *port-group-name*	
Configure port link type as hybrid		**port link-type hybrid**	Required
Configure the hybrid port(s) to permit the specified IP subnet-based VLANs to pass through		**port hybrid vlan** *vlan-id-list* { **tagged** | **untagged** }	Required
Associate the hybrid port(s) with the specified IP subnet-based VLAN		**port hybrid ip-subnet-vlan vlan** *vlan-id*	Required

NOTE:

After you configure a command on a Layer 2 aggregate interface, the system starts applying the configuration to the aggregate interface and its aggregation member ports. If the system fails to do that on the aggregate interface, it stops applying the configuration to the aggregation member ports. If it fails to do that on an aggregation member port, it skips the port and moves to the next port.

Displaying and maintaining VLAN

To do…	Use the command…	Remarks
Display VLAN information	**display vlan** [*vlan-id1* [**to** *vlan-id2*] | **all** | **dynamic** | **reserved** | **static**] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display VLAN interface information	**display interface vlan-interface** [*vlan-interface-id*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display hybrid ports or trunk ports on the device	**display port** { **hybrid** | **trunk** } [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display MAC address-to-VLAN entries	**display mac-vlan** { **all** | **dynamic** | **mac-address** *mac-address* | **static** | **vlan** *vlan-id* } [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display all interfaces with MAC-based VLAN enabled	**display mac-vlan interface** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display protocol information and protocol indexes of the specified VLANs	**display protocol-vlan vlan** { *vlan-id* [**to** *vlan-id*] | **all** } [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display protocol-based VLAN information on specified interfaces	**display protocol-vlan interface** { *interface-type interface-number* [**to** *interface-type interface-number*] | **all** } [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display IP subnet-based VLAN information and IP subnet indexes of specified VLANs	**display ip-subnet-vlan vlan** { *vlan-id* [**to** *vlan-id*] | **all** } [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the IP subnet-based VLAN information and IP subnet indexes of specified ports	**display ip-subnet-vlan interface** { *interface-list* | **all** } [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Clear statistics on a port	**reset counters interface vlan-interface** [*vlan-interface-id*]	Available in user view

Isolate-user-VLAN configuration

Overview

An isolate-user-VLAN uses a two-tier VLAN structure. In this approach, the following types of VLANs, isolate-user-VLAN and secondary VLAN, are configured on the same device.

The following are the characteristics of the isolate-user-VLAN implementation:

- Isolate-user-VLANs are mainly used for upstream data exchange. An isolate-user-VLAN can be associated with multiple secondary VLANs. Because the upstream device identifies only the isolate-user-VLAN and not the secondary VLANs, network configuration is simplified and VLAN resources are saved.
- You can isolate the Layer 2 traffic of different users by assigning the ports connected to them to different secondary VLANs.
- The dynamic MAC addresses entries learned in the isolate-user-VLAN are automatically synchronized to all the secondary VLANs, and the dynamic MAC address entries learned in a secondary VLAN are automatically synchronized to the isolate-user-VLAN.

As shown in Figure 38, the isolate-user-VLAN function is enabled on Switch B. VLAN 10 is the isolate-user-VLAN. VLANs 2, 5, and 8 are secondary VLANs associated with VLAN 10 and are invisible to Switch A.

Figure 38 An isolate-user-VLAN example

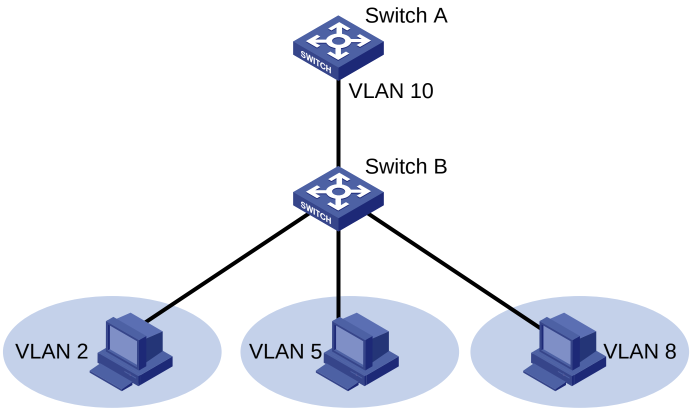

Configuring isolate-user-VLAN

Configure the isolate-user-VLAN through the following steps:

1. Configure the isolate-user-VLAN.
   - Assign non-trunk ports to the isolate-user-VLAN and configure these ports as upstream ports.
   - To enable users in the isolate-user-VLAN to communicate with other networks at Layer 3, configure a VLAN interface for the isolate-user-VLAN, and configure an IP address for the isolate-user-VLAN interface.
   - To enable Layer 3 communication among secondary VLANs associated with the same isolate-user-VLAN, you must enable local proxy ARP on the upstream device.

2. Configure the secondary VLANs.
   - Assign non-trunk ports to each secondary VLAN and configure these ports as downstream ports.
   - To enable users in the isolate-user-VLAN to communicate with other networks at Layer 3, configure VLAN interfaces for the secondary VLANs. Do not configure IP addresses for the secondary VLAN interfaces.
3. Associate the isolate-user-VLAN with the specified secondary VLANs.

Configuring an isolate-user-VLAN

Follow these steps to configure an isolate-user-VLAN:

To do...		Use the command	Remarks
Enter system view		**system-view**	—
Create a VLAN and enter VLAN view		**vlan** *vlan-id*	—
Configure the VLAN as an isolate-user-VLAN		**isolate-user-vlan enable**	Required
Return to system view		**quit**	—
Assign ports to the isolate-user-VLAN and ensure that at least one port takes the isolate-user-VLAN as its PVID	Access port	For how to assign an access port to a VLAN, see the chapter "VLAN configuration."	Required Use either approach.
	Hybrid port	For how to assign a hybrid port to a VLAN, see the chapter "VLAN configuration."	
Configure the isolate-user-VLAN type of the port as upstream		**port isolate-user-vlan promiscuous**	Optional By default, no isolate-user-VLAN type is specified for the port.
Return to system view		**quit**	—
Create the isolate-user-VLAN interface and enter the isolate-user-VLAN interface view		**interface vlan-interface** *vlan-interface-id*	• This configuration is required when users in the isolate-user-VLAN need to communicate with other networks at Layer 3. • This configuration is optional when users in the isolate-user-VLAN do not need to communicate with other networks at Layer 3. The *vlan-interface-id* argument must take the isolate-user-VLAN ID.

To do…	Use the command	Remarks
Configure an IP address for the isolate-user-VLAN interface	**ip address** *ip-address* { *mask* | *mask-length* } [**sub**]	• This configuration is required when users in the isolate-user-VLAN need to communicate with other networks at Layer 3. • This configuration is optional when users in the isolate-user-VLAN do not need to communicate with other networks at Layer 3. By default, the isolate-user-VLAN ID is not configured with any IP address.

Configuring secondary VLANs

Follow these steps to configure secondary VLANs:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Create secondary VLANs		**vlan** { *vlan-id1* [**to** *vlan-id2*] | **all** }	Required
Isolate ports in the same secondary VLAN at Layer 2		**isolated-vlan enable**	Optional By default, ports in the same secondary VLAN can communicate at Layer 2. This configuration takes effect only after you associate the secondary VLANs with an isolate-user-VLAN.
Return to system view		**quit**	—
Assign ports to each secondary VLAN and ensure that at least one port in a secondary VLAN takes the secondary VLAN as its PVID	Access port	For how to assign an access port to a VLAN, see the chapter "VLAN configuration."	Required Use either approach
	Hybrid port	For how to assign a hybrid port to a VLAN, see the chapter "VLAN configuration."	
Configure the isolate-user-VLAN type of the port as downstream		**port isolate-user-vlan host**	Optional By default, no isolate-user-VLAN type is specified for the port.
Return to system view		**quit**	—

To do...	Use the command...	Remarks
Create a secondary VLAN interface and enter secondary VLAN interface view	**interface vlan-interface** *vlan-interface-id*	• This configuration is required when users in the isolate-user-VLAN need to communicate with other networks at Layer 3. • This configuration is optional when users in the isolate-user-VLAN do not need to communicate with other networks at Layer 3. The *vlan-interface-id* argument must take the secondary VLAN ID.

Associating secondary VLANs with an isolate-user-VLAN

Follow these steps to associate secondary VLANs with an isolate-user-VLAN:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Associate the specified secondary VLANs with the specified isolate-user-VLAN	**isolate-user-vlan** *isolate-user-vlan-id* **secondary** *secondary-vlan-id* [**to** *secondary-vlan-id*]	Required

NOTE:

After associating an isolate-user-VLAN with the specified secondary VLANs, you cannot perform any of the following operations:

- Adding or removing an access port to or from an involved VLAN.
- Deleting an involved VLAN.
- Isolating ports in the same secondary VLAN at Layer 2.
- Changing the isolate-user-VLAN type of a port.

To perform the preceding configurations, cancel the association first.

Displaying and maintaining isolate-user-VLAN

To do...	Use the command...	Remarks
Display the mapping between an isolate-user-VLAN and its secondary VLANs and information about these VLANs	**display isolate-user-vlan** [*isolate-user-vlan-id*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

Isolate-user-VLAN configuration example

Network requirements

As shown in Figure 39,

- Connect Device A to downstream devices Device B and Device C.
- Configure VLAN 5 on Device B as an isolate-user-VLAN, assign the uplink port GigabitEthernet 1/0/5 to VLAN 5, and associate VLAN 5 with secondary VLANs VLAN 2 and VLAN 3. Assign GigabitEthernet 1/0/2 to VLAN 2 and GigabitEthernet 1/0/1 to VLAN 3.
- Configure VLAN 6 on Device C as an isolate-user-VLAN, assign the uplink port GigabitEthernet 1/0/5 to VLAN 6, and associate VLAN 6 with secondary VLANs VLAN 3 and VLAN 4. Assign GigabitEthernet 1/0/3 to VLAN 3 and GigabitEthernet 1/0/4 to VLAN 4.
- As far as Device A is concerned, Device B has only VLAN 5 and Device C has only VLAN 6.

Figure 39 Network diagram for isolate-user-VLAN configuration

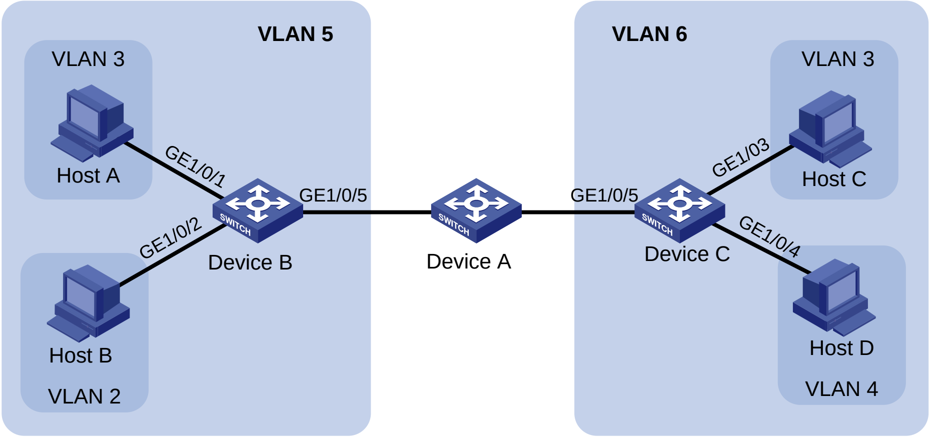

Configuration procedure

The following part provides only the configuration on Device B and Device C.

1. Configure Device B.

# Configure the isolate-user-VLAN.

```
<DeviceB> system-view
[DeviceB] vlan 5
[DeviceB-vlan5] isolate-user-vlan enable
[DeviceB-vlan5] quit
```

# Configure the secondary VLANs.

```
[DeviceB] vlan 2 to 3
```

# Configure the uplink port GigabitEthernet 1/0/5 to operate in promiscuous mode.

```
[DeviceB] interface gigabitethernet 1/0/5
[DeviceB-GigabitEthernet1/0/5] port isolate-user-vlan promiscuous
[DeviceB-GigabitEthernet1/0/5] quit
```

# Assign downlink ports GigabitEthernet 1/0/1 to VLAN 3 and GigabitEthernet 1/0/2 to VLAN 2, and configure the ports to operate in host mode.

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

```
[DeviceB-GigabitEthernet1/0/1] port access vlan 3
[DeviceB-GigabitEthernet1/0/1] port isolate-user-vlan host
[DeviceB-GigabitEthernet1/0/1] quit
[DeviceB] interface gigabitethernet 1/0/2
[DeviceB-GigabitEthernet1/0/2] port access vlan 2
[DeviceB-GigabitEthernet1/0/2] port isolate-user-vlan host
[DeviceB-GigabitEthernet1/0/2] quit
```

# Associate the isolate-user-VLAN with the secondary VLANs.

```
[DeviceB] isolate-user-vlan 5 secondary 2 to 3
```

2. Configure Device C.

# Configure the isolate-user-VLAN.

```
<DeviceC> system-view
[DeviceC] vlan 6
[DeviceC-vlan6] isolate-user-vlan enable
[DeviceC-vlan6] quit
```

# Configure the secondary VLANs.

```
[DeviceC] vlan 3 to 4
```

# Configure the uplink port GigabitEthernet 1/0/5 to operate in promiscuous mode.

```
[DeviceC] interface gigabitethernet 1/0/5
[DeviceC-GigabitEthernet1/0/5] port isolate-user-vlan promiscuous
[DeviceC-GigabitEthernet1/0/5] quit
```

# Configure downlink ports GigabitEthernet 1/0/3 to VLAN 3 and GigabitEthernet 1/0/4 to VLAN 4, and configure the ports to operate in host mode.

```
[DeviceC] interface gigabitethernet 1/0/3
[DeviceC-GigabitEthernet1/0/3] port access vlan 3
[DeviceC-GigabitEthernet1/0/3] port isolate-user-vlan host
[DeviceC-GigabitEthernet1/0/3] quit
[DeviceC] interface gigabitethernet 1/0/4
[DeviceC-GigabitEthernet1/0/4] port access vlan 4
[DeviceC-GigabitEthernet1/0/4] port isolate-user-vlan host
[DeviceC-GigabitEthernet1/0/4] quit
```

# Associate the isolate-user-VLAN with the secondary VLANs.

```
[DeviceC] isolate-user-vlan 6 secondary 3 to 4
```

Verification

# Display the isolate-user-VLAN configuration on Device B.

```
[DeviceB] display isolate-user-vlan
 Isolate-user-VLAN VLAN ID : 5
 Secondary VLAN ID : 2-3

 VLAN ID: 5
 VLAN Type: static
 Isolate-user-VLAN type : isolate-user-VLAN
 Route Interface: not configured
 Description: VLAN 0005
 Name: VLAN 0005
```

```
Tagged   Ports: none
Untagged Ports:
   GigabitEthernet1/0/1            GigabitEthernet1/0/2            GigabitEthernet1/0/5

VLAN ID: 2
VLAN Type: static
Isolate-user-VLAN type : secondary
Route Interface: not configured
Description: VLAN 0002
Name: VLAN 0002
Tagged   Ports: none
Untagged Ports:
   GigabitEthernet1/0/2            GigabitEthernet1/0/5

VLAN ID: 3
VLAN Type: static
Isolate-user-VLAN type : secondary
Route Interface: not configured
Description: VLAN 0003
Name: VLAN 0003
Tagged   Ports: none
Untagged Ports:
   GigabitEthernet1/0/1            GigabitEthernet1/0/5
```

Voice VLAN configuration

Overview

A voice VLAN is configured for voice traffic. After assigning the ports that connect to voice devices to a voice VLAN, the system automatically configures quality of service (QoS) parameters for voice traffic, to improve the transmission priority of voice traffic and ensure voice quality.

NOTE:

Common voice devices include IP phones and integrated access devices (IADs). Only IP phones are used in the voice VLAN configuration examples in this document.

OUI addresses

A device determines whether a received packet is a voice packet by evaluating its source MAC address. A packet whose source MAC address complies with the Organizationally Unique Identifier (OUI) address of the voice device is regarded as voice traffic.

You can configure the OUI addresses of a device in advance or use the default OUI addresses. Table 14 lists the default OUI address for each vendor's devices.

Table 14 The default OUI addresses of different vendors

Number	OUI address	Vendor
1	0001-e300-0000	Siemens phone
2	0003-6b00-0000	Cisco phone
3	0004-0d00-0000	Avaya phone
4	00d0-1e00-0000	Pingtel phone
5	0060-b900-0000	Philips/NEC phone
6	00e0-7500-0000	Polycom phone
7	00e0-bb00-0000	3Com phone

NOTE:

- In general, as the first 24 bits of a MAC address (in binary format), an OUI address is a globally unique identifier that IEEE assigns to a vendor. In this document, however, OUI addresses are addresses that the system uses to determine whether a received packet is a voice packet. They are the results of the AND operation of the arguments *mac-address* and *oui-mask* in the **voice vlan mac-address** command.
- You can remove the default OUI address of a device manually and then add new ones manually.

Voice VLAN assignment modes

A port can be assigned to a voice VLAN in one of the following modes:

- In automatic mode, the system matches the source MAC address carried in the untagged packets sent when an IP phone is powered on against the device's OUI addresses. If the system finds a match, it automatically assigns the receiving port to the voice VLAN, issues ACL rules, and configures the packet precedence. You can configure a voice VLAN aging time on the device. The system will remove a port from the voice VLAN if no packet is received from the port during the aging time. The system automatically assigns ports to, or removes ports from, a voice VLAN. Automatic mode is suitable for scenarios where PCs and IP phones connected in series access the network through the device and ports on the device transmit both voice traffic and data traffic at the same time, as shown in Figure 40. When the voice VLAN works normally, when the system reboots, the system reassigns ports in automatic voice VLAN assignment mode to the voice VLAN after the reboot, ensuring that existing voice connections can work normally. In this case, voice traffic streams do not trigger port assignment to the voice VLAN.

Figure 40 PCs and IP phones connected in series access the network

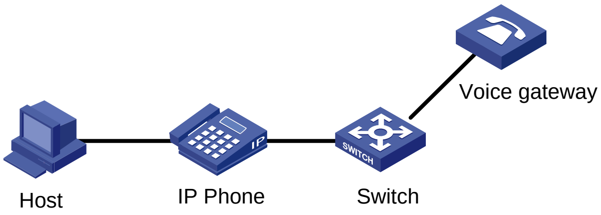

- In manual mode, you must manually assign an IP phone accessing port to a voice VLAN. Then, the system matches the source MAC addresses carried in the packets against the device's OUI addresses. If the system finds a match, it issues ACL rules and configures the packet precedence. In this mode, you must manually assign ports to, or remove ports from, a voice VLAN. Manual mode is suitable for scenarios where only IP phones access the network through the device, and ports on the device transmit only voice traffic, as shown in Figure 41. In this mode, ports assigned to a voice VLAN transmit voice traffic exclusively, which prevents the impact of data traffic on the transmission of voice traffic.

Figure 41 Only IP phones access the network

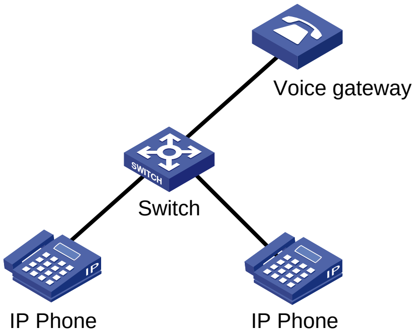

Both modes forward tagged packets according to their tags.

Table 15 and Table 16 list the configurations required for ports of different link types to support tagged or untagged voice traffic sent from IP phones when different voice VLAN assignment modes are configured.

- IP phones send tagged voice traffic

Table 15 Required configurations on ports of different link types in order for the ports to support tagged voice traffic

Port link type	Voice VLAN assignment mode	Support for tagged voice traffic	Configuration requirements
Access	Automatic	No	—
	Manual		
Trunk	Automatic	Yes	Configure the PVID of the port, which cannot be the voice VLAN, and assign the port to its PVID.
	Manual		Make all the configurations required for the automatic mode, and assign the port to the voice VLAN.
Hybrid	Automatic	Yes	Configure the PVID of the port, which cannot be the voice VLAN, and configure the port to permit packets of its PVID to pass through untagged.
	Manual		Make all the configurations required for the automatic mode, and configure the port to permit packets of the voice VLAN to pass through tagged.

- IP phones send untagged voice traffic

When IP phones send untagged voice traffic, you can only configure the voice traffic receiving ports on the device to operate in manual voice VLAN assignment mode.

Table 16 Required configurations on ports of different link types in order for the ports to support tagged voice traffic

Port link type	Voice VLAN assignment mode	Support for untagged voice traffic	Configuration requirements
Access	Automatic	No	—
	Manual	Yes	Configure the PVID of the port as the voice VLAN.
Trunk	Automatic	No	—
	Manual	Yes	Configure the PVID of the port as the voice VLAN and assign the port to the voice VLAN.
Hybrid	Automatic	No	—
	Manual	Yes	Configure the PVID of the port as the voice VLAN and configure the port to permit packets of the voice VLAN to pass through untagged.

△ CAUTION:

- If an IP phone sends tagged voice traffic and its accessing port is configured with 802.1X authentication and guest VLAN, assign different VLAN IDs for the voice VLAN, the PVID of the connecting port, and the 802.1X guest VLAN.
- If an IP phone sends untagged voice traffic, to implement the voice VLAN feature, you must configure the PVID of the IP phone's accessing port as the voice VLAN. As a result, you cannot implement 802.1X authentication.

NOTE:

- The PVID of a port is VLAN 1 by default. You can change the PVID and assign a port to certain VLANs by using commands. For more information, see the chapter "VLAN configuration."
- Use the **display interface** command to display the PVID of a port and the VLANs to which the port is assigned.

Security mode and normal mode of voice VLANs

Depending on their inbound packet filtering mechanisms, voice VLAN-enabled ports operate in the following modes.

- Normal mode: Voice VLAN-enabled ports receive packets that carry the voice VLAN tag, and forward packets in the voice VLAN without comparing their source MAC addresses against the OUI addresses configured for the device. If the PVID of the port is the voice VLAN and the port works in manual VLAN assignment mode, the port forwards all received untagged packets in the voice VLAN. In normal mode, voice VLANs are vulnerable to traffic attacks. Malicious users might send large quantities of forged voice packets to consume the voice VLAN bandwidth, affecting normal voice communication.
- Security mode: Only voice packets whose source MAC addresses match the recognizable OUI addresses can pass through the voice VLAN-enabled inbound port, while all other packets are dropped.

In a safe network, you can configure the voice VLANs to operate in normal mode, reducing the consumption of system resources due to source MAC addresses checking.

TIP:

HP does not recommend that you transmit both voice traffic and non-voice traffic in a voice VLAN. If you must transmit both voice traffic and nonvoice traffic, ensure that the voice VLAN security mode is disabled.

Table 17 How a voice VLAN-enabled port processes packets in security and normal mode

Voice VLAN mode	Packet type	Packet processing mode
Security mode	Untagged packets	If the source MAC address of a packet matches an OUI address configured for the device, it is forwarded in the voice VLAN. Otherwise, it is dropped.
	Packets that carry the voice VLAN tag	
	Packets that carry other tags	Forwarded or dropped depending on whether the port allows packets of these VLANs to pass through.
Normal mode	Untagged packets	The port does not determine the source MAC addresses of inbound packets. In this way, both voice traffic and non-

Voice VLAN mode	Packet type	Packet processing mode
	Packets that carry the voice VLAN tag	
	Packets that carry other tags	Forwarded or dropped depending on whether the port allows packets of these VLANs to pass through.

Configuring a voice VLAN

Configuration prerequisites

Before you configure a voice VLAN, complete the following tasks:

- Create a VLAN
- Configure QoS priority settings for voice VLAN traffic on an interface before enabling voice VLAN on the interface.

If the configuration order is reversed, your priority configuration will fail. For more information, see "Configuring QoS priority settings for voice traffic on an interface."

- Configure the voice VLAN assignment mode.

For more information, see "Configuring a port to operate in automatic voice VLAN assignment mode" and "Configuring a port to operate in manual voice VLAN assignment mode."

NOTE:

- A port can belong to only one voice VLAN at a time.
- You cannot enable voice VLAN on a port where Link Aggregation Control Protocol (LACP) is enabled.

Configuring QoS priority settings for voice traffic on an interface

In voice VLAN applications, you can improve the quality of voice traffic by configuring the appropriate QoS priority settings, including the Class of Service (CoS) and Differentiated Services Code Point (DSCP) values, for voice traffic. Voice traffic carries its own QoS priority settings. You can configure the device either to modify or not to modify the QoS priority settings that the incoming voice traffic carries.

Follow these steps to configure QoS priority settings for voice traffic:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Enter interface view	**interface** *interface-type interface-number*	—
Configure the interface to trust the QoS priority settings in incoming voice traffic, but not to modify the CoS and DSCP values marked for incoming traffic of the voice VLAN	**voice vlan qos trust**	Required Use either command By default, an interface modifies the CoS value and the DSCP value

To do…	Use the command…	Remarks
Configure the interface to modify the CoS and DSCP values marked for incoming traffic of the voice VLAN into specified values	**voice vlan qos** *cos-value dscp-value*	marked for voice VLAN traffic into 6 and 46 respectively. The **voice vlan qos** command and the **voice vlan qos trust** command can overwrite each other, whichever is configured last.

NOTE:

Configure the QoS priority settings for voice traffic on an interface before you enable voice VLAN on the interface. If the configuration order is reversed, your priority trust setting will fail.

Configuring a port to operate in automatic voice VLAN assignment mode

Follow these steps to set a port to operate in automatic voice VLAN assignment mode:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Set the voice VLAN aging time	**voice vlan aging** *minutes*	Optional 1440 minutes by default. The voice VLAN aging time configuration is only applicable on ports in automatic voice VLAN assignment mode.
Enable the voice VLAN security mode	**voice vlan security enable**	Optional Enabled by default.
Add a recognizable OUI address	**voice vlan mac-address** *oui* **mask** *oui-mask* [**description** *text*]	Optional By default, each voice VLAN has default OUI addresses configured. For the default OUI addresses of different vendors, see Table 14.
Enter Ethernet interface view	**interface** *interface-type interface-number*	—
Configure the port to operate in automatic voice VLAN assignment mode	**voice vlan mode auto**	Optional Automatic voice VLAN assignment mode is enabled by default. The voice VLAN assignment modes on different ports are independent of one another.
Enable voice VLAN on the port	**voice vlan** *vlan-id* **enable**	Required Disabled by default

NOTE:

A protocol-based VLAN on a hybrid port can process only untagged inbound packets, whereas the voice VLAN in automatic mode on a hybrid port can process only tagged voice traffic. Do not configure a VLAN as both a protocol-based VLAN and a voice VLAN. For more information, see the chapter "VLAN configuration."

Configuring a port to operate in manual voice VLAN assignment mode

Follow these steps to set a port to operate in manual voice VLAN assignment mode:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Enable the voice VLAN security mode	**voice vlan security enable**	Optional Enabled by default.
Add a recognizable OUI address	**voice vlan mac-address** *oui* **mask** *oui-mask* [**description** *text*]	Optional By default, each voice VLAN has default OUI addresses configured. For the default OUI addresses of different vendors, see Table 14.
Enter interface view	**interface** *interface-type interface-number*	—
Configure the port to operate in manual voice VLAN assignment mode	**undo voice vlan mode auto**	Required Disabled by default
Assign the port (access, trunk, or hybrid) in manual voice VLAN assignment mode to the voice VLAN	For how to assign a port to a VLAN, see the chapter "VLAN configuration."	Required After you assign an access port to the voice VLAN, the voice VLAN becomes the PVID of the port automatically.
Configure the voice VLAN as the PVID of the port (trunk or hybrid)	For how to assign a port to a VLAN, see the chapter "VLAN configuration."	Optional This operation is required for untagged inbound voice traffic and prohibited for tagged inbound voice traffic.
Enable voice VLAN on the port	**voice vlan** *vlan-id* **enable**	Required

NOTE:

- You can configure different voice VLANs on different ports at the same time. However, you can configure one port with only one voice VLAN, and this voice VLAN must be a static VLAN that already exists on the device.
- You cannot enable voice VLAN on a port where Link Aggregation Control Protocol (LACP) is enabled.
- To make voice VLAN take effect on a port that is enabled with voice VLAN and operates in manual voice VLAN assignment mode, you must assign the port to the voice VLAN manually.

Displaying and maintaining voice VLAN

To do...	Use the command...	Remarks
Display the voice VLAN state	**display voice vlan state** [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the OUI addresses that the system supports	**display voice vlan oui** [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

Voice VLAN configuration examples

Automatic voice VLAN mode configuration example

Network requirements

As shown in Figure 42,

- The MAC address of IP phone A is 0011-1100-0001. The phone connects to a downstream device named PC A whose MAC address is 0022-1100-0002 and to GigabitEthernet 1/0/1 on an upstream device named Device A.
- The MAC address of IP phone B is 0011-2200-0001. The phone connects to a downstream device named PC B whose MAC address is 0022-2200-0002 and to GigabitEthernet 1/0/2 on Device A.
- Device A uses voice VLAN 2 to transmit voice packets for IP phone A, and uses voice VLAN 3 to transmit voice packets for IP phone B.
- Configure GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to work in automatic voice VLAN assignment mode. In addition, if one of them has not received any voice packet in 30 minutes, the port is removed from the corresponding voice VLAN automatically.

Figure 42 Network diagram for automatic voice VLAN assignment mode configuration

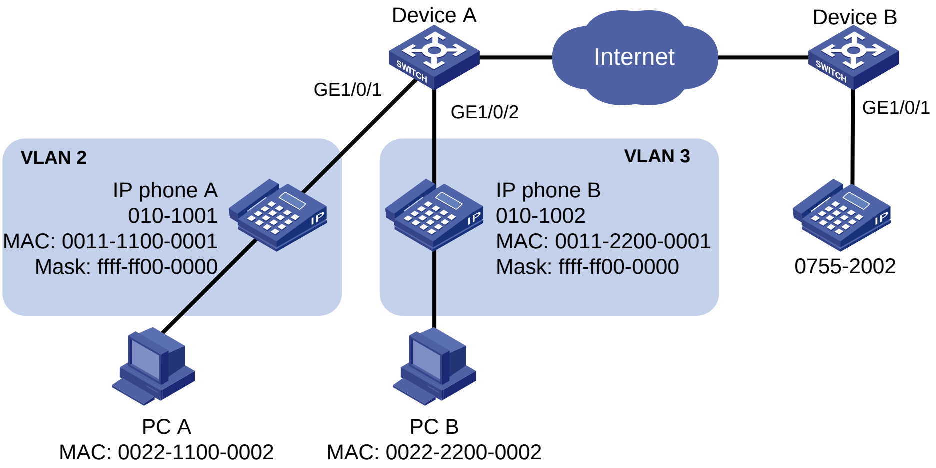

Configuration procedure

# Create VLAN 2 and VLAN 3.

```
<DeviceA> system-view
[DeviceA] vlan 2 to 3
```

# Set the voice VLAN aging time to 30 minutes.

```
[DeviceA] voice vlan aging 30
```

# GigabitEthernet 1/0/1 might receive both voice traffic and data traffic at the same time. To ensure the quality of voice packets and effective bandwidth use, configure voice VLANs to work in security mode to transmit only voice packets. By default, voice VLANs work in security mode. (Optional)

```
[DeviceA] voice vlan security enable
```

# Configure the allowed OUI addresses as MAC addresses prefixed by 0011-1100-0000 or 0011-2200-0000. Device A identifies packets whose MAC addresses match any of the configured OUI addresses as voice packets.

```
[DeviceA] voice vlan mac-address 0011-1100-0001 mask ffff-ff00-0000 description IP phone
A
[DeviceA] voice vlan mac-address 0011-2200-0001 mask ffff-ff00-0000 description IP phone
B
```

# Configure GigabitEthernet 1/0/1 as a hybrid port.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-type hybrid
```

# Configure GigabitEthernet 1/0/1 to operate in automatic voice VLAN assignment mode. By default, a port operates in automatic voice VLAN assignment mode. (Optional)

```
[DeviceA-GigabitEthernet1/0/1] voice vlan mode auto
```

# Configure VLAN 2 as the voice VLAN for GigabitEthernet 1/0/1.

```
[DeviceA-GigabitEthernet1/0/1] voice vlan 2 enable
[DeviceA-GigabitEthernet1/0/1] quit
```

# Configure GigabitEthernet 1/0/2.

```
[DeviceA] interface gigabitethernet 1/0/2
[DeviceA-GigabitEthernet1/0/2] port link-type hybrid
[DeviceA-GigabitEthernet1/0/2] voice vlan mode auto
[DeviceA-GigabitEthernet1/0/2] voice vlan 3 enable
```

Verification

# Display the OUI addresses, OUI address masks, and description strings.

```
<DeviceA> display voice vlan oui
Oui Address     Mask            Description
0001-e300-0000  ffff-ff00-0000  Siemens phone
0003-6b00-0000  ffff-ff00-0000  Cisco phone
0004-0d00-0000  ffff-ff00-0000  Avaya phone
0011-1100-0000  ffff-ff00-0000  IP phone A
0011-2200-0000  ffff-ff00-0000  IP phone B
00d0-1e00-0000  ffff-ff00-0000  Pingtel phone
0060-b900-0000  ffff-ff00-0000  Philips/NEC phone
00e0-7500-0000  ffff-ff00-0000  Polycom phone
00e0-bb00-0000  ffff-ff00-0000  3com phone
```

# Display the current states of voice VLANs.

```
<DeviceA> display voice vlan state
 Maximum of Voice VLANs: 8
 Current Voice VLANs: 2
 Voice VLAN security mode: Security
 Voice VLAN aging time: 30 minutes
 Voice VLAN enabled port and its mode:
 PORT                        VLAN        MODE        COS         DSCP
 --------------------------------------------------------------------
 GigabitEthernet1/0/1        2           AUTO        6           46
 GigabitEthernet1/0/2        3           AUTO        6           46
```

Manual voice VLAN assignment mode configuration example

Network requirements

As shown in Figure 43

- Create VLAN 2 and configure it as a voice VLAN that permits only voice traffic to pass through.
- The IP phones send untagged voice traffic. Configure GigabitEthernet 1/0/1 as a hybrid port.
- Configure GigabitEthernet 1/0/1 to operate in manual voice VLAN assignment mode. Configure GigabitEthernet 1/0/1 to allow voice traffic with an OUI address of 0011-2200-0000, a mask of ffff-ff00-0000, and a description string of test to be forwarded to the voice VLAN.

Figure 43 Network diagram for manual voice VLAN assignment mode configuration

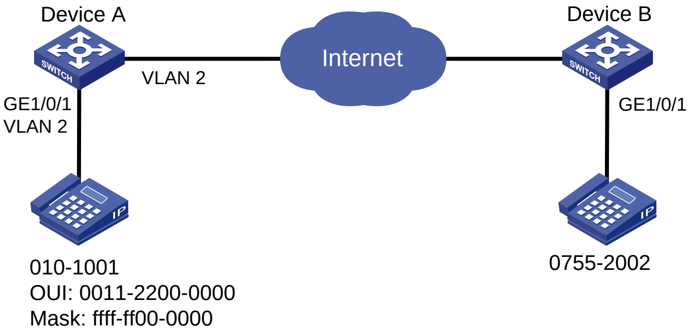

Configuration procedure

# Configure the voice VLAN to operate in security mode. (Optional. A voice VLAN operates in security mode by default.)

```
<DeviceA> system-view
[DeviceA] voice vlan security enable
```

# Add a recognizable OUI address 0011-2200-0000.

```
[DeviceA] voice vlan mac-address 0011-2200-0000 mask ffff-ff00-0000 description test
```

# Create VLAN 2.

```
[DeviceA] vlan 2
[DeviceA-vlan2] quit
```

# Configure GigabitEthernet 1/0/1 to operate in manual voice VLAN assignment mode.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] undo voice vlan mode auto
```

# Configure GigabitEthernet 1/0/1 as a hybrid port.

```
[DeviceA-GigabitEthernet1/0/1]port link-type hybrid
```

# Configure the voice VLAN (VLAN 2) as the PVID of GigabitEthernet 1/0/1 and configure GigabitEthernet 1/0/1 to permit the voice traffic of VLAN 2 to pass through untagged.

```
[DeviceA-GigabitEthernet1/0/1] port hybrid pvid vlan 2
```

```
[DeviceA-GigabitEthernet1/0/1] port hybrid vlan 2 untagged
```

# Enable voice VLAN on GigabitEthernet 1/0/1.

```
[DeviceA-GigabitEthernet1/0/1] voice vlan 2 enable
```

Verification

# Display the OUI addresses, OUI address masks, and description strings.

```
<DeviceA> display voice vlan oui
Oui Address     Mask            Description
0001-e300-0000  ffff-ff00-0000  Siemens phone
0003-6b00-0000  ffff-ff00-0000  Cisco phone
0004-0d00-0000  ffff-ff00-0000  Avaya phone
0011-2200-0000  ffff-ff00-0000  test
00d0-1e00-0000  ffff-ff00-0000  Pingtel phone
0060-b900-0000  ffff-ff00-0000  Philips/NEC phone
00e0-7500-0000  ffff-ff00-0000  Polycom phone
00e0-bb00-0000  ffff-ff00-0000  3com phone
```

# Display the current voice VLAN state.

```
<DeviceA> display voice vlan state
 Maximum of Voice VLANs: 8
 Current Voice VLANs: 1
 Voice VLAN security mode: Security
 Voice VLAN aging time: 1440 minutes
 Voice VLAN enabled port and its mode:
 PORT                        VLAN        MODE        COS         DSCP
 --------------------------------------------------------------------
 GigabitEthernet1/0/1        2           MANUAL      6           46
```

GVRP configuration

The Generic Attribute Registration Protocol (GARP) provides a generic framework for devices in a bridged LAN, such as end stations and switches, to register and deregister attribute values. The GARP VLAN Registration Protocol (GVRP) is a GARP application that registers and deregisters VLAN attributes. GVRP uses the operating mechanism of GARP to maintain and propagate dynamic VLAN registration information for GVRP devices on the network.

Introduction to GVRP

GARP

GARP provides a mechanism that allows participants in a GARP application to distribute, propagate, and register—with other participants in a LAN—the attributes specific to the GARP application, such as the VLAN or multicast address attributes.

How GARP works

Each port that participates in a GARP application (GVRP, for example) is a GARP participant.

Through the GARP mechanism, the attribute information of GARP participants is rapidly propagated across the entire LAN. A GARP participant registers and deregisters its attribute information with other GARP participants by sending and withdrawing declarations, and registers and deregisters the attribute information of other participants according to the declarations and withdrawals that it receives.

Figure 44 How GARP works

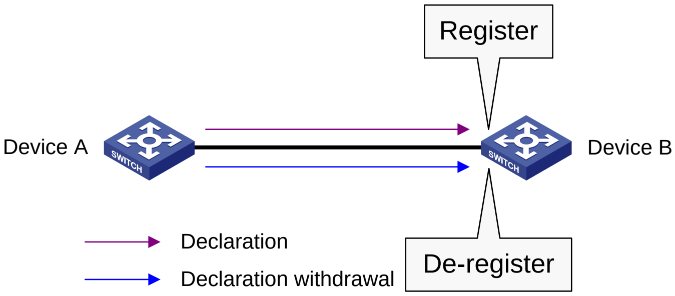

For example, GVRP registers and deregisters VLAN attributes in the following cases.

- When a port receives a declaration for a VLAN attribute, it registers the VLAN attribute carried in the declaration and joins the VLAN.
- When a port receives a withdrawal for a VLAN attribute, it deregisters the VLAN attribute carried in the withdrawal and leaves the VLAN.

GARP messages

A GARP participant exchanges information with other GARP participants by sending GARP messages, which are Join, Leave, and LeaveAll. These messages work together to ensure the registration and deregistration of attribute information. As a GARP application, GVRP also uses GARP messages for information exchange.

1. Join messages

A GARP participant sends Join messages when it must register its attributes (including manually configured attributes) with other participants, and when it receives Join messages from other participants. The types of Join messages are JoinEmpty and JoinIn.

- A GARP participant sends a JoinEmpty message to declare an attribute not registered on it.
- A GARP participant sends a JoinIn message to declare an attribute registered on it.

2. Leave messages

A GARP participant sends Leave messages to have its attributes deregistered on other participants. It also sends Leave messages when it deregisters attributes after it receives Leave messages from other GARP participants, and when attributes are manually deregistered on it. The types of Leave messages are LeaveEmpty and LeaveIn.

- A GARP participant sends a LeaveEmpty message to deregister an attribute not registered on it.
- A GARP participant sends a LeaveIn message to deregister an attribute registered on it.

3. LeaveAll messages

Each GARP participant starts a LeaveAll timer upon startup. When the LeaveAll timer expires, a GARP participant sends LeaveAll messages to deregister all attributes so that all attributes can be re-registered on the other GARP participants.

GARP timers

GARP defines timers to control the sending of GARP messages.

NOTE:

- The settings of GARP timers apply to all GARP applications, such as GVRP, on a LAN.
- On a GARP-enabled network, each port of a switch maintains its own Hold, Join, and Leave timers, but only one LeaveAll timer is maintained on each switch globally.
- The value ranges for the Hold, Join, Leave, and LeaveAll timers are dependent on one another. For more information, see Table 19.

1. Hold timer

The Hold timer sets the delay that a GARP participant waits before sending a Join or Leave message.

When an attribute value changes or a Join or Leave message arrives, the GARP participant does not sends the message immediately. Rather, it assembles Join and Leave messages in the least number of GARP PDUs, and sends them out when the Hold timer expires. This timer reduces the number of GARP PDUs and saves bandwidth.

2. Join timer

A GARP participant might declare an attribute twice to ensure reliable transmission. The Join timer sets the interval between the two declarations.

A GARP participant starts a Join timer when it declares an attribute value or receives a JoinIn message for the attribute value. If the GARP participant does not receive any declaration for the attribute value when the Join timer expires, it re-declares the attribute value.

NOTE:

All attributes of a GARP participant share the same Join timer. Set the Join timer long enough so that all attributes can be sent in one declaration.

3. Leave timer

A GARP participant starts a Leave timer when it receives a Leave message for an attribute value. If the GARP participant receives no Join message for the attribute value before the timer expires, it deregisters the attribute value.

4. LeaveAll timer

When a GARP application is enabled, a LeaveAll timer starts. The GARP participant sends a LeaveAll message when the timer expires. Then, the LeaveAll timer restarts to begin a new cycle. The LeaveAll timer and all other GARP timers also restart when the GARP participant receives a LeaveAll message.

NOTE:

- Because a LeaveAll message deregisters all attributes in the entire network, do not set the LeaveAll timer too short. At a minimum, the LeaveAll timer must be longer than the Leave timer.
- On a GARP-enabled network, a switch can send LeaveAll messages at the interval set by its LeaveAll timer or the LeaveAll timer of another device on the network, whichever is smaller. This is because each time a switch on the network receives a LeaveAll message, it resets its LeaveAll timer.

GARP message format

Figure 45 GARP message format

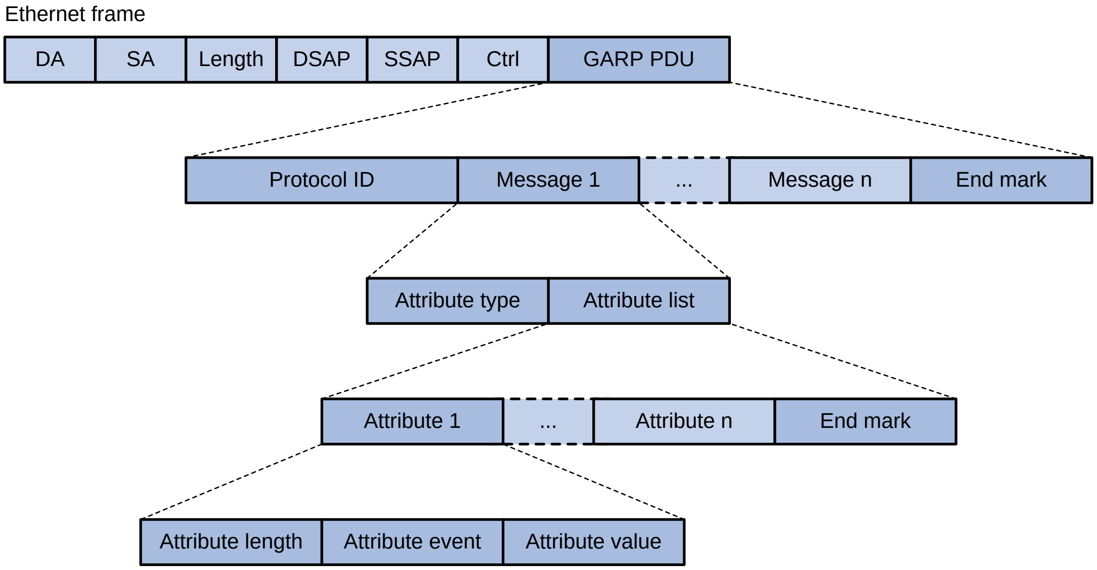

As shown in Figure 45, GARP messages use the IEEE 802.3 Ethernet frame format.

Table 18 Descriptions of the GARP message fields

Field	Description	Value
GARP PDU	GARP Protocol Data Unit	—
Protocol ID	Protocol identifier for GARP PDU	0x0001
Message	One or multiple messages, each of which contains an attribute type and an attribute list	—
End mark	Indicates the end of a GARP PDU	0x00

Field	Description	Value
Attribute type	Defined by the GARP application	0x01 for GVRP, which indicates the VLAN ID attribute
Attribute list	Contains one or multiple attributes	—
Attribute	Consists of an attribute length, an attribute event, and an attribute value	—
Attribute length	Length of an attribute, inclusive of the attribute length field	2 to 255 (in bytes)
Attribute event	Event that the attribute describes	• 0x00: LeaveAll event • 0x01: JoinEmpty event • 0x02: JoinIn event • 0x03: LeaveEmpty event • 0x04: LeaveIn event • 0x05: Empty event
Attribute value	Attribute value	VLAN ID for GVRP If the value of the attribute event field is 0x00 (LeaveAll event), the attribute value field is invalid.

The destination MAC addresses of GARP messages are multicast MAC addresses, and vary with GARP applications. For example, the destination MAC address of GVRP is 01-80-C2-00-00-21. A switch distributes GARP messages to different GARP applications according to the destination MAC addresses carried in GARP messages.

GVRP

GVRP overview

As a GARP application, GVRP enables a switch to propagate local VLAN registration information to other participant devices, and to dynamically update its local database with the VLAN registration information from other devices, including active VLAN members and the ports through which they can be reached. This ensures that all GVRP participants on a bridged LAN maintain the same VLAN registration information. The VLAN registration information propagated by GVRP includes both manually configured local static entries and dynamic entries from other devices.

GVRP registration modes

Manually created VLANs are called “static VLANs,” and VLANs created by GVRP are called “dynamic VLANs.” GVRP provides the following registration modes on a port: normal, fixed, and forbidden. The different registration modes determine how a port handles static and dynamic VLANs.

- Normal: Allows dynamic creation, registration, and deregistration of VLANs on the trunk port.
- Fixed: Allows manual creation and registration of VLANs, prevents VLAN deregistration, and registers all known VLANs on other ports on the trunk port.
- Forbidden: Deregisters all VLANs (except VLAN 1) and prevents any further VLAN creation or registration on the trunk port.

Protocols and standards

- IEEE 802.1Q, *Virtual Bridged Local Area Networks*

GVRP configuration task list

Complete these tasks to configure GVRP:

Task	Remarks
Configuring GVRP functions	Required
Configuring GARP timers	Optional

NOTE:

- GVRP configuration made in Ethernet interface view or Layer 2 aggregate interface view takes effect on the current interface only. GVRP configuration made in port group view takes effect on all the member ports in the group.
- GVRP configuration made on a member port in an aggregation group takes effect only after the port is removed from the aggregation group.

Configuring GVRP functions

Before enabling GVRP on a port, you must enable GVRP globally. In addition, you can configure GVRP only on trunk ports, and you must assign the involved trunk ports to all dynamic VLANs.

Follow these steps to configure GVRP functions on a trunk port:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enable GVRP globally		**gvrp**	Required Globally disabled by default
Enter Ethernet interface view, Layer 2 aggregate interface view, or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Perform either of the commands.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the link type of the ports as trunk		**port link-type trunk**	Required Access by default.
Assign the trunk ports to all VLANs		**port trunk permit vlan all**	Required By default, a trunk port is assigned to VLAN 1 only.

To do…	Use the command…	Remarks
Enable GVRP on the ports	**gvrp**	Required Disabled by default.
Configure the GVRP registration mode on the ports	**gvrp registration** { **fixed** | **forbidden** | **normal** }	Optional **normal** by default.

NOTE:

- For more information about the **port link-type trunk** and **port trunk permit vlan all** commands, see the chapter "VLAN configuration commands."
- GVRP is mutually exclusive with service loopback.
- In an MSTP network, GVRP can run on only the CIST. Blocked ports on the CIST cannot receive or send GVRP packets.
- Do not enable both GVRP and remote port mirroring. Otherwise, GVRP might register the remote probe VLAN to unexpected ports, which would cause the monitor port to receive undesired duplicates. For more information about port mirroring, see the *Network Management and Monitoring Configuration Guide*.
- Enabling GVRP on a Layer 2 aggregate interface enables both the aggregate interface and all selected member ports in the corresponding link aggregation group to participate in dynamic VLAN registration and deregistration.

Configuring GARP timers

The LeaveAll timer is configured in system view and takes effect on all ports, but the other GARP timers are configured on a port basis.

Follow these steps to configure GARP timers:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Configure the GARP LeaveAll timer		**garp timer leaveall** *timer-value*	Optional The default is 1000 centiseconds.
Enter Ethernet interface view, Layer 2 aggregate interface view, or port group view	Enter Ethernet or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Perform either of the commands. Depending on the view that you accessed, the subsequent configuration takes effect on a port or all ports in a port group.
	Enter port group view	**port-group manual** *port-group-name*	
Configure the Hold timer		**garp timer hold** *timer-value*	Optional 10 centiseconds by default
Configure the Join timer		**garp timer join** *timer-value*	Optional 20 centiseconds by default

To do…	Use the command…	Remarks
Configure the Leave timer	**garp timer leave** *timer-value*	Optional 60 centiseconds by default

As shown in Table 19, the value ranges for GARP timers are dependent on one another; use the following guidelines to configure GARP timers:

- If you want to set a value beyond the value range for a timer, you can change the value range by tuning the value of another related timer.
- If you want to restore the default settings of the timers, restore the Hold timer first, followed by the Join, Leave, and LeaveAll timers.

Table 19 Dependencies of GARP timers

Timer	Lower limit	Upper limit
Hold	10 centiseconds	No greater than half of the Join timer setting
Join	No less than two times the Hold timer setting	Less than half of the leave timer setting
Leave	Greater than two times the Join timer setting	Less than the LeaveAll timer setting
LeaveAll	Greater than the Leave timer setting	32,765 centiseconds

Displaying and maintaining GVRP

To do…	Use the command…	Remarks
Display statistics about GARP on ports	**display garp statistics** [**interface** *interface-list*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display GARP timers on ports	**display garp timer** [**interface** *interface-list*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the local VLAN information that GVRP maintains on ports	**display gvrp local-vlan interface** *interface-type interface-number* [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the current GVRP state in the specified VLANs on ports	**display gvrp state interface** *interface-type interface-number* **vlan** *vlan-id* [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display GVRP statistics on ports	**display gvrp statistics** [**interface** *interface-list*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the global GVRP state	**display gvrp status** [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the information about dynamic VLAN operations on ports	**display gvrp vlan-operation interface** *interface-type interface-number* [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

To do…	Use the command…	Remarks
Clear the GARP statistics on ports	**reset garp statistics** [**interface** *interface-list*]	Available in user view

GVRP configuration examples

GVRP normal registration mode configuration example

Network requirements

As shown in Figure 46:

- Device A and Device B are connected through their GigabitEthernet 1/0/1 ports.
- Enable GVRP and configure the normal registration mode on ports to enable the registration and deregistration of dynamic and static VLAN information between the two devices.

Figure 46 Network diagram for GVRP normal registration mode configuration

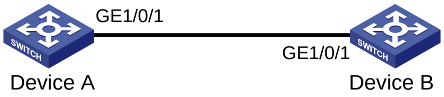

Configuration procedure

1. Configure Device A.

# Enable GVRP globally.

```
<DeviceA> system-view
[DeviceA] gvrp
```

# Configure port GigabitEthernet 1/0/1 as a trunk port, and assign it to all VLANs.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-type trunk
[DeviceA-GigabitEthernet1/0/1] port trunk permit vlan all
```

# Enable GVRP on trunk port GigabitEthernet 1/0/1.

```
[DeviceA-GigabitEthernet1/0/1] gvrp
[DeviceA-GigabitEthernet1/0/1] quit
```

# Create VLAN 2 (a static VLAN).

```
[DeviceA] vlan 2
[DeviceA-vlan2] quit
```

2. Configure Device B.

# Enable GVRP globally.

```
<DeviceB> system-view
[DeviceB] gvrp
```

# Configure port GigabitEthernet 1/0/1 as a trunk port, and assign it to all VLANs.

```
[DeviceB] interface gigabitethernet 1/0/1
[DeviceB-GigabitEthernet1/0/1] port link-type trunk
[DeviceB-GigabitEthernet1/0/1] port trunk permit vlan all
```

# Enable GVRP on trunk port GigabitEthernet 1/0/1.

```
[DeviceB-GigabitEthernet1/0/1] gvrp
[DeviceB-GigabitEthernet1/0/1] quit
```

# Create VLAN 3 (a static VLAN).

```
[DeviceB] vlan 3
[DeviceB-vlan3] quit
```

3. Verify the configuration.

Use the **display gvrp local-vlan** command to display the local VLAN information that GVRP maintains on ports. For example:

# Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device A.

```
[DeviceA] display gvrp local-vlan interface gigabitethernet 1/0/1
 Following VLANs exist in GVRP local database:
  1(default),2-3
```

According to the output, information about VLAN 1, static VLAN information of VLAN 2 on the local device, and dynamic VLAN information of VLAN 3 on Device B are all registered through GVRP.

# Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device B.

```
[DeviceB] display gvrp local-vlan interface gigabitethernet 1/0/1
 Following VLANs exist in GVRP local database:
  1(default),2-3
```

According to the output, information about VLAN 1, static VLAN information of VLAN 3 on the local device, and dynamic VLAN information of VLAN 2 on Device A are all registered through GVRP.

GVRP fixed registration mode configuration example

Network requirements

As shown in Figure 47:

- Device A and Device B are connected through their GigabitEthernet 1/0/1 ports.
- Enable GVRP and configure the fixed registration mode on ports to enable the registration and deregistration of static VLAN information between the two devices.

Figure 47 Network diagram for GVRP fixed registration mode configuration

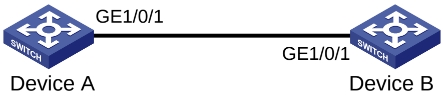

Configuration procedure

1. Configure Device A.

# Enable GVRP globally.

```
<DeviceA> system-view
[DeviceA] gvrp
```

# Configure port GigabitEthernet 1/0/1 as a trunk port, and assign it to all VLANs.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-type trunk
[DeviceA-GigabitEthernet1/0/1] port trunk permit vlan all
```

# Enable GVRP on GigabitEthernet 1/0/1 and set the GVRP registration mode to fixed on the port.

```
[DeviceA-GigabitEthernet1/0/1] gvrp
[DeviceA-GigabitEthernet1/0/1] gvrp registration fixed
[DeviceA-GigabitEthernet1/0/1] quit
```

# Create VLAN 2 (a static VLAN).

```
[DeviceA] vlan 2
[DeviceA-vlan2] quit
```

2. Configure Device B.

# Enable GVRP globally.

```
<DeviceB> system-view
[DeviceB] gvrp
```

# Configure port GigabitEthernet 1/0/1 as a trunk port, and assign it to all VLANs.

```
[DeviceB] interface gigabitethernet 1/0/1
[DeviceB-GigabitEthernet1/0/1] port link-type trunk
[DeviceB-GigabitEthernet1/0/1] port trunk permit vlan all
```

# Enable GVRP on GigabitEthernet 1/0/1, and set the GVRP registration mode to fixed on the port.

```
[DeviceB-GigabitEthernet1/0/1] gvrp
[DeviceB-GigabitEthernet1/0/1] gvrp registration fixed
[DeviceB-GigabitEthernet1/0/1] quit
```

# Create VLAN 3 (a static VLAN).

```
[DeviceB] vlan 3
[DeviceB-vlan3] quit
```

3. Verify the configuration.

Use the **display gvrp local-vlan** command to display the local VLAN information that GVRP maintains on ports. For example:

# Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device A.

```
[DeviceA] display gvrp local-vlan interface gigabitethernet 1/0/1
 Following VLANs exist in GVRP local database:
  1(default), 2
```

According to the output, information about VLAN 1 and static VLAN information of VLAN 2 on the local device are registered through GVRP, but dynamic VLAN information of VLAN 3 on Device B is not.

# Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device B.

```
[DeviceB] display gvrp local-vlan interface gigabitethernet 1/0/1
 Following VLANs exist in GVRP local database:
  1(default), 3
```

According to the output, information about VLAN 1 and static VLAN information of VLAN 3 on the local device are registered through GVRP, but dynamic VLAN information of VLAN 2 on Device A is not.

GVRP forbidden registration mode configuration example

Network requirements

As shown in Figure 48:

- Device A and Device B are connected through their GigabitEthernet 1/0/1 ports.

- Enable GVRP and configure the forbidden registration mode on ports to prevent the registration and deregistration of all VLANs but VLAN 1 between the two devices.

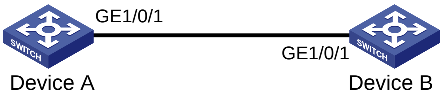

Figure 48 Network diagram for GVRP forbidden registration mode configuration

Configuration procedure

1. Configure Device A.

# Enable GVRP globally.

```
<DeviceA> system-view
[DeviceA] gvrp
```

# Configure port GigabitEthernet 1/0/1 as a trunk port, and assign it to all VLANs.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-type trunk
[DeviceA-GigabitEthernet1/0/1] port trunk permit vlan all
```

# Enable GVRP on GigabitEthernet 1/0/1, and set the GVRP registration mode to forbidden on the port.

```
[DeviceA-GigabitEthernet1/0/1] gvrp
[DeviceA-GigabitEthernet1/0/1] gvrp registration forbidden
[DeviceA-GigabitEthernet1/0/1] quit
```

# Create VLAN 2 (a static VLAN).

```
[DeviceA] vlan 2
[DeviceA-vlan2] quit
```

2. Configure Device B.

# Enable GVRP globally.

```
<DeviceB> system-view
[DeviceB] gvrp
```

# Configure port GigabitEthernet 1/0/1 as a trunk port, and assign it to all VLANs.

```
[DeviceB] interface gigabitethernet 1/0/1
[DeviceB-GigabitEthernet1/0/1] port link-type trunk
[DeviceB-GigabitEthernet1/0/1] port trunk permit vlan all
```

# Enable GVRP on GigabitEthernet 1/0/1, and set the GVRP registration mode to forbidden on the port.

```
[DeviceB-GigabitEthernet1/0/1] gvrp
[DeviceB-GigabitEthernet1/0/1] gvrp registration forbidden
[DeviceB-GigabitEthernet1/0/1] quit
```

# Create VLAN 3 (a static VLAN).

```
[DeviceB] vlan 3
[DeviceB-vlan3] quit
```

3. Verify the configuration.

Use the **display gvrp local-vlan** command to display the local VLAN information that GVRP maintains on ports. For example:

# Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device A.

```
[DeviceA] display gvrp local-vlan interface gigabitethernet 1/0/1
```

```
 Following VLANs exist in GVRP local database:
  1(default)
```

According to the output, information about VLAN 1 is registered through GVRP, but static VLAN information of VLAN 2 on the local device and dynamic VLAN information of VLAN 3 on Device B are not.

# Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device B.

```
[DeviceB] display gvrp local-vlan interface gigabitethernet 1/0/1
 Following VLANs exist in GVRP local database:
  1(default)
```

According to the output, information about VLAN 1 is registered through GVRP, but static VLAN information of VLAN 3 on the local device and dynamic VLAN information of VLAN 2 on Device A are not.

QinQ configuration

NOTE:

Throughout this document, customer network VLANs (CVLANs), also called "*inner VLANs*", refer to the VLANs that a customer uses on the private network; service provider network VLANs (SVLANs), also called "*outer VLANs*", refer to the VLANs that a service provider uses to carry VLAN tagged traffic for customers.

Introduction to QinQ

802.1Q-in-802.1Q (QinQ) is a flexible, easy-to-implement Layer 2 VPN technology based on IEEE 802.1Q. QinQ enables the edge switch on a service provider network to insert an outer VLAN tag in the Ethernet frames from customer networks, so that the Ethernet frames travel across the service provider network (public network) with double VLAN tags. QinQ enables a service provider to use a single SVLAN to serve customers who have multiple CVLANs.

Background and benefits

The IEEE 802.1Q VLAN tag uses 12 bits for VLAN IDs. A switch supports a maximum of 4094 VLANs. This is far from enough for isolating users in actual networks, especially in metropolitan area networks (MANs).

By tagging tagged frames, QinQ expands the available VLAN space from 4094 to 4094 × 4094. QinQ delivers the following benefits:

- Releases the stress on the SVLAN resource.
- Enables customers to plan their CVLANs without conflicting with SVLANs.
- Provides an easy-to-implement Layer 2 VPN solution for small-sized MANs or intranets.
- Enables the customers to keep their VLAN assignment schemes unchanged when the service provider upgrades the service provider network.

How QinQ works

The switches in the public network forward a frame only according to its outer VLAN tag and obtain its source MAC address into the MAC address table of the outer VLAN. The inner VLAN tag of the frame is transmitted as the payload.

Figure 49 Typical QinQ application scenario

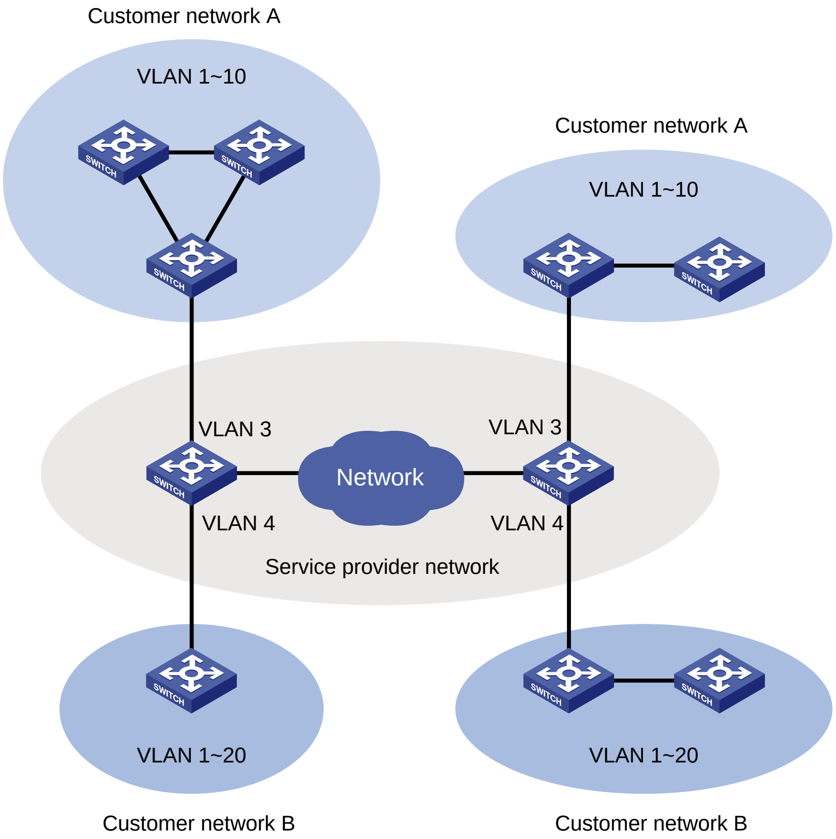

As shown in Figure 49, customer network A has CVLANs 1 through 10, and customer network B has CVLANs 1 through 20. The service provider assigns SVLAN 3 for customer network A, and assigns SVLAN 4 for customer network B. When a tagged Ethernet frame from customer network A arrives at the edge of the service provider network, the edge switch tags the frame with outer VLAN 3. When a tagged Ethernet frame from customer network B arrives at the edge of the service provider network, the edge switch tags it with outer VLAN 4. As a result, no overlap of VLAN IDs among customers exists, and traffic from different customers can be identified separately.

QinQ frame structure

A QinQ frame is transmitted double-tagged over the service provider network. The inner VLAN tag is the CVLAN tag, and the outer one is the SVLAN tag that the service provider has allocated to the customer.

Figure 50 Single-tagged Ethernet frame header versus double-tagged Ethernet frame header

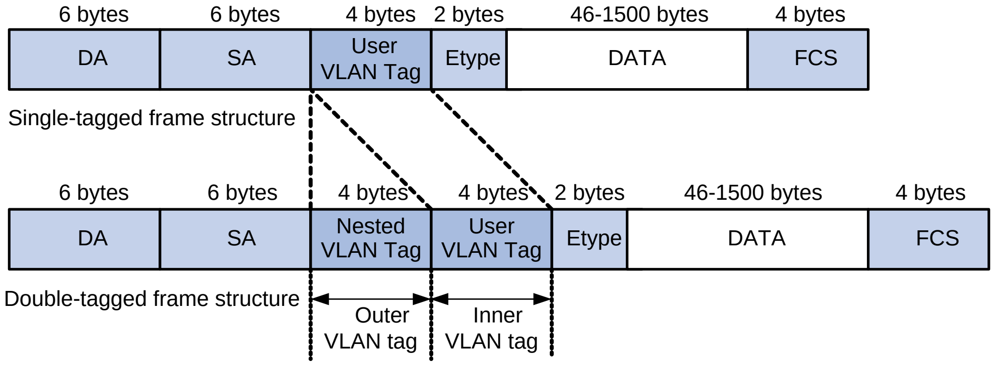

NOTE:

Ensure that all ports on the path of a QinQ packet allow 1508-byte or larger frames to pass through. The minimum size of a QinQ packet is 1508 bytes, which comprises two four-byte VLAN tags and one 1500-byte standard Ethernet frame.

Implementations of QinQ

HP provides the following QinQ implementations: basic QinQ and selective QinQ.

1. Basic QinQ

Basic QinQ enables a port to tag any incoming frames with its PVID, regardless of whether they have been tagged or not. If an incoming frame has been tagged, it becomes a double-tagged frame. If not, it becomes a frame tagged with the port's PVID.

2. Selective QinQ

Selective QinQ is more flexible than basic QinQ. In addition to all the functions of basic QinQ, selective QinQ enables a port to perform the following per-CVLAN actions for incoming frames:

- Tagging frames from different CVLANs with different SVLAN tags.
- Marking the outer VLAN 802.1p priority based on the existing inner VLAN 802.1p priority.

Besides being able to separate the service provider network from the customer networks, selective QinQ provides abundant service features and enables more flexible networking.

Modifying the TPID in a VLAN tag

A VLAN tag uses the tag protocol identifier (TPID) field to identify the protocol type of the tag. The value of this field, as defined in IEEE 802.1Q, is 0x8100.

Figure 51 VLAN tag structure of an Ethernet frame

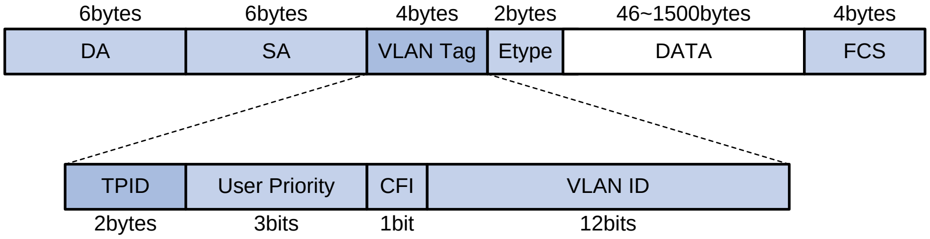

The switch determines whether a received frame carries a SVLAN or CVLAN tag by checking the TPID value. For example, if a frame carries a SVLAN tag with TPID value 0x9100 and a CVLAN tag with TPID value 0x8100, and the configured TPID value of the SVLAN tag is 0x9100 and that of the CVLAN tag is 0x8200, the switch considers that the frame carries only the SVLAN tag but not the CVLAN tag.

The systems of different vendors might set the TPID of the outer VLAN tag of QinQ frames to different values. For compatibility with these systems, modify the TPID value so that the QinQ frames, when sent to the public network, carry the TPID value identical to the value of a particular vendor, allowing interoperability with the switches of that vendor.

The TPID in an Ethernet frame has the same position as the protocol type field in a frame without a VLAN tag. To avoid problems in packet forwarding and handling in the network, do not set the TPID value to any of the reserved values.

Table 20 Reserved protocol type values

Protocol type	Value
ARP	0x0806
PUP	0x0200
RARP	0x8035
IP	0x0800
IPv6	0x86DD
PPPoE	0x8863/0x8864
MPLS	0x8847/0x8848
IPX/SPX	0x8137
IS-IS	0x8000
LACP	0x8809
802.1X	0x888E
Cluster	0x88A7
Reserved	0xFFFD/0xFFFE/0xFFFF

Protocols and standards

IEEE 802.1Q: *IEEE standard for local and metropolitan area networks: Virtual Bridged Local Area Networks*

QinQ configuration task list

Complete the follows tasks to configure QinQ:

Task		Remarks
Configuring basic QinQ	Enabling basic QinQ	Required
	Configuring VLAN transparent transmission	Optional
Configuring selective	Configuring an outer VLAN tagging policy	Optional

Task		Remarks
QinQ	Configuring an inner-outer VLAN 802.1p priority mapping	Optional
Configuring the TPID value in VLAN tags		Optional

NOTE:

- QinQ requires configurations only on the service provider network.
- QinQ configurations made in Ethernet interface view take effect on the current interface only. Those made in Layer 2 aggregate interface view take effect on the current aggregate interface and all the member ports in the aggregation group. Those made in port group view take effect on all member ports in the current port group.
- Configure both basic and selective QinQ on the ports that connect customer networks.

Configuring basic QinQ

Enabling basic QinQ

A basic QinQ-enabled port tags an incoming packet with its PVID.

Follow these steps to enable basic QinQ:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Enable basic QinQ		**qinq enable**	Required Disabled by default.

Configuring VLAN transparent transmission

When basic QinQ is enabled on a port, all packets passing through the port are tagged with the port's PVID. However, by configuring the VLAN transparent transmission function on a port, you can specify the port not to add its PVID to packets carrying specific inner VLAN tags when they pass through it, so that these packets are transmitted in the service provider network with single tags.

Follow these steps to configure VLAN transparent transmission:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command.

To do...		Use the command...	Remarks
	Enter port group view	**port-group manual** *port-group-name*	
Configure VLAN transparent transmission on the ports		**qinq transparent-vlan** *vlan-list*	Required By default, VLAN transparent transmission is not configured.

NOTE:

When you are configuring transparent transmission for a VLAN, you must configure all the switches on the transmission path to permit packets of this VLAN to pass through.

Configuring selective QinQ

Configuring an outer VLAN tagging policy

Basic QinQ can only tag received frames with the PVID of the receiving port. Selective QinQ allows adding different outer VLAN tags based on different inner VLAN tags.

The A5120 EI Switch Series supports the configuration of basic QinQ and selective QinQ at the same time on a port. When both features are enabled on the port, frames that meet the selective QinQ condition are handled with selective QinQ before the remaining frames are handled with basic QinQ.

Follow these steps to configure an outer VLAN tagging policy:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Enter interface view or port group view	Enter Ethernet or Layer-2 aggregate interface view	**interface** *interface-type interface-number*	Required Use either command
	Enter port group view	**port-group manual** *port-group-name*	
Enter QinQ view and configure the SVLAN tag for the port to add		**qinq vid** *vlan-id*	Required By default, the SVLAN tag to be added is the PVID of the receiving port.
Tag frames of the specified CVLANs with the current SVLAN		**raw-vlan-id inbound** { **all** | *vlan-list* }	Required

CAUTION:

- An inner VLAN tag corresponds to only one outer VLAN tag.
- To change an outer VLAN tag, you must delete the old outer VLAN tag configuration and configure a new outer VLAN tag.

Configuring an inner-outer VLAN 802.1p priority mapping

Through QoS policies, the A5120 EI switches achieve the following inner-outer VLAN 802.1p priority mapping modes:

- Marking the 802.1p priorities in outer VLAN tags according to the inner VLAN IDs or the 802.1p priorities in the inner VLAN tags.
- Copying the 802.1p priority in the inner VLAN tags to the outer VLAN tags.

Follow these steps to mark the 802.1p priorities in outer VLAN tags according to the inner VLAN IDs or the 802.1p priorities in the inner VLAN tags:

To do...		Use the command...	Remarks
Enter system view		**system-view**	—
Create a class and enter class view		**traffic classifier** *classifier-name* [**operator** { **and** | **or** }]	Required By default, the operator of a class is AND.
Configure a match criterion	Configure a match criterion to match the specified inner VLAN IDs	**if-match customer-vlan-id** *vlan-id-list*	Use either command.
	Configure a match criterion to match the specified inner VLAN tag priorities	**if-match customer-dot1p** *8021p-list*	
Return to system view		**quit**	—
Create a traffic behavior and enter traffic behavior view		**traffic behavior** *behavior-name*	Required
Configure the action of specifying the outer VLAN tag priorities of packets	Configure the action of marking the 802.1p priorities in outer VLAN tags	**remark dot1p** *8021p*	Use either command. Choose to configure inner-outer VLAN 802.1p priority mapping or copying as needed.
	Configure the action of copying the 802.1p priorities in the inner VLAN tags to the outer VLAN tags	**remark dot1p customer-dot1p-trust**	
Return to system view		**quit**	—
Create a QoS policy and enter QoS policy view		**qos policy** *policy-name*	Required
Associate the traffic class with the traffic behavior defined earlier		**classifier** *classifier-name* **behavior** *behavior-name*	Required
Return to system view		**quit**	—
Enter Ethernet interface view or port group	Enter Layer 2 Ethernet port view	**interface** *interface-type interface-number*	Required Use either command.

To do...		Use the command...	Remarks
view of the customer network-side port	Enter port group view	**port-group manual** *port-group-name*	• Settings made in Layer 2 Ethernet interface view take effect only on the current port. • Settings made in port group view take effect on all ports in the port group.
Enable basic QinQ		**qinq enable**	Required
Apply the QoS policy to the incoming traffic		**qos apply policy** *policy-name* **inbound**	Required

Configuring the TPID value in VLAN tags

Follow these steps to configure the TPID value:

To do...	Use the command...	Remarks
Enter system view	**system-view**	—
Configure the TPID value in the CVLAN tag	**qinq ethernet-type** *hex-value*	Optional 0x8100 by default The configuration applies to all ports.

NOTE:

- On a port with basic QinQ and customer-side QinQ not enabled, the switch judges whether a frame is VLAN tagged based on the SVLAN TPID value on the port; on a port with basic QinQ or customer-side QinQ enabled, the switch judges whether a frame is VLAN tagged based on the CVLAN TPID value globally configured.
- The TPID value configured on the A5120 EI Switch Series applies to both the CVLAN tags and the SVLAN tags.

QinQ configuration examples

Basic QinQ configuration example

Network requirements

As shown in Figure 52,

- Provider A and Provider B are edge switches on the service provider network and are connected through trunk ports. They belong to SVLAN 10 and 50.
- Customer A1, Customer A2, Customer B1, and Customer B2 are edge switches on the customer network.
- Third-party switches with a TPID value of 0x8200 are deployed between Provider A and Provider B.

Make the configuration to satisfy the following requirements:

- Frames of VLAN 200 through VLAN 299 can be exchanged between Customer A1 and Customer A2 through VLAN 10 of the service provider network.

- Frames of VLAN 250 through VLAN 350 can be exchanged between Customer B1 and Customer B2 through VLAN 50 of the service provider network.

Figure 52 Network diagram for Basic QinQ configuration

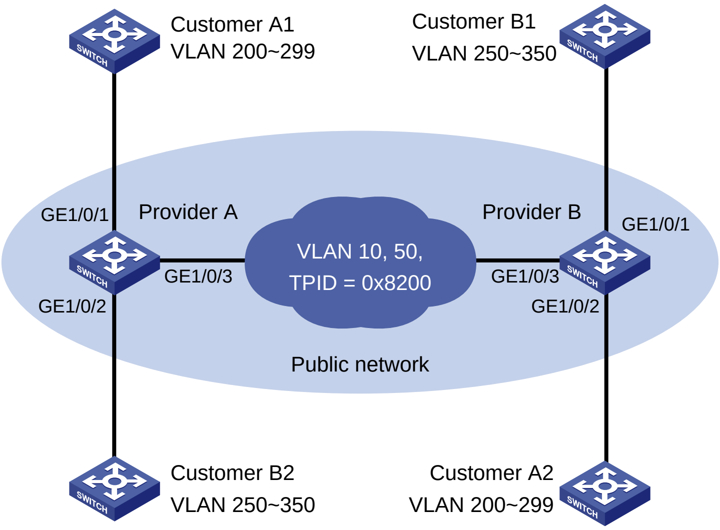

Configuration procedure

NOTE:

Be sure that you have configured the switches in the service provider network to allow QinQ packets to pass through.

1. Configure Provider A.

- Configure GigabitEthernet 1/0/1

# Configure VLAN 10 as the PVID of GigabitEthernet 1/0/1.

```
<ProviderA> system-view
[ProviderA] interface gigabitethernet 1/0/1
[ProviderA-GigabitEthernet1/0/1] port access vlan 10
```

# Enable basic QinQ on GigabitEthernet 1/0/1.

```
[ProviderA-GigabitEthernet1/0/1] qinq enable
[ProviderA-GigabitEthernet1/0/1] quit
```

- Configure GigabitEthernet 1/0/2

# Configure GigabitEthernet 1/0/2 as a hybrid port and configure VLAN 50 as the PVID of the port.

```
[ProviderA] interface gigabitethernet 1/0/2
[ProviderA-GigabitEthernet1/0/2] port link-type hybrid
[ProviderA-GigabitEthernet1/0/2] port hybrid pvid vlan 50
[ProviderA-GigabitEthernet1/0/2] port hybrid vlan 50 untagged
```

# Enable basic QinQ on GigabitEthernet 1/0/2.

```
[ProviderA-GigabitEthernet1/0/2] qinq enable
[ProviderA-GigabitEthernet1/0/2] quit
```

- Configure GigabitEthernet 1/0/3

# Configure GigabitEthernet 1/0/3 as a trunk port to permit frames of VLAN 10 and 50 to pass through.

```
[ProviderA] interface gigabitethernet 1/0/3
[ProviderA-GigabitEthernet1/0/3] port link-type trunk
[ProviderA-GigabitEthernet1/0/3] port trunk permit vlan 10 50
```

# Set the TPID value in the outer tag to 0x8200.

```
[ProviderA-GigabitEthernet1/0/3] quit
[ProviderA] qinq ethernet-type 8200
```

2. Configure Provider B.

- Configure GigabitEthernet 1/0/1

# Configure VLAN 50 as the PVID of GigabitEthernet 1/0/1.

```
<ProviderB> system-view
[ProviderB] interface gigabitethernet 1/0/1
[ProviderB-GigabitEthernet1/0/1] port access vlan 50
```

# Enable basic QinQ on GigabitEthernet 1/0/1.

```
[ProviderB-GigabitEthernet1/0/1] qinq enable
[ProviderB-GigabitEthernet1/0/1] quit
```

- Configure GigabitEthernet 1/0/2

# Configure GigabitEthernet 1/0/2 as a hybrid port and configure VLAN 10 as the PVID of the port.

```
[ProviderB] interface gigabitethernet 1/0/2
[ProviderB-GigabitEthernet1/0/2] port link-type hybrid
[ProviderB-GigabitEthernet1/0/2] port hybrid pvid vlan 10
[ProviderB-GigabitEthernet1/0/2] port hybrid vlan 10 untagged
```

# Enable basic QinQ on GigabitEthernet 1/0/2.

```
[ProviderB-GigabitEthernet1/0/2] qinq enable
[ProviderB-GigabitEthernet1/0/2] quit
```

- Configure GigabitEthernet 1/0/3

# Configure GigabitEthernet 1/0/3 as a trunk port to permit frames of VLAN 10 and 50 to pass through.

```
[ProviderB] interface gigabitethernet 1/0/3
[ProviderB-GigabitEthernet1/0/3] port link-type trunk
[ProviderB-GigabitEthernet1/0/3] port trunk permit vlan 10 50
```

# Set the TPID value in the outer tag to 0x8200.

[ProviderB-GigabitEthernet1/0/3] quit

```
[ProviderB] qinq ethernet-type 8200
```

3. Configure third-party switches.

Configure the third-party switches between Provider A and Provider B as follows: configure the port that connects GigabitEthernet 1/0/3 of Provider A and the port that connects GigabitEthernet 1/0/3 of Provider B to allow tagged frames of VLAN 10 and 50 to pass through.

Selective QinQ Configuration Example

Network requirements

As shown in Figure 53,

- Provider A and Provider B are edge devices on the service provider network and are connected through trunk ports. They belong to SVLAN 1000 and SVLAN 2000 separately.

- Customer A, Customer B and Customer C are edge devices on the customer network.
- Third-party devices with a TPID value of 0x8200 are deployed between Provider A and Provider B.

Make configuration to achieve the following:

- VLAN 10 frames of Customer A and Customer B can be forwarded to each other across SVLAN 1000;
- VLAN 20 frames of Customer A and Customer C can be forwarded to each other across SVLAN 2000.

Figure 53 Network diagram

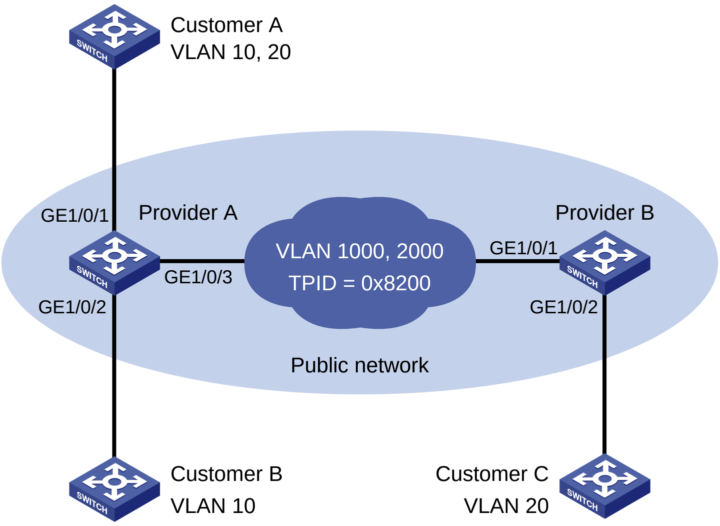

Configuration procedure

NOTE:

Be sure that you have configured the devices in the service provider network to allow QinQ packets to pass through.

1. Configure Provider A.

- Configure GigabitEthernet 1/0/1

# Configure GigabitEthernet 1/0/1 as a hybrid port to permit frames of VLAN 1000 and VLAN 2000 to pass through, and configure GigabitEthernet 1/0/1 to send packets of these VLANs with tags removed.

```
<ProviderA> system-view
[ProviderA] interface gigabitethernet 1/0/1
[ProviderA-GigabitEthernet1/0/1] port link-type hybrid
[ProviderA-GigabitEthernet1/0/1] port hybrid vlan 1000 2000 untagged
```

# Tag CVLAN 10 frames with SVLAN 1000.

```
[ProviderA-GigabitEthernet1/0/1] qinq vid 1000
[ProviderA-GigabitEthernet1/0/1-vid-1000] raw-vlan-id inbound 10
[ProviderA-GigabitEthernet1/0/1-vid-1000] quit
```

# Tag CVLAN 20 frames with SVLAN 2000.

```
[ProviderA-GigabitEthernet1/0/1] qinq vid 2000
[ProviderA-GigabitEthernet1/0/1-vid-2000] raw-vlan-id inbound 20
[ProviderA-GigabitEthernet1/0/1-vid-2000] quit
```

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

- Configure GigabitEthernet 1/0/2

# Configure GigabitEthernet 1/0/2 as a hybrid port to permit frames of VLAN 1000 to pass through, and configure GigabitEthernet 1/0/2 to send packets of VLAN 1000 with tag removed.

```
[ProviderA] interface gigabitethernet 1/0/2
[ProviderA-GigabitEthernet1/0/2] port link-type hybrid
[ProviderA-GigabitEthernet1/0/2] port hybrid vlan 1000 untagged
```

# Tag CVLAN 10 frames with SVLAN 1000.

```
[ProviderA-GigabitEthernet1/0/2] qinq vid 1000
[ProviderA-GigabitEthernet1/0/2-vid-1000] raw-vlan-id inbound 10
[ProviderA-GigabitEthernet1/0/2-vid-1000] quit
[ProviderA-GigabitEthernet1/0/2] quit
```

- Configure GigabitEthernet 1/0/3

# Configure GigabitEthernet 1/0/3 as a trunk port to permit frames of VLAN 1000 and VLAN 2000 to pass through.

```
[ProviderA] interface gigabitethernet 1/0/3
[ProviderA-GigabitEthernet1/0/3] port link-type trunk
[Sysname-GigabitEthernet1/0/3] port trunk permit vlan 1000 2000
```

# Set the TPID value in the outer tag to 0x8200.

```
[ProviderA-GigabitEthernet1/0/3] quit
[ProviderA] qinq ethernet-type 8200
```

2. Configure Provider B.

- Configure GigabitEthernet 1/0/1

# Configure GigabitEthernet 1/0/1 as a trunk port to permit frames of VLAN 1000 and VLAN 2000 to pass through.

```
<ProviderB> system-view
[ProviderB] interface gigabitethernet 1/0/1
[ProviderB-GigabitEthernet1/0/1] port link-type trunk
[ProviderB-GigabitEthernet1/0/1] port trunk permit vlan 1000 2000
```

- Configure GigabitEthernet 1/0/2

# Configure GigabitEthernet 1/0/2 as a hybrid port to permit frames of VLAN 2000 to pass through, and configure GigabitEthernet 1/0/2 to send packets of VLAN 2000 with tag removed.

```
[ProviderB] interface gigabitethernet 1/0/2
[ProviderB-GigabitEthernet1/0/2] port link-type hybrid
[ProviderB-GigabitEthernet1/0/2] port hybrid vlan 2000 untagged
```

# Tag CVLAN 20 frames with SVLAN 2000.

```
[ProviderB-GigabitEthernet1/0/2] qinq vid 2000
[ProviderB-GigabitEthernet1/0/2-vid-2000] raw-vlan-id inbound 20
```

# Set the TPID value in the outer tag to 0x8200.

```
[ProviderA-GigabitEthernet1/0/3] quit
[ProviderA] qinq ethernet-type 8200
```

3. Configure third-party devices.

Configure the third-party devices between Provider A and Provider B as follows: configure the port that connects GigabitEthernet 1/0/3 of Provider A and the port that connects GigabitEthernet 1/0/1 of Provider B to allow tagged frames of VLAN 1000 and VLAN 2000 to pass through.

LLDP configuration

Overview

Background

In a heterogeneous network, a standard configuration exchange platform ensures that different types of network devices from different vendors can discover one another and exchange configuration for the sake of interoperability and management.

The IETF drafted the Link Layer Discovery Protocol (LLDP) in IEEE 802.1AB. The protocol operates on the data link layer to exchange device information between directly connected devices. With LLDP, a device sends local device information (including its major functions, management IP address, device ID, and port ID) as TLV (type, length, and value) triplets in Link Layer Discovery Protocol Data Units (LLDPDUs) to the directly connected devices. At the same time, the device stores the device information received in LLDPDUs sent from the LLDP neighbors in a standard management information base (MIB). LLDP enables a network management system to quickly and identify Layer-2 network topology change.

NOTE:

For more information about MIBs, see the *Network Management and Monitoring Configuration Guide*.

Basic concepts

LLDPDU formats

LLDP sends device information in LLDP data units (LLDPDUs). LLDPDUs are encapsulated in Ethernet II or Subnetwork Access Protocol (SNAP) frames.

1. Ethernet II-encapsulated LLDPDU format

Figure 54 Ethernet II-encapsulated LLDPDU format

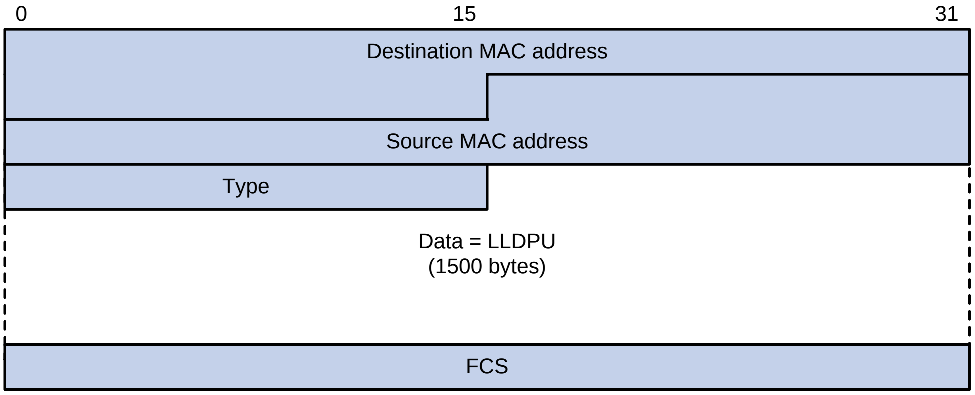

Table 21 Description of the fields in an Ethernet II-encapsulated LLDPDU

Field	Description
Destination MAC address	The MAC address to which the LLDPDU is advertised. It is fixed to 0x0180-C200-000E, a multicast MAC address.
Source MAC address	The MAC address of the sending port. If the port does not have a MAC address, the MAC address of the sending bridge is used.
Type	The Ethernet type for the upper layer protocol. It is 0x88CC for LLDP.
Data	LLDP data unit (LLDPDU)
FCS	Frame check sequence, a 32-bit CRC value used to determine the validity of the received Ethernet frame

2. SNAP-encapsulated LLDPDU format

Figure 55 SNAP-encapsulated LLDPDU format

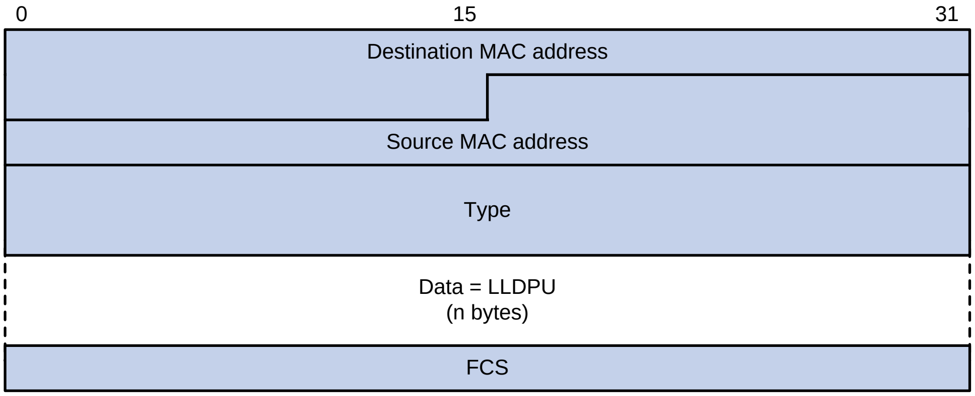

Table 22 Description of the fields in a SNAP-encapsulated LLDPDU

Field	Description
Destination MAC address	The MAC address to which the LLDPDU is advertised. It is fixed to 0x0180-C200-000E, a multicast MAC address.
Source MAC address	The MAC address of the sending port. If the port does not have a MAC address, the MAC address of the sending bridge is used.
Type	The SNAP type for the upper layer protocol. It is 0xAAAA-0300-0000-88CC for LLDP.
Data	LLDPDU
FCS	Frame check sequence, a 32-bit CRC value used to determine the validity of the received Ethernet frame

LLDPDUs

LLDP uses LLDPDUs to exchange information. An LLDPDU comprises multiple type, length, and value (TLV) sequences. Each TLV carries a type of device information, as shown in Figure 56.

Figure 56 LLDPDU encapsulation format

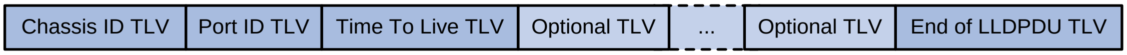

An LLDPDU can carry up to 28 types of TLVs. Mandatory TLVs include Chassis ID TLV, Port ID TLV, Time To Live TLV, and End of LLDPDU TLV. Other TLVs are optional.

TLVs

TLVs are type, length, and value sequences that carry information elements. The type field identifies the type of information, the length field measures the length of the information field in octets, and the value field contains the information itself.

LLDPDU TLVs fall into the categories of basic management TLVs, organizationally (IEEE 802.1 and IEEE 802.3) specific TLVs, and LLDP-MED (media endpoint discovery) TLVs. Basic management TLVs are essential to device management. Organizationally specific TLVs and LLDP-MED TLVs are used for enhanced device management; they are defined by standardization or other organizations and are optional to LLDPDUs.

1. Basic management TLVs

Table 21 lists the basic management TLV types. Some of them are mandatory to LLDPDUs, that is, must be included in every LLDPDU.

Table 23 Basic management TLVs

Type	Description	Remarks
Chassis ID	Specifies the bridge MAC address of the sending device	Mandatory
Port ID	Specifies the ID of the sending port If the LLDPDU carries LLDP-MED TLVs, the port ID TLV carries the MAC address of the sending port or the bridge MAC if the port does not have a MAC address. If the LLDPDU carries no LLDP-MED TLVs, the port ID TLV carries the port name.	
Time To Live	Specifies the life of the transmitted information on the receiving device	
End of LLDPDU	Marks the end of the TLV sequence in the LLDPDU	
Port Description	Specifies the port description of the sending port	Optional
System Name	Specifies the assigned name of the sending device	
System Description	Specifies the description of the sending device	
System Capabilities	Identifies the primary functions of the sending device and the enabled primary functions	
Management Address	Specifies the management address, and the interface number and OID (object identifier) associated with the address	

2. IEEE 802.1 organizationally specific TLVs

Table 24 IEEE 802.1 organizationally specific TLVs

Type	Description
Port VLAN ID	Specifies the port's VLAN identifier (PVID). An LLDPDU carries only one TLV of this type.

Type	Description
Port And Protocol VLAN ID	Indicates whether the device supports protocol VLANs and, if so, what VLAN IDs these protocols will be associated with. An LLDPDU can carry multiple different TLVs of this type.
VLAN Name	Specifies the textual name of any VLAN to which the port belongs. An LLDPDU can carry multiple different TLVs of this type.
Protocol Identity	Indicates protocols supported on the port. An LLDPDU can carry multiple different TLVs of this type.

NOTE:

HP A5120 EI Switch Series supports only receiving protocol identity TLVs.

3. IEEE 802.3 organizationally specific TLVs

Table 25 IEEE 802.3 organizationally specific TLVs

Type	Description
MAC/PHY Configuration/Status	Contains the bit-rate and duplex capabilities of the sending port, support for auto negotiation, enabling status of auto negotiation, and the current rate and duplex mode.
Power Via MDI	Contains the power supply capability of the port, including the Power over Ethernet (PoE) type, which can be Power Sourcing Equipment (PSE) or Powered Device (PD), PoE mode, whether PSE power supply is supported, whether PSE power supply is enabled, and whether the PoE mode is controllable.
Link Aggregation	Indicates the aggregation capability of the port (whether the link is capable of being aggregated), and the aggregation status (whether the link is in an aggregation).
Maximum Frame Size	Indicates the supported maximum frame size. It is now the MTU of the port.
Power Stateful Control	Indicates the power state control configured on the sending port, including the power type of the PSE or PD, PoE sourcing and receiving priority, and PoE sourcing and receiving power.

NOTE:

The Power Stateful Control TLV is defined in IEEE P802.3at D1.0. The later versions no longer support this TLV. HP devices send this type of TLVs only after receiving them.

LLDP-MED TLVs

LLDP-MED TLVs provide multiple advanced applications for voice over IP (VoIP), such as basic configuration, network policy configuration, and address and directory management. LLDP-MED TLVs provide a cost-effective and easy-to-use solution for deploying voice devices in Ethernet. LLDP-MED TLVs are shown in Table 26:

Table 26 LLDP-MED TLVs

Type	Description
LLDP-MED Capabilities	Allows a network device to advertise the LLDP-MED TLVs that it supports
Network Policy	Allows a network device or terminal device to advertise the VLAN ID of the specific port, the VLAN type, and the Layer 2 priorities for specific applications
Extended Power-via-MDI	Allows a network device or terminal device to advertise power supply capability. This TLV is an extension of the Power Via MDI TLV.
Hardware Revision	Allows a terminal device to advertise its hardware version
Firmware Revision	Allows a terminal device to advertise its firmware version
Software Revision	Allows a terminal device to advertise its software version
Serial Number	Allows a terminal device to advertise its serial number
Manufacturer Name	Allows a terminal device to advertise its vendor name
Model Name	Allows a terminal device to advertise its model name
Asset ID	Allows a terminal device to advertise its asset ID. The typical case is that the user specifies the asset ID for the endpoint to facilitate directory management and asset tracking.
Location Identification	Allows a network device to advertise the appropriate location identifier information for a terminal device to use in the context of location-based applications

Management address

The network management system uses the management address of a device to identify and manage the device for topology maintenance and network management. The management address TLV encapsulates the management address.

How LLDP works

Operating modes of LLDP

LLDP can operate in one of the following modes:

- TxRx mode. A port in this mode sends and receives LLDPDUs.
- Tx mode. A port in this mode only sends LLDPDUs.
- Rx mode. A port in this mode only receives LLDPDUs.
- Disable mode. A port in this mode does not send or receive LLDPDUs.

Each time the LLDP operating mode of a port changes, its LLDP protocol state machine re-initializes. To prevent LLDP from being initialized too frequently at times of frequent changes to the operating mode, you can configure a re-initialization delay. With this delay configured, before a port can initialize LLDP, it must wait for the specified interval after the LLDP operating mode changes.

Transmitting LLDPDUs

An LLDP-enabled port operating in TxRx mode or Tx mode sends LLDPDUs to its directly connected devices both periodically and when the local configuration changes. To prevent LLDPDUs from overwhelming the network during times of frequent changes to local device information, an interval is introduced between two successive LLDPDUs.

This interval is shortened to 1 second in either of the following cases:

- A new neighbor is discovered, in other words, a new LLDPDU is received and carries device information new to the local device.
- The LLDP operating mode of the port changes from Disable or Rx to TxRx or Tx.

This is the fast sending mechanism of LLDP. With this mechanism, a specific number of LLDPDUs are sent successively at 1-second intervals, to help LLDP neighbors discover the local device as soon as possible. Then, the normal LLDPDU transmit interval resumes.

Receiving LLDPDUs

An LLDP-enabled port that is operating in TxRx mode or Rx mode checks the validity of TLVs carried in every received LLDPDU. If valid, the information is saved and an aging timer is set for it based on the time to live (TTL) value in the Time To Live TLV carried in the LLDPDU. If the TTL value is zero, the information ages out immediately.

Protocols and standards

- IEEE 802.1AB-2005, *Station and Media Access Control Connectivity Discovery*
- ANSI/TIA-1057, *Link Layer Discovery Protocol for Media Endpoint Devices*

LLDP configuration task list

Complete these tasks to configure LLDP:

Task		Remarks
Performing basic LLDP configuration	Enabling LLDP	Required
	Setting the LLDP operating mode	Optional
	Setting the LLDP re-initialization delay	Optional
	Enabling LLDP polling	Optional
	Configuring the advertisable TLVs	Optional
	Configuring the management address and its encoding format	Optional
	Setting other LLDP parameters	Optional
	Setting an encapsulation format for LLDPDUs	Optional
Configuring CDP compatibility		Optional
Configuring LLDP trapping		Optional

NOTE:

LLDP-related configurations made in Ethernet interface view take effect only on the current port, and those made in port group view take effect on all ports in the current port group.

Performing basic LLDP configuration

Enabling LLDP

To make LLDP take effect on certain ports, you must enable LLDP both globally and on these ports.

Follow these steps to enable LLDP:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enable LLDP globally		**lldp enable**	Required By default, LLDP is globally enabled.
Enter Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Enable LLDP		**lldp enable**	Optional By default, LLDP is enabled on a port.

Setting the LLDP operating mode

LLDP can operate in one of the following modes.

- TxRx mode. A port in this mode sends and receives LLDPDUs.
- Tx mode. A port in this mode only sends LLDPDUs.
- Rx mode. A port in this mode only receives LLDPDUs.
- Disable mode. A port in this mode does not send or receive LLDPDUs.

Follow these steps to set the LLDP operating mode:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Set the LLDP operating mode		**lldp admin-status** { **disable** | **rx** | **tx** | **txrx** }	Optional TxRx by default

Setting the LLDP re-initialization delay

When LLDP operating mode changes on a port, the port initializes the protocol state machines after a certain delay. By adjusting the LLDP re-initialization delay, you can avoid frequent initializations caused by frequent changes to the LLDP operating mode on a port.

Follow these steps to set the LLDP re-initialization delay for ports:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Set the LLDP re-initialization delay	**lldp timer reinit-delay** *delay*	Optional 2 seconds by default

Enabling LLDP polling

With LLDP polling enabled, a device searches for local configuration changes periodically. Upon detecting a configuration change, the device sends LLDPDUs to inform the neighboring devices of the change.

Follow these steps to enable LLDP polling:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Enable LLDP polling and set the polling interval		**lldp check-change-interval** *interval*	Required Disabled by default

Configuring the advertisable TLVs

Follow these steps to configure the advertisable LLDPDU TLVs on the specified port or ports:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	

To do…	Use the command…	Remarks
Configure the advertisable TLVs	**lldp tlv-enable** { **basic-tlv** { **all** | **port-description** | **system-capability** | **system-description** | **system-name** } | **dot1-tlv** { **all** | **port-vlan-id** | **protocol-vlan-id** [*vlan-id*] | **vlan-name** [*vlan-id*] } | **dot3-tlv** { **all** | **link-aggregation** | **mac-physic** | **max-frame-size** | **power** } | **med-tlv** { **all** | **capability** | **inventory** | **location-id** { **civic-address** *device-type country-code* { *ca-type ca-value* }&<1–10> | **elin-address** *tel-number* } | **network-policy** | **power-over-ethernet** } }	Optional By default, all types of LLDP TLVs except the location identification TLVs are advertisable on a Layer 2 Ethernet port.

Configuring the management address and its encoding format

LLDP encodes management addresses in numeric or character string format in management address TLVs.

By default, management addresses are encoded in numeric format. If a neighbor encoded its management address in character string format, you must configure the encoding format of the management address as string on the connecting port to guarantee normal communication with the neighbor.

Follow these steps to configure a management address to be advertised and its encoding format on one or a group of ports:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Allow LLDP to advertise the management address in LLDPDUs and configure the advertised management address		**lldp management-address-tlv** [*ip-address*]	Optional By default, the management address is sent through LLDPDUs. For a Layer 2 Ethernet port, the management address is the main IP address of the lowest-ID VLAN carried on the port. If none of the carried VLANs is assigned an IP address, no management address will be advertised.
Configure the encoding format of the management address as character string		**lldp management-address-format string**	Optional By default, the management address is encapsulated in the numeric format.

Setting other LLDP parameters

The Time To Live TLV carried in an LLDPDU determines how long the device information carried in the LLDPDU can be saved on a recipient device.

By setting the TTL multiplier, you can configure the TTL of locally sent LLDPDUs, which determines how long information about the local device can be saved on a neighbor device. The TTL is expressed using the following formula:

TTL = Min (65535, (TTL multiplier × LLDPDU transmit interval))

As the expression shows, the TTL can be up to 65535 seconds. TTLs greater than 65535 will be rounded down to 65535 seconds.

Follow these steps to change the TTL multiplier:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—
Set the TTL multiplier	**lldp hold-multiplier** *value*	Optional 4 by default.
Set the LLDPDU transmit interval	**lldp timer tx-interval** *interval*	Optional 30 seconds by default
Set LLDPDU transmit delay	**lldp timer tx-delay** *delay*	Optional 2 seconds by default
Set the number of LLDPDUs sent each time fast LLDPDU transmission is triggered	**lldp fast-count** *count*	Optional 3 by default

NOTE:

To ensure that the LLDP neighbors can receive LLDPDUs to update information about the current device before it ages out, configure both the LLDPDU transmit interval and delay to be less than the TTL.

Setting an encapsulation format for LLDPDUs

LLDPDUs can be encapsulated in the following formats: Ethernet II or SNAP frames.

- With Ethernet II encapsulation configured, an LLDP port sends LLDPDUs in Ethernet II frames and processes only incoming, Ethernet II encapsulated LLDPDUs.
- With SNAP encapsulation configured, an LLDP port sends LLDPDUs in SNAP frames and processes only incoming, SNAP encapsulated LLDPDUs.

By default, Ethernet II frames encapsulate LLDPDUs. If the neighbor devices encapsulate LLDPDUs in SNAP frames, configure the encapsulation format for LLDPDUs as SNAP to guarantee normal communication with the neighbors.

Follow these steps to set the encapsulation format for LLDPDUs to SNAP:

To do…	Use the command…	Remarks
Enter system view	**system-view**	—

To do…		Use the command…	Remarks
Enter Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Set the encapsulation format for LLDPDUs to SNAP		**lldp encapsulation snap**	Required Ethernet II encapsulation format applies by default.

NOTE:

LLDP-CDP (Cisco Discovery Protocol) packets use only SNAP encapsulation.

Configuring CDP compatibility

NOTE:

For more information about voice VLAN, see the chapter "Voice VLAN configuration."

To make your device work with Cisco IP phones, you must enable CDP compatibility.

If your LLDP-enabled device cannot recognize CDP packets, it does not respond to the requests of Cisco IP phones for the voice VLAN ID configured on the device. As a result, a requesting Cisco IP phone sends voice traffic without any tag to your device, and, as a result, your device cannot differentiate the voice traffic from other types of traffic.

With CDP compatibility enabled, your device can receive and recognize CDP packets from a Cisco IP phone and respond with CDP packets, which carry the voice VLAN configuration TLVs. According to the voice VLAN configuration TLVs, the IP phone automatically configures the voice VLAN. As a result, the voice traffic is confined in the configured voice VLAN, and differentiated from other types of traffic.

Configuration prerequisites

Before you configure CDP compatibility, complete the following tasks:

- Globally enable LLDP.
- Enable LLDP on the port connecting to an IP phone and configure the port to operate in TxRx mode.

Configuring CDP compatibility

CDP-compatible LLDP operates in one of the following modes:

- TxRx: CDP packets can be transmitted and received.
- Disable: CDP packets can be neither transmitted nor received.

LLDP traps are sent periodically, and the interval is configurable. To make CDP-compatible LLDP take effect on certain ports, first enable CDP-compatible LLDP globally, and then configure CDP-compatible LLDP to operate in TxRx mode.

Follow these steps to enable LLDP to be compatible with CDP:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enable CDP compatibility globally		**lldp compliance cdp**	Required Disabled by default.
Enter Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Configure CDP-compatible LLDP to operate in TxRx mode		**lldp compliance admin-status cdp txrx**	Required Disable mode by default

△ CAUTION:

The maximum TTL value that CDP allows is 255 seconds. To make CDP-compatible LLDP work properly with Cisco IP phones, be sure that the product of the TTL multiplier and the LLDPDU transmit interval is less than 255 seconds.

Configuring LLDP trapping

LLDP trapping notifies the network management system (NMS) of events such as newly-detected neighboring devices and link malfunctions.

To prevent excessive LLDP traps from being sent when the topology is unstable, you can set a minimum trap sending interval for LLDP.

Follow these steps to configure LLDP trapping:

To do…		Use the command…	Remarks
Enter system view		**system-view**	—
Enter Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view	**interface** *interface-type interface-number*	Required Use either command.
	Enter port group view	**port-group manual** *port-group-name*	
Enable LLDP trapping		**lldp notification remote-change enable**	Required Disabled by default
Quit to system view		**quit**	—
Set the interval to send LLDP traps		**lldp timer notification-interval** *interval*	Optional 5 seconds by default

Displaying and maintaining LLDP

To do…	Use the command…	Remarks
Display the global LLDP information or the information contained in the LLDP TLVs to be sent through a port	**display lldp local-information** [**global** | **interface** *interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the information contained in the LLDP TLVs sent from neighboring devices	**display lldp neighbor-information** [**brief** | **interface** *interface-type interface-number* [**brief**] | **list** [**system-name** *system-name*]] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display LLDP statistics	**display lldp statistics** [**global** | **interface** *interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display LLDP status of a port	**display lldp status** [**interface** *interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display types of advertisable optional LLDP TLVs	**display lldp tlv-config** [**interface** *interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

LLDP configuration examples

Basic LLDP configuration example

Network requirements

As shown in Figure 57, the NMS and Switch A are located in the same Ethernet. An MED device and Switch B are connected to GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 of Switch A.

Enable LLDP on the ports of Switch A and Switch B to monitor the link between Switch A and Switch B and the link between Switch A and the MED device on the NMS.

Figure 57 Network diagram for basic LLDP configuration

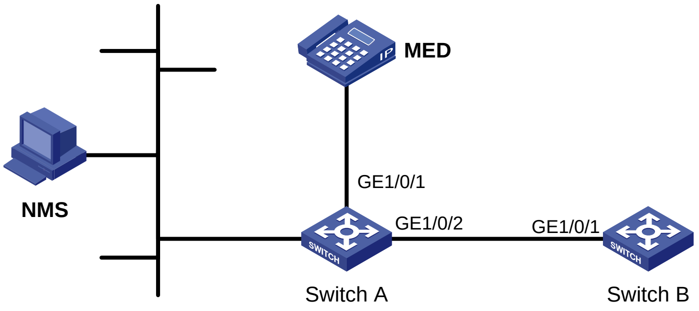

Configuration procedure

1. Configure Switch A

# Enable LLDP globally. (You can skip this step because LLDP is enabled globally by default.)

```
<SwitchA> system-view
[SwitchA] lldp enable
```

# Enable LLDP on GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2. (You can skip this step because LLDP is enabled on ports by default.) Set the LLDP operating mode to Rx.

```
 [SwitchA] interface gigabitethernet 1/0/1
[SwitchA-GigabitEthernet1/0/1] lldp enable
[SwitchA-GigabitEthernet1/0/1] lldp admin-status rx
[SwitchA-GigabitEthernet1/0/1] quit
[SwitchA] interface gigabitethernet 1/0/2
[SwitchA-GigabitEthernet1/0/2] lldp enable
[SwitchA-GigabitEthernet1/0/2] lldp admin-status rx
[SwitchA-GigabitEthernet1/0/2] quit
```

2. Configure Switch B.

# Enable LLDP globally. (You can skip this step because LLDP is enabled globally by default.)

```
<SwitchB> system-view
[SwitchB] lldp enable
```

# Enable LLDP on GigabitEthernet1/0/1. (You can skip this step because LLDP is enabled on ports by default.) Set the LLDP operating mode to Tx.

```
 [SwitchB] interface gigabitethernet 1/0/1
[SwitchB-GigabitEthernet1/0/1] lldp enable
[SwitchB-GigabitEthernet1/0/1] lldp admin-status tx
[SwitchB-GigabitEthernet1/0/1] quit
```

3. Verify the configuration

# Display the global LLDP status and port LLDP status on Switch A.

```
[SwitchA] display lldp status
Global status of LLDP: Enable
The current number of LLDP neighbors: 2
The current number of CDP neighbors: 0
LLDP neighbor information last changed time: 0 days,0 hours,4 minutes,40 seconds
Transmit interval              : 30s
Hold multiplier                : 4
Reinit delay                   : 2s
Transmit delay                 : 2s
Trap interval                  : 5s
Fast start times               : 3

Port 1 [GigabitEthernet1/0/1]:
Port status of LLDP            : Enable
Admin status                   : Rx_Only
Trap flag                      : No
Polling interval               : 0s

Number of neighbors:             1
Number of MED neighbors        : 1
Number of CDP neighbors        : 0
Number of sent optional TLV    : 0
Number of received unknown TLV : 0

Port 2 [GigabitEthernet1/0/2]:
Port status of LLDP            : Enable
Admin status                   : Rx_Only
Trap flag                      : No
Polling interval               : 0s

Number of neighbors:             1
Number of MED neighbors        : 0
Number of CDP neighbors        : 0
Number of sent optional TLV    : 0
Number of received unknown TLV : 3
```

As the sample output shows, GigabitEthernet 1/0/1 of Switch A connects to an MED device, and GigabitEthernet 1/0/2 of Switch A connects to a non-MED device. Both ports operate in Rx mode, and they only receive LLDPDUs.

# Remove the link between Switch A and Switch B and then display the global LLDP status and port LLDP status on Switch A.

```
[SwitchA] display lldp status
Global status of LLDP: Enable
The current number of LLDP neighbors: 1
The current number of CDP neighbors: 0
LLDP neighbor information last changed time: 0 days,0 hours,5 minutes,20 seconds
Transmit interval              : 30s
Hold multiplier                : 4
Reinit delay                   : 2s
Transmit delay                 : 2s
Trap interval                  : 5s
Fast start times               : 3

Port 1 [GigabitEthernet1/0/1]:
Port status of LLDP            : Enable
Admin status                   : Rx_Only
Trap flag                      : No
Polling interval               : 0s

Number of neighbors             : 1
Number of MED neighbors         : 1
Number of CDP neighbors         : 0
Number of sent optional TLV     : 0
Number of received unknown TLV  : 5

Port 2 [GigabitEthernet1/0/2]:
Port status of LLDP            : Enable
Admin status                   : Rx_Only
Trap flag                      : No
Polling interval               : 0s
```

```
Number of neighbors              : 0
Number of MED neighbors          : 0
Number of CDP neighbors          : 0
Number of sent optional TLV      : 0
Number of received unknown TLV   : 0
```

As the sample output shows, GigabitEthernet 1/0/2 of Switch A does not connect to any neighboring devices.

CDP-compatible LLDP configuration example

Network requirements

As shown in Figure 58:

- GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 of Switch A are each connected to a Cisco IP phone.
- Configure voice VLAN 2 on Switch A. Enable CDP compatibility of LLDP on Switch A to allow the Cisco IP phones to automatically configure the voice VLAN, confining their voice traffic within the voice VLAN and isolating the voice traffic from other types of traffic.

Figure 58 Network diagram for CDP-compatible LLDP configuration

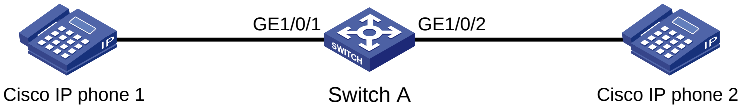

Configuration procedure

1. Configure a voice VLAN on Switch A

# Create VLAN 2.

```
<SwitchA> system-view
[SwitchA] vlan 2
[SwitchA-vlan2] quit
```

# Set the link type of GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to trunk and enable voice VLAN on them.

```
[SwitchA] interface gigabitethernet 1/0/1
[SwitchA-GigabitEthernet1/0/1] port link-type trunk
[SwitchA-GigabitEthernet1/0/1] voice vlan 2 enable
[SwitchA-GigabitEthernet1/0/1] quit
[SwitchA] interface gigabitethernet 1/0/2
[SwitchA-GigabitEthernet1/0/2] port link-type trunk
[SwitchA-GigabitEthernet1/0/2] voice vlan 2 enable
[SwitchA-GigabitEthernet1/0/2] quit
```

2. Configure CDP-compatible LLDP on Switch A

# Enable LLDP globally and enable LLDP to be compatible with CDP globally.

```
[SwitchA] lldp enable
[SwitchA] lldp compliance cdp
```

# Enable LLDP. (You can skip this step because LLDP is enabled on ports by default.) Configure LLDP to operate in TxRx mode, and configure CDP-compatible LLDP to operate in TxRx mode on GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2.

```
[SwitchA] interface gigabitethernet 1/0/1
[SwitchA-GigabitEthernet1/0/1] lldp enable
[SwitchA-GigabitEthernet1/0/1] lldp admin-status txrx
[SwitchA-GigabitEthernet1/0/1] lldp compliance admin-status cdp txrx
[SwitchA-GigabitEthernet1/0/1] quit
[SwitchA] interface gigabitethernet 1/0/2
[SwitchA-GigabitEthernet1/0/2] lldp enable
[SwitchA-GigabitEthernet1/0/2] lldp admin-status txrx
[SwitchA-GigabitEthernet1/0/2] lldp compliance admin-status cdp txrx
[SwitchA-GigabitEthernet1/0/2] quit
```

3. Verify the configuration

# Display the neighbor information on Switch A.

```
[SwitchA] display lldp neighbor-information

CDP neighbor-information of port 1[GigabitEthernet1/0/1]:
  CDP neighbor index : 1
  Chassis ID         : SEP00141CBCDBFE
  Port ID            : Port 1
  Sofrware version   : P0030301MFG2
  Platform           : Cisco IP Phone 7960
  Duplex             : Full

CDP neighbor-information of port 2[GigabitEthernet1/0/2]:
  CDP neighbor index : 2
  Chassis ID         : SEP00141CBCDBFF
  Port ID            : Port 1
  Sofrware version   : P0030301MFG2
  Platform           : Cisco IP Phone 7960
  Duplex             : Full
```

As the sample output shows, Switch A has discovered the IP phones connected to GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2, and has obtained their LLDP device information.

Support and other resources

Contacting HP

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

http://www.hp.com/support

Before contacting HP, collect the following information:

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

Subscription service

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

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

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

Related information

Documents

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

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

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

Websites

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

Conventions

This section describes the conventions used in this documentation set.

Command conventions

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

GUI conventions

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

Symbols

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

Network topology icons

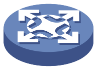	Represents a generic network device, such as a router, switch, or firewall.
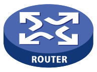	Represents a routing-capable device, such as a router or Layer 3 switch.
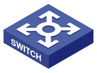	Represents a generic switch, such as a Layer 2 or Layer 3 switch, or a router that supports Layer 2 forwarding and other Layer 2 features.

Port numbering in examples

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

Index